1

The Technology

1

Blockchains, Bitcoin, and Decentralized Computing Platforms

At their core, blockchains are decentralized databases, maintained by a distributed network of computers. They blend together a variety of different technologies—including peer-to-peer networks, public-private key cryptography, and consensus mechanisms—to create a novel type of database. We provide here a short description of how blockchains work, and unpack and contextualize their key technological components.

UNTIL THE BIRTH of the Internet, computers suffered in isolation. They were islands, lacking a way to connect to one another except by using cumbersome cables. That all changed in the late 1950s. With the Soviets successfully launching *Sputnik* into space, and with fears of the Cold War mounting, researchers at the Rand Corporation began to explore a new computing paradigm—in hopes of developing a system that would be able to withstand a nuclear catastrophe.[1] In August 1964, after years of research, Paul Baran, one of the Rand researchers, reported a breakthrough. By relying on a technology called packet switching, Baran was able to send fragments of information from one computer to another and have these fragments reassembled, almost like magic.[2]

Armed with Baran's research, the Advanced Research Projects Agency (ARPA) at the U.S. Department of Defense used this new technology to create the first network of computers, ARPAnet, later renamed DARPAnet after “Defense” was added to the beginning of the agency’s name, helping researchers and academics to share files and exchange resources with one another. Over the course of the next several decades, the power of this new network grew, as additional layers of technology—such as TCP/IP (the Transmission Control Program and Internet Protocol) and domain name services (DNSs)—were developed to make it easier to identify computers on the network and ensure that information was being appropriately routed. Computers were no longer isolated.[3] They were now being stitched together by using thin layers of code.

Public-Private Key Encryption and Digital Signatures

As DARPAnet was getting off the ground, a second revolution was brewing. New cryptographic algorithms were creating new means for individuals and machines to swap messages, files, and other information in a secure and authenticated way. In 1976, Whitfield Diffie and Marty Hellman, two cryptographers from Stanford University, ingeniously invented the concept of “public-private key cryptography,” solving one of cryptography’s fundamental problems—the need for secure key distribution—while at the same time laying out a theoretical foundation for authenticated digital signatures.[4]

Before the advent of public-private key encryption, sending private messages was difficult. Encrypted messages traveled over insecure channels, making them vulnerable to interception. To send an encrypted message, the message would need to be scrambled by using a “key” (also known as a cipher), resulting in an impenetrable string of text. When the scrambled message arrived at its intended destination, the recipient would use the same key to decode the encrypted text, revealing the underlying message.[5]

One significant limitation of these early cryptographic systems was that the key was central to maintaining the confidentiality of any message sent. Parties using these systems had to agree on a key before exchanging messages, or the key somehow had to be communicated to the receiving party. Because of these limitations, keys could easily be compromised. If a third party gained access to a key, they could intercept a communication and decode an encrypted message.[6]

Public-private key cryptography solved this problem by enabling the sending of encrypted messages without the need for a shared key. Under Diffie and Hellman’s model, both parties would agree on a shared pubic key and each party would generate a unique private key.[7] The private key acted as a secret password, which parties did not need to share, whereas the public key served as a reference point that could be freely communicated. By combining the public key with one party's private key, and then combining the outcome with the private key of the other party, Diffie and Hellman realized that it was possible to generate a shared secret key that could be used to both encrypt and decrypt messages.[8]

In 1978, shortly after Diffie and Hellman publicly released their groundbreaking work, a team of cryptographers from MIT—Ron Rivest, Adi Shamir, and Len Adleman—built on Diffie and Hellman's research. They developed an algorithm, known as the RSA algorithm (after the last initials of the developers), in order to create a mathematically linked set of public and private keys generated by multiplying together two large prime numbers. These cryptographers figured out that it was relatively straightforward to multiply two large prime numbers together but exceptionally difficult—even for powerful computers—to calculate which prime numbers were used (a process called prime factorization).[9]

By taking advantage of this mathematical peculiarity, the RSA algorithm made it possible for people to broadcast their public keys widely, knowing that it would be nearly impossible to uncover the underlying private keys.[10] For example, if Alice wanted to send sensitive information to Bob, she could encrypt the information using her own public key and Bob's public key and publicly publish the encrypted message. With the RSA algorithm, and because of the use of prime factorization, only Bob's private key would be able to decrypt the message.

The application of public-private key cryptography extended beyond just encrypting messages. As Diffie and Hellman recognized, by building new cryptosystems where "enciphering and deciphering were governed by distinct keys," public-private key cryptography could underpin secure and authenticated digital signatures that were highly resistant to forgery—thus replacing the need for written signatures that "require paper instruments and contracts."[11]

For instance, by using the RSA algorithm, a sending party could attach to a message a "digital signature" generated by combining the message with the sending party's private key.[12] Once sent, the receiving party could use the sending party's public key to check the authenticity and integrity of the message. By using public-private key encryption and digital signatures, if Alice wanted to send a private message to Bob, she could encrypt the message by using her own private key and Bob's public key and then sign the message by using her private key. Bob could then use Alice's public key to verify that the message originated from Alice and had not been altered during transmission. Bob could then safely decrypt the message by using his private key and Alice's public key.[13]

Public-private key encryption sparked the imagination of a new generation of academics, mathematicians, and computer scientists, who began to envision new systems that could be constructed using these new cryptographic techniques. By relying on public-private key cryptography and digital signatures, it became theoretically possible to build electronic cash, pseudonymous reputation, and content distribution systems, as well as new forms of digital contracts.[14]

The Commercial Internet and Peer-to-Peer Networks

In the years following the birth of the Internet and the invention of public-private key cryptography, the computing revolution spread. With the cost of computers rapidly decreasing, these once esoteric machines graduated from the basements of large corporations and government agencies onto our desks and into our homes. After Apple released its iconic personal computer, the Apple II, a wide range of low-cost computers flooded the market. Seemingly overnight, computers seeped into our daily lives.

By the mid-1990s, the Internet had entered a phase of rapid expansion and commercialization. DARPAnet had grown beyond its initial academic setting and, with some updates, was transformed into the modern Internet. Fueled by a constellation of private Internet service providers (ISPs), millions of people across the globe were exploring the contours of "cyberspace," interacting with new software protocols that enabled people to send electronic messages (via the simple mail transfer protocol, SMTP), transfer files (via the file transfer protocol, FTP), and distribute and link to media hosted on one another's computers (via the hypertext transfer protocol, HTTP). In a matter of years, the Internet had transformed from a government and academic backwater to a new form of infrastructure—one that, as the *New York Times* reported, did "for the flow of information . . . what the transcontinental railroad did for the flow of goods a century ago."[15]

At first, Internet services were predominantly structured using a "client-server" model. Servers, owned by early "dot-com" companies, would run one or more computer programs, hosting websites and providing various types of applications, which Internet users could access through their clients. Information generally flowed one way—from a server to a client. Servers could share their resources with clients, but clients often could not share their resources with the server or other clients connected to the same Internet service.[16]

These early client-server systems were relatively secure but often acted as bottlenecks. Each online service had to maintain servers that were expensive to set up and operate. If a centrally managed server shut down, an entire service could stop working, and, if a server received too many requests from users, it could become overwhelmed, making the service temporarily unavailable.[17]

By the turn of the twenty-first century, new models for delivering online services had emerged. Instead of relying on a centralized server, parties began experimenting with peer-to-peer (P2P) networks, which relied on a decentralized infrastructure where each participant in the network (typically called a "peer" or a "node") acted as both a supplier and consumer of informational resources.[18] This new model gained mainstream popularity, with the launch of Napster. By running Napster's software, anyone could download music files from other users (acting as a client) while simultaneously serving music files to others (acting as a server). Using this approach, at its peak, Napster knitted together millions of computers across the globe, creating a massive music library.[19]

Napster's popularity, however, was short lived. Underlying the peer-to-peer network was a centrally controlled, continually updated index of all music available on the network. This index directed members to the music files they wanted, acting as a linchpin for the entire network.[20]

Although necessary for the network's operation, this centralized index proved to be Napster's downfall. Following lawsuits against Napster, courts found it liable for secondary copyright infringement, in part because it maintained this index. Napster was forced to manage the files available to peers on the network more carefully, and it scrubbed its index of copyright-protected music. Once this was implemented, the popularity of Napster waned and its users dispersed.[21]

Following Napster's defeat, a second generation of peer-to-peer networks emerged, bringing file sharing to an even larger audience. New peer-to-peer networks, such as Gnutella and BitTorrent, enabled people to share information about files located on their personal computers, without the need for centralized indices.[22] With Gnutella, users could find files by sending a search request, which was passed along from computer to computer on the network until the requested file was found on another peer's computer.[23] BitTorrent took an alternative approach, introducing the idea of fragmenting files into small pieces that could be downloaded from multiple users simultaneously, thus often making file transfer more rapid and efficient. BitTorrent initiated and coordinated the transmission of these chunks using small ".torrent" files that could be hosted on different servers,[24] thus avoiding the need for one overarching centralized service.

With the advent of these second-generation decentralized peer-to-peer networks, a new mode for content delivery had begun to solidify, untethering the exchange of information from large online operators. These decentralized networks lacked a discernible center, and fewer intermediaries supported these networks. Unlike Napster, these networks became nearly impossible to shut down.[25]

Digital Currencies

The idea of resilient, decentralized peer-to-peer networks resonated with a pocket of cryptographers and other technologists fascinated with advances in public-private key cryptography. These self-proclaimed "cypherpunks" realized the power of peer-to-peer networks and encryption, viewing both as tools to counteract erosions of personal freedom and liberty.[26]

Cypherpunks believed that without proper checks and balances, the deployment of modern information technology would narrow the sphere of personal privacy, resulting in pervasive government and corporate surveillance.[27] According to cryptographer David Chaum, founder of the International Association for Cryptologic Research, computing technology, over time, would rob individuals of their ability to monitor and control their information, which governments and corporations would collect and use "to infer individuals' life-styles, habits, whereabouts, and associations from data collected in ordinary consumer transactions."[28]

To counteract these perceived risks, cypherpunks advocated for the mass deployment of cryptographic tools, which they believed would preserve personal privacy while simultaneously undermining the hegemony of governments across the globe. They sought to democratize access to cryptography, building secure messaging systems, digital contracts, privacy-compliant identity systems, and "tamper-proof boxes."[29] By writing free and "widely dispersed" software that could not be "destroyed" or "shut down," they hoped to construct an "open society" that could escape the bonds of governmental or corporate control.[30]

The essential substrate of cypherpunks' dream was anonymous cash and other untraceable payment systems. Starting in 1983, cypherpunks and other cryptographers began exploring the use of public-private key cryptography to build new monetary systems. That year, Chaum proposed a system to enable the creation and transfer of electronic cash that would not require users to hand over personal information.[31] This system eventually turned into DigiCash, a company that Chaum launched in 1994.[32]

DigiCash relied on public-private key cryptography to issue a digital currency, using a digital signature system invented by Chaum (called blind signatures) to validate transactions between parties.[33] The company acted as a central clearinghouse, fixing the supply of money and processing DigiCash transactions. However, like Napster, DigiCash had a technical limitation. It operated via a client-server model, which required that Chaum's company double-check and validate every transaction on the network. The success of DigiCash was intimately tied to, and entirely dependent on, the fate of one company. When that company went bankrupt in 1998, DigiCash crumbled with it.[34]

The idea of creating an anonymous digital currency, however, exhibited a luster that was hard to dull. In the wake of DigiCash, a growing number of cypherpunks, including Hal Finney, Wai Dai, and Nick Szabo, embarked on a decade-long quest to build an anonymous digital currency that lacked centralized control.[35] These cypherpunks knew that, to create such a system, they would need to deploy one or more technologies that both control the supply of a digital currency and maintain a secure and authenticated record of who owned what at what time. Digital currency is just a series of bits stored in the memory of one or more machines. As opposed to dollar bills or metal coins, it does not have a physical instantiation. Hence, like any other digital resource, a unit of digital currency can be endlessly copied and reproduced. Because of these inherent features, digital currencies create obvious avenues for fraud. Without a central clearinghouse or any other intermediary capable of validating transactions and updating account balances, anyone in possession of a unit of digital cash would have the ability to send funds to two parties simultaneously, creating a "double spending" problem.[36] For example, if Bob owned $5 worth of digital currency, he could transfer that amount to both Alice and John at the same time, thereby illegitimately spending a total of $10.

Any decentralized payment system would need to solve this double-spending problem and would need to do so in a way that did not rely on any centralized intermediary. The total amount of the currency in circulation at any given time would need to be fixed, or controlled by a software protocol, so as to prevent individuals from devaluing the currency by generating additional unauthorized funds.[37] The system would also need to incorporate a secure and nonrepudiable record of transactions to keep track of all the digital currency flowing through the system. Without these essential characteristics, it would be impossible to validate who owned what amounts of digital currency at any given point in time without relying on a trusted authority or clearinghouse.

Bitcoin

In late 2008, one or more anonymous developers named Satoshi Nakamoto solved this problem by fusing together public-private key cryptography, digital signatures, and peer-to-peer technologies to create a new distributed database, which came to be known as a blockchain. Using a blockchain, Nakamoto built a decentralized digital currency that could operate without the need for a centralized middleman.

Unlike Chaum's DigiCash, which relied on a centralized operator, Nakamoto's system, outlined in a short nine-page article entitled "Bitcoin: A Peer-to-Peer Electronic Cash System,"[38] relied on a network of computers to validate and maintain a record of all Bitcoin transactions. Under this model, transactions were recorded in a common data store, and the underlying Bitcoin software controlled the supply of the digital currency and coordinated transaction validation, thereby eliminating the need for centralized control.[39]

Since its launch in 2009, Bitcoin has become one of the largest payment systems in the world, and yet its technical underpinnings are, for many, still as mysterious as its founding. One way to conceptualize how Bitcoin works is to think about e-mail. Today, an e-mail address enables us to send and receive electronic messages from anyone connected to the Internet in just a few seconds. E-mail addresses often are not tied to our individual identity;

they can be pseudonymous and act as a reference point to receive electronic messages. While many users rely on third-party operators to manage e-mail, the underlying protocol for sending and receiving messages is a free, open, and interoperable protocol that can be used without having to ask permission from anyone. Access to an e-mail inbox is maintained by a unique password, enabling people to control their e-mail accounts, either via a web interface such as Gmail or through an e-mail client such as Microsoft Outlook or Thunderbird.

Bitcoin is similar. As with e-mail, Bitcoin is an open and interoperable protocol not centrally controlled by any one party.[40] Bitcoin relies on public-private key cryptography to enable people to create pseudonymous Bitcoin accounts without asking permission from anyone. With a Bitcoin account, people can receive and send bitcoin to anyone around the world, in a matter of minutes, by executing and digitally signing a Bitcoin "transaction" with a private key. After a transaction is signed, members of the Bitcoin network verify that the transaction is valid and subsequently update the balances of relevant Bitcoin accounts.

People generally interact with the Bitcoin network by using a "wallet." Just like an e-mail client, Bitcoin wallets help people on the Bitcoin network manage their accounts. People store wallets on personal computers or maintain them using online applications, often maintained by third parties, making bitcoin readily accessible through a web browser or an everyday smartphone. For increased security, some people store their wallets offline on a USB flash drive or another form of secure hardware (often known as "cold wallets").[41] Like e-mail, Bitcoin transactions are unrestricted and flow freely across national borders.[42] No central party controls the transmission of the digital currency, and no one needs to authorize a transaction or pre-approve membership on the network. Anyone with a Bitcoin account can send or receive bitcoin in both large and small denominations (as low as 0.00000001 bitcoin, or about $0.0001758995 today).[43]

Records of transactions on the Bitcoin network are stored in the Bitcoin blockchain governed by underlying free and open-source software known as the Bitcoin protocol. Instead of swapping music or media files, computers on the Bitcoin network exchange information about new transactions occurring on the network. The Bitcoin protocol incorporates a mechanism that helps members of the network reach *consensus* as to whether a Bitcoin transaction is valid and whether it should be recorded to the Bitcoin blockchain.[44]

Unlike physical coins and currency, which pass between hands without leaving a trace, all Bitcoin transactions are recorded to the shared Bitcoin blockchain and are publicly auditable.[45] Anyone who chooses to join the Bitcoin network can download or review a full copy of the Bitcoin blockchain and trace through Bitcoin transactions.[46] Because of Bitcoin's transparent and open nature, the Bitcoin blockchain has become widely distributed and currently resides on thousands of computers in more than ninety-seven countries. Copies of the Bitcoin blockchain can be found scattered across large industrialized nations, such as the United States and China, and smaller jurisdictions, such as Cambodia and Belize.[47] Because the Bitcoin blockchain is redundantly stored across the globe and because of the payment network's reliance on a peer-to-peer network, Bitcoin is resilient and exceptionally difficult to shut down. So long as one computer maintains a copy of the Bitcoin blockchain, the Bitcoin network will continue to exist. Even in the case of a catastrophic event or an attempt by a local jurisdiction to shut down the network, the Bitcoin blockchain can be copied and replicated in a matter of a few hours (with a high-speed Internet connection).[48]

In many ways, the Bitcoin blockchain can be regarded as a tamper-resistant "book" with identical copies stored on a number of computers across the globe. Anyone can add new content to the book, and once new content has been added, all existing copies of the book are updated on computers running the Bitcoin protocol.

Unlike a book, however, Bitcoin is not organized by pages. Rather, bundles of Bitcoin transactions are grouped together into separate "blocks," which Bitcoin's protocol links together to form a sequential, timestamped "chain."[49] Each block stores information about transfers of bitcoin from one member of the network to another, along with other information that may be appended to each transaction (such as a poem, a prayer, a reference to an image, or some other file). Each block also contains a "header" used to organize the shared database.[50]

The core components of a block's header are a unique fingerprint (or a *hash*) of all transactions contained in that block, along with a timestamp and—importantly—a hash of the previous block. Hashes are generated using standard cryptographic hashing functions invented by the U.S. National Security Agency (NSA),[51] providing a way to represent the bundle of transactions in a block as a string of characters and numbers that are uniquely associated with that block's transactions.[52]

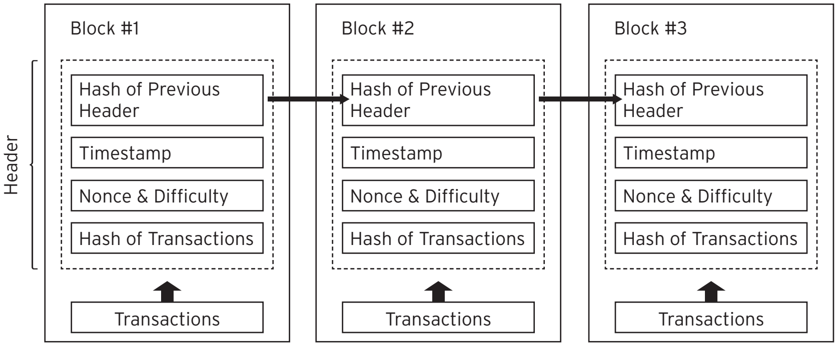

FIGURE 1.1 Simplified Bitcoin blockchain

While a book relies on page numbers to order its internal contents—making it possible for anyone to assemble a book in its appropriate order—the Bitcoin blockchain depends on data stored in each block's header to organize the shared database, which includes a hash of the previous block and a timestamp, creating a sequentially organized chain (see Figure 1.1).

To protect the security and integrity of the Bitcoin blockchain, Nakamoto built an ingenious system that makes it difficult to add information to the Bitcoin blockchain and even harder to modify or delete information once it has been saved. Storing information in the Bitcoin blockchain takes work and can only be achieved through collective effort. The Bitcoin protocol establishes a strict procedure for adding new blocks to the shared database, and all blocks are verified to ensure that they contain valid transactions and a valid hash.[53]

While generating a hash for any given block does not need to be challenging, the Bitcoin protocol purposefully makes this task difficult by requiring that a block's hash begin with a specified number of leading zeros. It does this by using a mathematical guessing game usually given the slight misnomer of "proof of work."[54] To generate a valid hash with the required number of leading zeros, parties on the network need to solve a mathematical puzzle and ensure that their solution contains at least as many leading zeros as required by the Bitcoin protocol at that point in time.[55]

Any computer trying to generate a valid hash must run through repeated calculations to meet the protocol's stringent requirements. Finding a solution to this mathematical puzzle does not require arbitrary computations; it is merely a game of trial and error, which is often referred to as "mining."[56]

The Bitcoin protocol adjusts the difficulty of the mathematical puzzle, depending on the number of miners on the Bitcoin network playing the proof of work game, to ensure that the network adds a new block approximately every ten minutes. The more nodes on the Bitcoin network attempting to solve the puzzle, the harder it becomes to generate a valid hash with an appropriate number of leading zeros.[57]

Once a miner finds a valid hash for a given block, the miner broadcasts the solution to the Bitcoin network. Once broadcast, other nodes in the network run a simple calculation to make sure that the resulting hash meets the Bitcoin protocol's specifications.[58] If valid, the block is added to the Bitcoin blockchain and stored on the local hard drives of active nodes. Through this process the network reaches *consensus* as to who owns what amount of bitcoin at that time. Occasionally, the Bitcoin network "forks," or splits into multiple copies, when different portions of the network append a different block to the blockchain. At times, forks occur as a result of a network split, possibly due to a malicious attack. At other times, forks occur when an updated version of the client running the Bitcoin network is released and a number of nodes connected to the network fail to update their software—as a result of negligence or because of an actual refusal to adopt the technical changes embodied in a new codebase.[59]

When the Bitcoin blockchain forks, the database's structure begins to resemble a tree rather than a linear chain. To make sure that the network eventually converges toward the same "branch" of the tree, the Bitcoin protocol implements a particular rule (a *fork choice*) stipulating that, in the case of a fork, miners should always pick the longest chain—that is, the branch with the most confirmed blocks as measured by computational power required to validate these blocks.[60]

This rule enables the Bitcoin protocol to preserve consensus throughout the network. If a majority of the network agrees on a particular chain of transactions, that chain is presumed valid. Bitcoin holders thus trust that, at any given time, those controlling a majority of the computational power supporting the Bitcoin network are acting in accordance with the protocol's rules, verifying transactions and recording new blocks to the longest chain.

Proof of work guessing game is useful not just for ensuring the orderly storage of records in the Bitcoin blockchain. This consensus algorithm also prevents people from creating fake transactions or otherwise altering the records stored in the Bitcoin blockchain. Because the header of each block incorporates a hash of the preceding block's header, anyone trying to modify the content stored in a block will inevitably break the chain. Even a small alteration will give rise to a new, unique hash tied to the altered block, and will necessarily trigger a change to the hashes of all subsequent blocks.[61]

Anyone willing to modify even a single record in the Bitcoin blockchain would have to go through the computationally expensive task of generating new hashes for every subsequent block in the Bitcoin blockchain. The more transactions that occur on the network—and the more blocks appended to the Bitcoin blockchain—the harder it becomes to retroactively modify previously recorded transactions. Moreover, because the Bitcoin blockchain operates via consensus, a would-be attacker or group of attackers would need to rewrite the transaction history of the Bitcoin blockchain at a pace that is faster than the majority of honest nodes supporting the network.

The most plausible way to change a record in the Bitcoin blockchain would be for a group of attackers to engage in a "51% attack" and effectively take over the network so that they can approve transactions at a rate that outpaces the rest of the network. Given the scale of Bitcoin's distributed network, such an attack is becoming increasingly remote. Starting in 2015, Bitcoin miners, according to *The Economist,* had "13,000 times more combined number-crunching power than the world's 500 biggest supercomputers."[62] Given the growth of the network, orchestrating such an attack today could thus cost hundreds of millions of dollars, if not billions, an operation that would prove too costly for most private parties or independent hacking coalitions.[63]

To balance the cost of engaging in Bitcoin mining, Nakamoto implemented a clever incentivization scheme to encourage people to maintain and secure the Bitcoin blockchain. Every time a miner generates a valid hash for a new block of transactions, the Bitcoin network will credit that miner's account with a specific amount of bitcoin—known as a "block reward"—along with transaction fees.[64] Miners on the Bitcoin network thus have an economic incentive to validate transactions and engage in the proof of work guessing game. By devoting their computational resources, they can earn a block reward and associated transaction fees for each block they validate and sell the digital currency to others on an open market.

Because the Bitcoin protocol is only programmed to allocate 21 million bitcoin, Nakamoto encouraged early adoption. The block reward progressively decreases over time—halving once every approximately four years, beginning when the network was launched in January 2009 and continuing until approximately 2140.[65] Miners supporting the network in its early days thus had the opportunity to earn more bitcoin.

By combining the proof of work consensus algorithm and the block reward incentivization scheme, Nakamoto developed a scheme capable of solving the double spending problem by building a decentralized system that could limit the supply of bitcoin and process transactions without the need for a central clearinghouse. Ultimately, Nakamoto created what can be regarded as a "state transition system." Every ten minutes, the Bitcoin network updates its "state," calculating the balances of all existing Bitcoin accounts. The proof of work consensus algorithm serves as a "state transition function" that takes the current state of the Bitcoin network and updates it with a set of new Bitcoin transactions.[66]

To evaluate whether a user has enough bitcoin to execute a transaction, the Bitcoin protocol searches through all previous transactions, starting from Bitcoin's first block (the "genesis block"). If a user has enough bitcoin, the transaction is deemed valid and will be bundled into a block. Once a miner generates a valid hash with proof of work, which other miners confirm, the "state" of the Bitcoin network is updated, including the balances of the accounts involved in the block's transactions.[67] If, however, a user does not have enough bitcoin, the transaction will be rejected by the network. Miners will not bundle the transaction into a block, and the invalid transaction will not impact the network's state.

Through this technical design, even though Bitcoin lacks a centralized clearinghouse, users gain assurance that the balance of every Bitcoin account is accurate at any given time. The protocol has been implemented in such a way as to enable trusted peer-to-peer interactions between people who do not know, and therefore do not necessarily trust, one another. This is why Bitcoin—and blockchain technology more generally—has been described as a trustless system.[68] Instead of relying on a centralized trusted authority or middleman, people only need to trust the underlying code and the miners supporting the Bitcoin blockchain.

Ethereum

With Bitcoin's launch and rapid growth, an increasing number of programmers began to explore the technology for uses outside of just digital currencies. In the wake of Bitcoin, hundreds of blockchains and new digital currencies appeared, seemingly overnight.[69] Centralized wallet services emerged to help people join the Bitcoin network.[70] Exchanges were launched that made it possible to trade bitcoin for traditional fiat currencies such as the U.S. dollar, euro, and Chinese yuan.[71] The price of bitcoin skyrocketed, hitting a high of $1,200 (albeit briefly) in 2013.[72] Interest in Bitcoin grew, and venture capitalists and traditional businesses, such as Microsoft and Dell, explored Bitcoin as a payment option.[73]

However, the more people who considered Bitcoin, the more its limitations became apparent. Bitcoin excelled as a platform to facilitate the exchange of digital currency, but without updating the underlying protocol, it could not be used for much more. The Bitcoin network was slow—it could only reach consensus and validate transactions roughly every ten minutes—and therefore questions emerged as to how much information the Bitcoin blockchain could store. Bitcoin's decentralized structure made its protocol hard to update and improve, and the network lacked formal governance, relying on the efforts of a small group of developers who slowly revise and fix bugs in the underlying software.[74]

New blockchain-based projects were launched with the hope of addressing these limitations. They sought to leverage the power of a blockchain not just to store information related to the transfer of a digital currency but to build a medium to host decentralized applications (or "dapps") that rely on a blockchain, at least partially, for their underlying functionality.[75]

At a generalized level, because a blockchain is a data storage system, the technology may be deployed for far more than just storing data related to Bitcoin transactions. Blockchains are equipped to store or reference other forms of information, including what are essentially small computer programs—which technologists often refer to as *smart contracts*.[76]

The first blockchain to enable the creation and deployment of sophisticated smart contracts was the Ethereum blockchain. Announced in February 2014 and launched roughly a year and a half later, this second-generation blockchain-based network built on the efforts of Bitcoin but added richer functionality to enable parties to deploy smart contracts on a blockchain just as one would deploy code for a website on a server today.[77]

Like Bitcoin, Ethereum is a peer-to-peer network governed by a free and open-source protocol. Ethereum implements a native digital currency (ether), which is allocated to miners supporting the network and can be transferred just like bitcoin. The Ethereum blockchain uses a similar proof of work mechanism to update the state of the Ethereum blockchain.[78]

However, unlike Bitcoin, Ethereum is faster and has a greater range of capabilities when it comes to smart contracts. The Ethereum blockchain is updated roughly every twelve seconds as opposed to every ten minutes.[79] Ethereum also implements a Turing-complete programming language called Solidity, which makes it possible for anyone to write smart contracts and deploy decentralized applications. With Solidity, it is theoretically possible to execute a range of complex computations on a peer-to-peer network.[80]

As opposed to Bitcoin, which only has one type of account, on the Ethereum network, there are two different kinds of accounts: one for everyday users of the network (known as an "externally owned account") and one for smart contract applications (known as a "contract account"). A contract account has a public address on the Ethereum network but does not come with a private key. It stores the compiled bitcode of a particular smart contract and can collect and distribute ether, record data to the Ethereum blockchain, process information, and possibly also trigger the execution of other smart contracts. An externally controlled account is different. As with Bitcoin, this account is assigned a public address. Anyone with access to the account's private key can send ether to other members of the network and can interact with smart contracts stored in a contract account.[81]

The part of the Ethereum protocol responsible for processing smart contracts is the Ethereum Virtual Machine (EVM). From a practical standpoint, the EVM can be thought of as a decentralized virtual machine running a number of smart contract programs. As a general rule, anyone can trigger the execution of a smart contract by sending an ether transaction to the corresponding contract account, thereby setting Ethereum's wheels in motion.[82]

Contracts interact by either "receiving" or "sending" messages. A "message" is an object containing a particular quantity of ether, an array of data, the address of the sender, and a destination address (which can be either another contract account or an externally owned account). When a contract receives a message, it has the option of returning a message to the original sender—acting, to a large extent, like a standard computer function.[83]

By design, every operation processed by the Ethereum Virtual Machine is executed by every active node on the Ethereum network. Through this implementation, any contract on the Ethereum EVM can trigger any other contract at almost zero cost. To prevent abuse, the Ethereum protocol charges a small fee—referred to as "gas"—for each computational step.[84]

To avoid excessive price fluctuations, the price of gas is not fixed but is dynamically adjusted by miners based on the market price of ether. The Ethereum protocol also implements a floating limit on the number of operations that can be contained in a block, forcing miners on the Ethereum network to charge a gas fee that is commensurate with the cost of the transaction for others on the network.[85]

The open and decentralized nature of Ethereum allows smart contracts to be deployed pseudonymously and to operate in a largely autonomous manner. Because all active nodes on Ethereum run the code of every smart contract, the code is not controlled by—and cannot be halted by—any single party. In a sense, a smart contract operates like an autonomous agent, automatically reacting to inputs received from externally owned accounts or other smart contract programs executed on the network.[86]

Hundreds of thousands of smart contracts have been deployed since Ethereum's launch.[87] These smart contracts are capable of processing basic logic ("if this, then that") and can be used to generate and transfer tokens (associated with physical or digital assets), verify signatures, record votes, and implement new blockchain-based governance systems.[88]

Given the underlying design of the network, however, running code via the Ethereum EVM is slow and potentially expensive.[89] Despite these limitations, Ethereum is paving the way for a new paradigm of computing—one where software applications are no longer controlled by a central authority but rather operate autonomously on a decentralized, peer-to-peer network.

Decentralized File Sharing and Overlay Networks

To extend the capabilities of Bitcoin, Ethereum, and other blockchains, new decentralized protocols are being developed, making it possible for blockchains to manage the transfer of additional assets (beyond just digital currency) and enabling smart contracts to interact with, and potentially control, other digital files. These protocols serve as "overlay networks," extending the power and usefulness of these new data structures.[90]

For example, protocols such as Color Coin enable parties to use the Bitcoin network to create tokens that represent a range of valuable assets. Using the Color Coin protocol, parties can send a transaction for a nominal amount of bitcoin (or another digital currency) and append some metadata to the transaction to indicate that the transaction in fact represents the transfer of a tangible or digital asset, such as a stock certificate, a title to a copyrighted work, or a vote. In other words, using the Color Coin protocol, the transfer of 0.0000000001 bitcoin may constitute the transfer of a share of Google's stock.[91]

More ambitiously, new decentralized file-sharing protocols enable people to store files on a peer-to-peer network and control access to those files using smart contracts, creating new tools to build robust and complex decentralized applications. Because the Bitcoin blockchain can only store a limited amount of information per transaction, and because the Ethereum blockchain charges for each computational step in a smart contract program, it is often prohibitively expensive to build decentralized applications that rely on a blockchain for file storage.

New distributed storage platforms and content distribution services, such as Swarm and Filecoin (powered by the IPFS protocol), are trying to address these limitations to support more advanced blockchain-based uses. These new systems aim to provide secure and resilient peer-to-peer storage for blockchain-based networks, with no central administration, zero downtime, and the ability to operate even if members of the network leave.[92]

Like Bitcoin and Ethereum, these overlay networks come with built-in incentive mechanisms to encourage adoption and use. Members of these networks receive compensation for storing data and serving it to third parties. Those with extra bandwidth or space on their hard drive can participate on these networks and voluntarily agree to store small portions of files (called chunks or shards), which are reassembled on demand by these decentralized file-sharing protocols.[93]

By pooling bandwidth and storage resources, the technologists behind these decentralized systems believe that, through the power of numbers, they will gain a leg up on existing online services. Instead of building an online application by renting space from a traditional cloud service provider such as Amazon, Microsoft, or IBM, or relying on large centralized intermediaries such as Facebook and Google, to host data, new decentralized applications will rely on blockchain technology and peer-to-peer networks for coordination and storage. If successful, blockchains and these new decentralized protocols may underpin an entirely new Internet architecture (sometimes referred to as "Web 3.0").[94]

Permissioned Blockchains

Because blockchains are decentralized and enable the deployment of potentially autonomous code, efforts also are under way to harness the power of blockchains in a more controlled and predictable manner. Bitcoin, Ethereum, and other "permissionless" blockchains are open and accessible to everyone. Anyone with an Internet connection can download the open source software governing these blockchains and participate in the network, without revealing their true identity or asking for prior permission. The open, decentralized, and pseudonymous nature of permissionless blockchains causes concern when deployed in heavily regulated areas such as banking and finance, which require that financial institutions track and vet parties and report suspicious activity.[95]

Alongside Bitcoin and Ethereum, a number of alternative "permissioned" blockchains have emerged. These blockchains rely on a peer-to-peer network, but they are not open for anyone to join. Rather, a central authority or consortium selects the parties permitted to engage in a blockchain-based network, imposing limits on who can access or record information to the shared database.[96] Consortium members ultimately control membership, thus creating an environment where each party on the network is known or somewhat trusted.

Existing permissioned blockchains often are purpose driven. For instance, the Ripple protocol uses a permissioned blockchain to facilitate the exchange of currencies and other stores of value, such as U.S. dollars, euros, and even gold. Underlying the Ripple protocol is an alternative consensus mechanism that differs from Bitcoin's and Ethereum's proof of work in that it relies on collectively trusted subnetworks within the larger Ripple network to process and validate transactions.[97]

Currently, one notable advantage of permissioned blockchains is speed. In an open and permissionless network, such as Ethereum and Bitcoin, active nodes need to reach consensus as to the validity of every transaction. These networks can only process transactions every ten minutes in the case of Bitcoin and every twelve seconds in the case of Ethereum, lagging behind modern databases, which store information in milliseconds. Because permissioned blockchains tend to be operated by a smaller number of preselected participants, they can implement alternative ways to validate and approve transactions, often in a faster manner.[98]

Despite creating benefits in terms of speed and predictability, permissioned blockchains ultimately suffer from one important drawback. Trust is fickle. With permissioned blockchains, there is no guarantee that parties will not collude to tamper with the underlying blockchain in ways that may ultimately harm other network participants. If only a handful of parties can validate and record information to a blockchain, these parties become single points of failure (and control), which could be compromised by technical failures, corruption, or hacking. These security limitations could potentially result in catastrophic consequences, especially if permissioned blockchains grow to support important social or economic systems.[99]

Security concerns are often brushed aside by those developing permissioned blockchains. Some argue that permissioned blockchains will supplant and eventually eliminate the need for public blockchains such as Bitcoin and Ethereum. Others view them as a temporary solution—much like intranets, which dominated corporate America in the mid-1990s.[100]

At least in the short term, the future of blockchain technology likely will involve both permissioned and permissionless blockchains, working together in harmony. Fast-paced transactions could occur on permissioned blockchains, while—for the sake of security—the state of these networks could be secured by open and permissionless blockchains. In effect, one can envision a future where open and permissionless blockchains serve as a backbone for a broader ecosystem comprised of different permissioned blockchains focused on specific uses.[101]

As of right now, the truly innovative aspect of blockchain technology does not lie in private and permissioned blockchains; it rests with those that are public and permissionless. Public blockchains can be deployed for lawful purposes, but they can also support the emergence of a variety of unlawful systems that are difficult to halt and control.

Because of the unique characteristics of public and permissionless blockchains, they are the focus of this book.[102] We turn now to a description of the unique characteristics of these systems, and an explanation of how new applications are employing these decentralized technologies to establish a new set of cryptographically secured rules, which we term *lex cryptographica.*

2

Characteristics of Blockchains

Blockchain technology constitutes a new infrastructure for the storage of data and the management of software applications, decreasing the need for centralized middlemen. While databases often sit invisibly behind the scenes, their significance cannot be understated. Databases serve as a backbone for every platform, website, app, or other online service. Up to this point, databases have for the most part been maintained by centralized intermediaries, such as large Internet companies or cloud computing operators such as Amazon, Microsoft, and Google. Blockchains are changing this dynamic, powering a new generation of disintermediated peer-to-peer applications, which are less dependent on centralized control.

WHILE COMPLEX, blockchains exhibit a set of core characteristics, which flow from the technology's reliance on a peer-to-peer network, public-private key cryptography, and consensus mechanisms. Blockchains are disintermediated and transnational. They are resilient and resistant to change, and enable people to store nonrepudiable data, pseudonymously, in a transparent manner. Most—if not all—blockchain-based networks feature market-based or game-theoretical mechanisms for reaching consensus, which can be used to coordinate people or machines. These characteristics, when combined, enable the deployment of autonomous software and explain why blockchains serve as a powerful new tool to facilitate economic and social activity that otherwise would be difficult to achieve.

At the same time, these characteristics represent the technology's greatest limitations. The disintermediated and transnational nature of blockchains makes the technology difficult to govern and makes it difficult to implement changes to a blockchain's underlying software protocol. Because blockchains are pseudonymous and have a tamper-resistant data structure supported by decentralized consensus mechanisms, they can be used to coordinate socially unacceptable or criminal conduct, including conduct facilitated by autonomous software programs. Moreover, because blockchains are transparent and traceable, they are prone to being co-opted by governments or corporations, transforming the technology into a powerful tool for surveillance and control.

Disintermediation and Transnational Networks

Today, as noted earlier, online services are delivered primarily via a client-server model. To interact on the Internet, users rely on trusted authorities or middlemen, which assume a variety of roles. Some of these middlemen are responsible for creating marketplaces between buyers and sellers (as in the case of eBay and Uber). Some are responsible for storing and maintaining repositories of information collected from disparate parties across the web (as in the case of Facebook, YouTube, and Wikipedia). Others serve as authenticated sources of specific goods or services (as in the case of PayPal and Spotify).

Blockchains operate under a different hierarchical structure. They are supported by a network of computers, linked together via an overarching software protocol. At a generalized level, no single party controls a blockchain, and blockchains do not rely on one centralized party for their maintenance or operation. These shared databases operate globally and extend across national borders. Because they do not come with any centralized authority or gatekeeper, anyone with an Internet connection can retrieve information stored on a blockchain simply by downloading freely available open source software.[1]

These characteristics give blockchains the potential to support increasingly disintermediated and global services, allowing parties to engage more directly with one another for a variety of reasons. New services can rely on a blockchain to store information, transfer value, or coordinate social or economic activity, with less of a need to pass through centralized choke points.

The distributed and transnational nature of blockchains, however, comes with tradeoffs: the larger and more distributed the blockchain-based network, the more complex and challenging it is to manage. Most blockchain-based protocols are open source software, developed by loosely connected teams of software developers, who often work on these systems on a vocational basis.[2] These programmers may be skilled and technically proficient, but they often operate outside of formal organizational structures or legal entities, which are responsible for the operations and maintenance of large-scale systems.[3]

Open source software, like every other piece of software, contains bugs. Despite their best intentions, open source developers may be incapable of patching errors at a sufficiently high rate for blockchain-based protocols to mature into highly reliable systems. If blockchains are used to power protocols responsible for transferring value and structuring social and economic activity, a lack of formal governance may limit a blockchain's overall usefulness.[4]

The distributed nature and transnational scope of blockchains also raise important jurisdictional concerns. From a technical perspective, national borders are largely irrelevant for the operation of blockchain-based networks. And in fact, many blockchain-based services and applications aim to operate worldwide. These systems span the globe, fueling questions about how governments can regulate and, if necessary, constrain blockchains or associated decentralized applications.

Resiliency and Tamper Resistance

Blockchains are characterized not just by their global and transnational nature. They also store data in a unique manner, at least as compared to existing data structures. Given the distributed nature of a blockchain, along with consensus mechanisms (such as proof of work) and one-way hashing algorithms, once information has been recorded to a blockchain, it becomes exceptionally hard to change or delete. No single party has the power to modify or roll back information stored on a blockchain, and no single party can halt the execution of a smart contract once it has been deployed, unless provided for in the code.[5]

On popular blockchain-based networks, the entire blockchain is replicated across thousands of different computers scattered across the globe. These computers store exact or nearly exact copies of the blockchain, and the underlying software protocol ensures that copies are consistently updated whenever a party connects to the network.

By using this approach, if one copy of a blockchain fails or is somehow corrupted, the event has little impact on the broader network, making blockchains difficult to shut down and censor. If a single computer on a network has a complete copy of a blockchain, that blockchain will remain available for others to access and use. As long as there is an Internet connection, a blockchain can be replicated, and the network can be rebuilt. Even if Internet connectivity is shut down in one region of the world, because of a governmental measure or a cataclysmic event, the rest of the network supporting a blockchain will still retain the ability to store new information and access previously recorded data. As soon as Internet connectivity is restored, parties in those previously excluded regions can update their personal copies of a blockchain and continue to participate in the network, picking up where others have left off.[6]

Beyond being resilient, blockchains store information that is highly resistant to change. Those seeking to alter a blockchain must either wage an expensive takeover or engage in a complex, and often contentious, public debate to convince other network participants to implement a change underlying protocol. For example, with blockchains relying on a proof of work consensus mechanism, parties seeking to modify a blockchain would need to deploy sufficient computational resources to generate blocks faster than other honest parties supporting the network—a task that is costly on large blockchain-based networks like Bitcoin and Ethereum. Alternatively, a majority of miners (as measured by the miners' computational power) must collectively decide to update a blockchain's underlying protocol to unwind previously recorded transactions or to block certain accounts or smart contracts.

A network's voluntary decision to update a protocol thus carries with it political and social dimensions. As with every political system, reaching consensus among disparate groups, with different individual preferences and motives, is a difficult and often time-consuming process. Parties aiming to alter a blockchain must make their case as to why a blockchain should be modified or amended and reach out to miners (for example, via social media or in-person meetings). If a majority of nodes supporting the network do not agree on a change, a blockchain will remain the same.

The technical design of blockchains therefore favors the status quo, making blockchain-based networks highly resistant to change. Nodes supporting a blockchain-based network ultimately have the power to decide whether to alter its state. If network participants aim to build an "immutable" database—as has been the case with Bitcoin so far—the data stored on a blockchain may never change once it has been recorded, unless it is compromised by malicious parties.

The tamper-resistant and resilient nature of blockchains also creates complications for governments and regulators. Unless a government can succeed in taking over a blockchain or can convince miners and other relevant stakeholders to modify a blockchain's protocol, any data or programs stored on a blockchain cannot be altered, creating incentives for the technology to be used for unlawful or illicit purposes.

Transparent and Nonrepudiable Data

Because blockchains rely on peer-to-peer networks and digital signatures, the data they store is both transparent and nonrepudiable. Information maintained on a blockchain is authenticated, and metadata and other contextual information about blockchain-based transactions are available for others to view. Anyone can download a blockchain and assess whether a given account was involved in a transaction or—as in the case of Ethereum—whether an account interacted with a smart contract.

In effect, a blockchain serves as an auditable trail of activity occurring on a peer-to-peer network. Although information stored on a blockchain may be encrypted, contextual information about what accounts are engaging in transactions or interacting with smart contracts is, in most cases, publicly available for anyone to view.

All transaction data stored on a blockchain is not just auditable but also authenticated and nonrepudiable. Because blockchains rely on public-private key encryption and digital signatures, once a transaction occurs on a blockchain-based network, parties subject to that transaction will have a hard time denying involvement. Before engaging in a transaction with a smart contract or another member of a blockchain-based network, a party must sign the transaction with their private key. The digital signature serves as evidence that an account initiated a transaction, narrowing the ability of the holder of a blockchain-based account to refute the fact that a transaction occurred, unless a party can prove that the private key associated with the account was somehow compromised.[7]

The combination of transparency and nonrepudiability, together with the resilient and tamper-resistant nature of data stored on a blockchain, helps create trust in the network. Digital signatures provide a high degree of assurance that parties to a peer-to-peer transaction intended to be bound by its terms. A blockchain's transparency means that parties can subsequently review a blockchain and verify that a transaction has indeed occurred. Because data recorded on a blockchain is tamper resistant and resilient, blockchains provide solace that information related to a transaction has not been altered in an opportunistic way—and will not be in the future.

These characteristics also enable parties to rely on a blockchain to publicly disseminate authenticated information. Parties on a blockchain-based network can choose to reveal their public address to prove that they are the source of information or to prove that they engaged in certain transactions. By doing so, the public can verify that the party was, with a high probability, the source of the recorded information.[8] For instance, in late 2016, questions swirled about whether Julian Assange, the prominent cypherpunk and founder of Wikileaks, was still alive. Conspiracy theories about his death bounced around the Internet, including stories posted to online communities like Reddit and 8chan. Assange managed to counteract these rumors without making any public appearance. He used a Bitcoin address widely known to be associated with WikiLeaks to execute a series of transactions with the following hidden message: "We're Fine. 8 Chan Post [is] Fake."[9] The blockchain provided the necessary infrastructure to prove the integrity of the message and the authenticity of its source, in a way that could not easily be repudiated.

By implication, private actors and governmental authorities have the option to store authenticated information on a blockchain, making the information available worldwide to anyone with an Internet connection. Government records no longer need to be stored on paper or in centralized silos; they can be digitized and stored on a blockchain for the benefit of the public.

Pseudonymity

Blockchains are further characterized by their pseudonymity. By relying on digital signatures and public-private key cryptography, blockchains make it possible for a person to store information or engage in transactions without revealing one's true identity.[10] With blockchains, parties can interact with one another even if they do not trust each other—provided that they trust the underlying technical infrastructure and the rules embedded in a blockchain's protocol.

Pseudonymity, however, carries tradeoffs and creates incentives for parties to engage in unlawful social and economic activity. For example, pseudonymity may embolden parties to use a blockchain-based digital currency to pay for drugs or other unlawful goods. The digital currency could also be used to launder money or engage in tax evasion.[11]

Today, these risks have constraints. On networks like Bitcoin and Ethereum, anyone who has access to a blockchain-based network can rely on information about events occurring on the network to uncover a party's identity. Contextual information related to blockchain-based transactions can be combed through to deanonymize individuals. For example, by analyzing the flow of money between network participants through a process called "transaction graph analysis," researchers from the University of San Diego and George Mason University successfully identified clusters of merchants and customers relying on the Bitcoin blockchain.[12] Likewise, researchers at the University of Maryland and Cornell have scrutinized transactions related to several smart contracts operating on the Ethereum network, shedding light on their operation.[13]

Over time, however, blockchains may become increasingly anonymous, making transaction graph analyses and comparable tracing techniques increasingly difficult. Services already have sprung up to mix and scramble Bitcoin transactions to mask parties' identities. Recently launched blockchains, such as Zcash and Monero, are hiding the source, destination, and amount of digital currency transferred within these blockchain-based networks by using advanced cryptography such as zero-knowledge proofs and ring signatures.[14]

If these obfuscation and anonymization techniques gain widespread adoption and operate as claimed,[15] the risk posed by blockchain-based networks will likely expand. These networks may no longer be pseudonymous but rather could morph into global and truly anonymous networks that facilitate low-cost and low-friction exchanges.

Incentivization and Cost Structures

The protocols governing some of the most advanced blockchains also have complex incentivization and market-based schemes for engaging in transactions and running smart contracts. By using block rewards and transaction fees, blockchains incorporate payoff structures designed to reward parties that maintain a blockchain-based network.[16] These incentivization structures influence how parties process transactions for a blockchain and impact the types of transactions and smart contracts that parties deploy.

For example, the proof of work consensus mechanisms and the incentivization structures of Bitcoin and Ethereum have encouraged the consolidation of mining. To validate transactions for these blockchains, a miner must dedicate processing power to verify transactions and rapidly churn through potential solutions to solve the mathematical puzzle associated with each block. Blockchains such as Bitcoin and Ethereum dynamically adjust the difficulty of this puzzle as more and more processing power is added to the network.[17]

As these networks have increased in popularity, the difficulty of generating a valid hash has grown dramatically, thus decreasing the probability that an individual using an everyday computer will mine a block. A miner's probability of finding a valid hash for a block is roughly proportional to the percentage of a blockchain-based network's total computational resources that the miner contributes. Therefore, on large networks like Bitcoin, the probability that any individual could mine a block using nonspecialized hardware is low.[18]

Because of this dynamic, parties seeking to validate transactions for Bitcoin and Ethereum have organized themselves into "mining pools," combining their computational resources and deploying specialized hardware, such as application-specific integrated circuits (ASICs), to mine new blocks and share block rewards and fees—like waiters pooling tips at a restaurant. By joining forces, miners in a pool increase the likelihood that they will earn a block reward and can distribute any collected digital currency to the pool.[19]

These pools largely control the processing of transactions on Bitcoin and Ethereum. As of December 2017, four mining pools controlled over 50 percent of the Bitcoin network and two mining pools controlled more than 50 percent of Ethereum.[20] These pools thus have the power to control the operation of Bitcoin and Ethereum and shape their development.

Incentivization mechanisms and payoff structures also influence the decision-making processes of parties using a blockchain-based network to store information, transfer digital currency, or interact with smart contracts, because all of these operations carry with them certain fees. For example, when transferring bitcoin, parties can set the fee they are willing to pay miners to incorporate their transaction into a new block.[21] Suppose that Alice needs to send bitcoin to Bob. Alice can pay miners a higher fee to help ensure that they rapidly process the transaction. Conversely, if Alice can wait, she may only pay miners a small fee, hoping that miners will eventually decide to process her transaction.

A similar cost structure underpins the Ethereum network. Each computational step of a smart contract has a cost, influencing the types of programs that people are willing to store on a blockchain. If a smart contract is costly to execute, people may choose not to interact with it because they do not want to pay miners to execute its underlying code. Instead, they may decide to build all or parts of a blockchain-based application by relying on overlay networks or more centralized alternatives.[22]

Initially, fees on popular blockchain-based networks remained relatively low, only costing users a couple of cents to store information, engage in a transaction, or execute a smart contract. As these networks gain broader adoption and the number of processed transactions increases, transaction fees may limit the operations of these blockchain-based systems or make them less attractive compared to centralized alternatives.

Indeed, at least for the Bitcoin network, its incentivization structure may spell its ultimate doom. The Bitcoin protocol is programmed to issue 21 million bitcoins, allocated to miners as block rewards roughly every ten minutes.[23] Once these allocations stop, there are questions about whether miners will retain sufficient incentive to operate the network or whether transaction fees will simply become too costly.

To decide whether to support a blockchain, miners often run a simple calculation. They multiply the expected market value of a block reward and associated fees with the probability of receiving this prize. Miners then compare this amount to the cost of purchasing the computational resources necessary to engage in the proof of work consensus protocol. In general, rational miners will only choose to support a blockchain-based network like Bitcoin if the expected reward exceeds the computational costs.[24]

If the block rewards or transaction fees are too low, miners may decide that it is no longer profitable to support the Bitcoin blockchain. In turn, the difficulty of Bitcoin's proof of work puzzles may decrease, opening up the network to malicious attacks that could result in a potential manipulation of the Bitcoin blockchain.

Conversely, if Bitcoin's transaction fees increase (as they are expected to do once the Bitcoin protocol stops issuing block rewards), sending bitcoin may become expensive, making it less likely that people will choose to rely on this network as opposed to more centralized alternatives, thereby causing interest in bitcoin to wane.[25]

The same holds true in the context of the Ethereum network. If the costs of running a smart contract outweigh its anticipated benefits, there will be little incentive for people to use Ethereum—making either centralized alternatives or competing blockchain-based solutions more attractive.

Consensus

Another core characteristic of blockchains is their ability to coordinate social activity and help people reach an agreement as to a particular state of affairs. Underlying each blockchain-based network is a consensus mechanism that governs how information can be added to the shared repository. Consensus mechanisms make it possible for a distributed network of peers to record information to a blockchain, in an orderly manner, without the need to rely on any centralized operator or middleman.[26] Because data recorded on a blockchain is visible to all and is hard to repudiate and retroactively modify, groups of people who do not know—and therefore do not trust—one another can rely on this new data structure to coordinate their activity, with less of a need for trusted authorities.

Blockchains, however, are capable of storing more than mere records about the transfer of digital currencies. They can store data, messages, votes, and other kinds of information that can be encoded in a digital format. For instance, the Bitcoin blockchain has been used as a repository for different kinds of information—from prayers and eulogies to messages and images ranging from the sophomoric to the sublime.[27] More generally, a blockchain can be regarded as a shared repository of information—an open, low-cost, resilient, and secure storage system that nobody owns but many people maintain.

Overlay networks enhance the capacity of blockchain-based applications to coordinate human and machine interactions. The ability to store information using overlay networks, combined with the ability to structure services using smart contracts and transfer value almost instantaneously via a blockchain, makes blockchains a new tool for managing the activities of loosely connected groups of individuals and machines. An organization can use a blockchain to reach consensus and use smart contracts to govern contractual relationships and facilitate payments between parties.

Because blockchains help people reach consensus, they may help solve some of the issues traditionally associated with shared common-pool resources—such as the free-rider problem or the tragedy of the commons.[28] Transparent data storage and smart contracts could be used by communities to help them reach an agreement and self-govern. For instance, by recording every interaction on a public blockchain and encoding rules linking these interactions to specific transactions—such as the assignment of tokens or the allotment of small payments of digital currency—blockchains can help commons-based communities govern themselves through decentralized incentive systems. While online communities will probably be the first ones to experiment with these new organizational structures, as the ease of using these tools decreases over time, they could eventually be used to implement organizational structures that also operate in the physical world.

Autonomy

Perhaps most profoundly, blockchains are characterized by their ability to facilitate the deployment of autonomous software that is not under the control of any one party. Today, code is generally maintained and executed by intermediaries on centralized servers. These operators ultimately retain control over the code's execution and have the power to stop code from executing if so desired.[29] If necessary, they can prevent a party from running a program that may cause damage or inflict harm.

Blockchains lack these limitations. By relying on a peer-to-peer network and a consensus mechanism, they facilitate the execution of computer code in a way that is entirely independent of any one party. Indeed, transfers of bitcoin are executed automatically on the Bitcoin network, so long as parties comply with the protocol's strict requirements. Once submitted to the network, Bitcoin transactions cannot be reversed, and no single party can halt their execution.

Similarly, on the Ethereum network, smart contract code is run in a distributed manner by all active nodes in the network using the Ethereum Virtual Machine. After a smart contract has been deployed, little can be done to change its underlying logic—unless the party deploying the smart contract has introduced a mechanism to do so. Because all nodes on the Ethereum network are responsible for running the smart contract code,[30] even if a handful of nodes refuse to execute a smart contract's code, these nodes cannot stop others from running the code, except by advocating for a change in the Ethereum protocol.

Blockchains thus enable the creation of autonomous software programs run through the collaborative effort of parties with different incentives and in different locations scattered across the globe, none of which can unilaterally affect the code's execution. Once deployed on a blockchain, these programs no longer need or necessarily heed their creators; they are run on a decentralized network, making it difficult to unwind or halt their execution.[31]

One important advantage of these autonomous systems is that—if properly designed—they can handle basic economic transactions at lower costs, with higher degrees of reliability and potentially greater speeds. These blockchain-based systems can reduce or even eliminate the need for human oversight, narrowing the possibility for parties to act opportunistically in ways that benefit the few at the expense of the many.

At the same time, the deployment of blockchain-based software creates systems that are highly deterministic. If, for example, a party mistakenly sends bitcoin to the wrong address, it can be difficult to unwind the transaction retroactively. Likewise, if the code of a smart contract on the Ethereum network is faulty, parties would need to reverse the transactions or initiate an after-the-fact legal action to secure the return of any exchanged value unless otherwise provided for in the code.[32]

Autonomy also creates opportunities for certain types of activities that can be both lawful and unlawful. As a general rule, because of their decentralized and transnational nature, blockchain-based systems exhibit a degree of *alegality*.[33] Autonomous systems do not need to abide by existing rules and jurisdictional constraints; they can be designed to bypass or simply ignore the laws of a particular jurisdiction. Once deployed on a blockchain, these systems will continue to operate—even if they are socially unacceptable, morally wrong, or potentially damaging to humans—so long as there are sufficient incentives for miners to support that blockchain.

The *alegal* nature of blockchains, combined with the autonomous nature of smart contracts, may prove attractive to criminals, who would be able to engage in binding transactions with one another, even if they do not trust each other. These systems can enable bad actors to coordinate their activities in a decentralized way, without the need to rely on any intermediary that could be easily infiltrated or compromised by law-enforcement officials.[34] When combined with cryptographically secure communication channels, blockchains can thus facilitate illicit activity and make such activity harder to stop or intercept.[35] With fewer intermediaries involved in these criminal operations, governments may struggle to find ways to stop these illegal acts.

The Dual Nature of Blockchains

When viewed as a whole, blockchains possess competing characteristics that wrap the technology in opportunities and contradictions. This ultimately means that blockchains can be used both for good and for bad. The technology can power new automated systems that operate globally and at low cost, bringing new efficiencies in the realm of finance, media, and law, as well as in the public sector. Blockchains can be used to prevent certain types of criminal activities while simultaneously making it easier for criminals to operate under the radar. The technology can make it harder to restrict the flow of information, undermining efforts by authoritarian regimes to censor their citizens while simultaneously enabling governments to track an increasing range of financial and nonfinancial transactions—opening up new avenues for surveillance and control. Indeed, we are already seeing the dual nature of blockchain technology in a series of use cases, which currently fall into three distinct categories.

First, as demonstrated by Bitcoin, blockchain technology enables the creation of decentralized, global value transfer systems that are both transnational and pseudonymous. By using a blockchain, parties can transfer digital currencies or other valuable assets without the need to rely on a centralized clearinghouse or trusted authority. These blockchain-based systems can decrease the cost of transferring value across the globe, serving as a new payment and financial layer for the Internet.

Second, blockchain technology allows for the development of autonomous systems, which are not controlled by any single party. By using smart contracts, people can build decentralized applications that enable value transfer and enable disparate groups to achieve consensus, potentially even pseudonymously. Instead of relying on standard legal agreements, parties to a contract can use a smart contract to stipulate certain contractual rights and obligations and build dynamic agreements that bind parties together in more concrete ways.[36] People can create virtual corporations and decentralized (autonomous) organizations to help disparate groups of individuals achieve consensus in a pseudonymous and nonhierarchical manner. Blockchain-based systems can manage Internet-connected devices, ushering in an age of machine-to-machine transactions.

Third, and finally, blockchain technology supports resilient, transparent, nonrepudiable, and tamper-resistant registries. Blockchains are storing important records in a sequential, time-stamped manner, by known and authenticated parties, which are accessible (and auditable) by anyone with an Internet connection. These records include title to land or other kinds of property and public sector information.

At the same time, blockchain-based systems are being developed to operate outside of the legal system, ignoring long-standing restrictions placed on existing markets and financial institutions. Decentralized digital currencies are being used to launder money and avoid financial regulations. Blockchains and smart contracts are powering gambling dens and decentralized marketplaces where people can trade counterfeit or illegal goods. They underpin decentralized exchanges that facilitate the transfer of millions of dollars' worth of digital tokens (some of which resemble securities) and unregistered options, as well as file-sharing systems and communication and social media networks that do not fit squarely with copyright laws and regulators preventing the dissemination of obscene or illicit material.

Blockchains and the Layers of the Internet

The immense potential of blockchains has led to proclamations that the technology is as important as—or may even replace—the Internet.[37] However, such statements are misguided. These decentralized databases piggyback on existing Internet technologies, allowing people to develop new application protocols and higher-level services with all or some of a blockchain's distinctive characteristics.

The Internet is made up of multiple protocols that, when combined, create different layers of communication.[38] Although there is no consensus as to what these layers might be, two models have acquired some recognition. The first is the OSI / ISO Basic Reference Model (or seven-layer networking model), elaborated in the 1980s as a result of international deliberation between large telecom operators and the International Standards Organization (ISO).[39] The second is the five-layers software model,[40] which is the product of a bottom-up process aimed at describing the role of existing Internet protocols. We present here—and expand on—the one that is the most useful for the purpose of this book, the five-layers model, also known as the TCP/IP model.

The TCP/IP Model

Under the TCP/IP model, the Internet is conceptualized as five separate, independent, and modular layers: the physical layer, the data link layer, the network layer, the transportation layer, and the application layer (see Figure 2.1).[41]

The application layer sits on top of the TCP/IP stack and consists of a set of protocols—such as HTTP, FTP, SMTP, and DNS—that enable people to share information, swap messages, transfer files, or resolve domain names into their corresponding IP addresses. These protocols underpin a variety of online services that people interact with on a daily basis.[42]

Underneath the application layer are both the network and transport layers. The network layer—governed by the Internet Protocol (IP)—is the "glue that holds the entire Internet together."[43] Computers connected via the Internet are assigned unique IP addresses to help packets of data navigate across the network, passing through a variety of computers until they reach the requested destination. The transport layer—primarily governed by the Transmission Control Protocol (TCP) and the User Datagram Protocol (UDP)—ensures that data packets sent through the IP layer are properly delivered, in the appropriate order.[44] While the IP layer governs the delivery of data packets, the transport layer is concerned with data's fragmentation and reassembly. Data received via the application layer protocols is first broken down into smaller packets (via the transport layer), which are then passed along to the IP layer to reach their destinations. Once a recipient receives the requested packets, the transport layer reassembles them into the right order and sends the reassembled data back to the application protocol.[45]

Below the IP and transport layers is the data link layer, which comprises all protocols (Ethernet, ATM, 802.11 protocols supporting WiFi systems) that interface with hardware connected to the Internet. This layer ensures that the Internet operates independently of any specialized hardware so that the Internet can evolve over time.[46]

The bottom layer of this model is the physical layer, namely the pipes and tubes of the Internet. This layer comprises all pieces of hardware necessary

Layers	Protocols
Application	HTTP (web), SMTP (email), FTP (file transfer)
Transport	TCP, UDP
Network	Internet Protocol
Link	Ethernet, ATM, 802.11 protocols for Wifi
Physical	Cable modems and satellite links

FIGURE 2.1 The TCP/IP model

for machines to transfer and receive information from the Internet—things like DSL and cable modems, T1 connections, and satellite links.[47]

How Blockchains Fit within the TCP/IP Model

Blockchain technology supports a range of application protocols capable of not just transmitting bits of information but also storing information and executing computational processes in a way that does not rely on any centralized operator. Protocols like Bitcoin ultimately rely on TCP/IP to operate[48] and can be viewed as new application protocols that sit on top of the transport layer (see Figure 2.2).

Traditional application protocols facilitate the transmission of data over the Internet, assuming that there would be a centralized party (acting as a server) hosting data, as well as individual users (each acting as a client). For example, the HTTP protocol facilitates requests to a web server that, once received, sends back information such as web pages and images to a requesting user. In much the same way, the SMTP protocol relies on a mail server to send messages back and forth between Internet users.[49] Newer protocols supporting peer-to-peer networks were designed to deliver information with less of a need to rely on centralized servers. Protocols like BitTorrent use TCP/IP, as well as centralized trackers or distributed hash tables (DHTs), to facilitate the exchange of data packets between distributed networks of peers.[50]

Blockchain-based application protocols work like the BitTorrent protocol in many ways, although they do not rely on centralized trackers or distributed

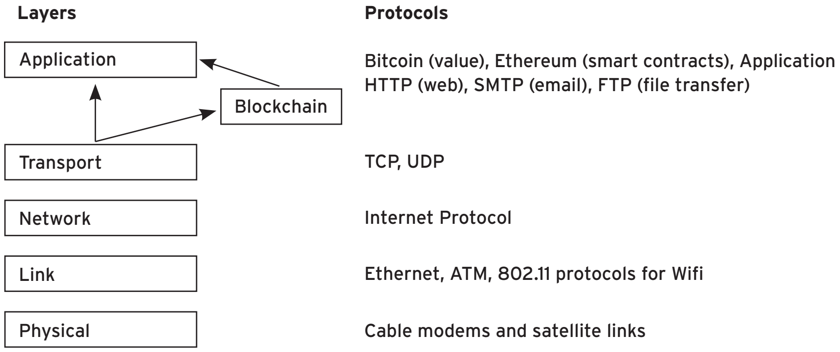
Layers
Protocols
Application
Blockchain
Bitcoin (value), Ethereum (smart contracts), Application
HTTP (web), SMTP (email), FTP (file transfer)
Transport
TCP, UDP
Network
Internet Protocol
Link
Ethernet, ATM, 802.11 protocols for Wifi
Physical
Cable modems and satellite links

FIGURE 2.2 Blockchains within the TCP/IP layers

hash tables to coordinate activity on the network. Instead, these protocols use a blockchain, a consensus mechanism, and a decentralized virtual machine to manage, validate data, and run computations on a peer-to-peer network. By implication, these protocols perform a greater range of functions necessary to build and support robust online services, including storing information and running code that can operate independently of any single party.

As with HTTP, online services are also being built on top of blockchain-based networks like Ethereum and Bitcoin, exhibiting degrees of centralization that fall on a spectrum. Some services operate in a fully autonomous manner, using the underlying blockchain or other overlay networks to store data, transfer value, and run computer processes. Other services only rely on a blockchain for one or more of these functionalities and leverage intermediaries or other centralized services delivered via HTTP to handle other essential functions. There are also centralized services that interact with a blockchain but operate independently from the blockchain's underlying peer-to-peer network. For instance, this is how centralized digital currency wallet providers are structured. They access information and interact with a blockchain-based network but do not require the services of the underlying peer-to-peer network for storage or computational power.

Lex Cryptographica

Blockchain-based application protocols and services that rely on blockchain technology to operate autonomously hold out the potential to create tensions with existing laws and regulations. These protocols and services have the capacity to implement their own systems of rules—*lex cryptographica*—which are enforced by the underlying protocol and smart contracts running on a blockchain-based network.

When the Internet first emerged, legal scholars David Johnson and David Post looked at this new distributed communication network and proclaimed that the Internet "undermin[ed] the feasibility—and legitimacy—of laws based on geographic boundaries."[51] No longer would the world be governed by nation-states, and no longer would governments be able to enact laws to establish fundamental rights, shape markets, or manage social interactions;[52] rather, national laws and regulations would dissipate into the bits and bytes of "cyberspace," replaced by rules defined by private actors.

Post and Johnson argued that laws were, as a general rule, inherently territorial and could operate only within specific geographic boundaries.[53] Governments were not omnipotent but only had the power to enact and enforce laws on individuals living, transacting, or otherwise operating within their jurisdictions.[54] Because the Internet was not governed by any single actor, and because it cut across multiple jurisdictions—creating a shifting and uncertain landscape of various laws, changing national rules, and conflicting regulations—they feared services accessed via the Internet would undermine the ability of governments to shape human behavior.[55] National courts would lack jurisdiction over potential wrongdoers, and states would become largely incapable of remedying online harms.[56]

Ultimately, however, these early prognostications about the unregulatability of the Internet were found to be overly broad. As Jack Goldsmith and Timothy Wu recognized, the Internet could be tamed because it was not fully distributed but rather had points of control that could shape and influence online activity. Online operators, located in physical space, operated within the jurisdiction of a state, and therefore could be required to comply with national laws.[57] By applying coercive force on these intermediaries, national governments could curtail the anarchic potential of the Internet and bring order to the online world.[58]

Over the past twenty years, the Internet has transformed from a digital "Wild West" to an increasingly regulated medium where large online operators bear the responsibility of abiding by and applying national laws.[59] Today, an increasing number of regulations apply to online service providers. Governments routinely pass laws forcing ISPs to filter Internet communications and block certain sites that violate the law, such as those hosting copyright-infringing material or child pornography.[60] Governments impose obligations on payment processors to prevent online gambling and deprive illegal services of accumulating revenue.[61] They impose rules on information intermediaries such as Google and Facebook, requiring that they police their services, report criminal activity, remove links to information that could potentially invade privacy rights, and censor attempts at cyberbullying.[62] Governments also carefully define what large e-commerce platforms can buy or sell, preventing these platforms from dealing in drugs, illegal arms, and other harmful or offensive items.[63] This list does not include the approaches adopted by countries like China and Russia, which impose extensive regulations on ISPs and other online operators to directly control the types of information that their citizens can view.[64]

And yet Johnson and Post had a point. While it was possible for governments to control certain aspects of the Internet, online activities could never be entirely constrained. With some research and technical know-how, citizens could always find ways to avoid certain rules by using services operating outside national boundaries. In effect, because of the lack of an overarching regulatory authority, the Internet created a legal vacuum, where online operators started using technical constructs—the code of their platforms—to define rules that shape how people can act and interact online.[65] This is what led Lawrence Lessig to declare that on the Internet "code is law."[66] Code can serve as a "salient regulator" defining our human experience.[67] It can constrain or enable behavior in ways that differ from traditional, state-enacted laws.

In many ways, blockchains walk us back to Post and Johnson's initial vision, supporting code-based rules that operate transnationally and are difficult to regulate and control. By relying on the unique characteristics of a blockchain, people can build systems that operate autonomously, governed by *lex cryptographica,* and designed in such a way that they cannot be altered by any single party. These systems enable new forms of social interactions and commercial activity, with less of a need for centralized coordination. They can leave room for people to interact and coordinate, or they can be implemented as a set of rigid and static rules that establish what people can or cannot do, leaving virtually no room for human intervention.

These systems can be designed to undermine and erode existing social structures or enhance and protect them. Like all other pieces of software, the design of blockchain-based protocols and services reflects discrete choices, which are not free from bias, influence, or politics.[68] As with other technologies, blockchains and smart contracts are capable of both circumventing and complementing the law—depending on the developers' desired outcome.

When faced with blockchain-based autonomous systems, governments may struggle to ensure the proper application of the law, because online services relying on *lex cryptographica* differ from online services that depend on intermediaries. Today, intermediaries ultimately control the services they provide and retain the power to intervene and unilaterally alter the rules governing their platforms if so desired. Because intermediaries often are identifiable, governments can force them to shut down or modify their rules without impacting other online services.[69]

Systems deployed on a blockchain—especially those relying on *lex cryptographica*—are not subject to the same kinds of limitations. By relying on decentralized peer-to-peer networks, blockchain-based systems can be designed to operate autonomously and potentially independent of the whims of centralized intermediaries by implementing code-based rules that are more persistent and often harder to change than those deployed by traditional centralized operators. As Michael Abramowicz describes, these blockchain-based systems can "serve as the foundation for more sophisticated types of decision making, allowing legal institutions to be created without voting or the designation of a central authority."[70] While there are ways to regulate these applications, the mechanisms to do so require controlling the way a blockchain-based protocol operates or regulating intermediaries operating at lower levels of the Internet stack.

In effect, with *lex cryptographica,* national laws get pushed to the edges. Individuals decide whether to interact with these autonomous systems, frustrating legal regimes focused on implementing rules on central parties that currently control or help facilitate online activity. If blockchain-based autonomous systems become increasingly used to provide online services, governments will need to adopt new techniques and approaches to shape or regulate these services. Traditional legal doctrines, especially those focused on regulating middlemen, will not easily translate to these new decentralized and autonomous systems, and the broader adoption of blockchain technologies may ultimately require the development of alternative mechanisms of regulation that better account for the distinctive characteristics of *lex cryptographica*.

Protocols and Power

If blockchain development continues apace, *lex cryptographica* could increasingly dictate and seep into our everyday lives, potentially affecting a greater range of online interactions. As more and more online platforms rely on blockchain technology, the power that these protocols exert over individuals—and society more broadly—will not evaporate; rather, power will shift to the code and programmers supporting these systems.

Before the advent of the Internet, rules were imposed by governments and public institutions through a hierarchical and bureaucratic model. Governmental authorities served as centralized points of control, delegating power to agencies, organizations, or individuals acting on behalf of higher-level officials. These bureaucratic organizations—as Max Weber described them—operated according to specific rules that constrained the discretionary power of governments and public administrations with written laws and regulations.[71]

Michel Foucault termed these societies "disciplinary societies,"[72] societies that control and shape the behavior of individuals by regulating the institutions around them—including schools, universities, factories, hospitals, asylums, and prisons. These "disciplinary institutions" ensured that every citizen would respect established rules and laws[73] by employing an elaborate system of checks and balances that required significant governmental oversight and surveillance.[74] Institutions, however, only had a limited ability to control the behavior of citizens, in that there was at that time a discernible distinction between the public sphere, ruled by bureaucratic rules, and the private sphere, which largely escaped the control of governmental institutions.[75]

As the Internet and digital technologies have continued to expand and mature, they have begun to shift society away from "disciplinary societies" toward what Gilles Deleuze has termed a "society of control."[76] In this new society, individuals are free to establish their own courses of action, with fewer constraints by previous forms of institutional enclosures and in ways that are less dependent on disciplinary institutions. For example, people can attend online classes to receive their education without the need for a physical university or school. They can work from anywhere around the globe on a piecemeal basis without depending on employment from a single factory or employer. With the advent of body sensors and self-measurement devices, people can even perform some medical diagnoses themselves without the need to visit a hospital.

While individuals operating in a "society of control" appear to have less of a need to follow the rules and procedures of disciplinary institutions that they may interact with, they are now subject to a much broader and more subtle form of control over their activities: a diffuse system of information gathering and code-based protocols that shape and mold behavior.[77] The disciplinary society of Foucault—governed by strict rules and centrally administered regulations[78]—has begun to shift toward the society of control envisioned by Deleuze, governed by a much more flexible and malleable system of continuous control and ubiquitous surveillance, administered via technical protocols.[79]

These protocols, shaped by governments and private corporations, dictate what people can or cannot do on a given online platform, and because they are automatically applied by the underlying technology, they often are less dependent on disciplinary institutions for enforcement. In other words, disciplinary actions and ex-post mechanisms of punishment are being replaced by a system of ex-ante regulation and continuous control, enabling governments and private actors to influence the activities of individuals—both in the public and private spheres—to ensure that they comply with the law.[80]

The mainstream adoption of the Internet and the growing reliance on online services for everyday tasks have facilitated the shift toward a society of control. Because most actions on the Internet leave a trace, governments and private institutions can increasingly shape what people do online and assess individuals' compliance with the rules of the platform and, in turn, the law.

Online services can deploy algorithms to shape human behavior. Google's search algorithm and Facebook's news feed algorithm spread and frame information in ways that influence individual decision making.[81] Algorithms trade stocks on Wall Street, identify potential tax evasion or other suspicious activities, assist doctors in the diagnosis of diseases, and help researchers with scientific discovery. They even help us decide where to have dinner and who our life partner should be.[82]

However, we are just at the beginning. As Tarleton Gillespie recognized, "Algorithms are inert, meaningless machines until paired with databases upon which to function."[83] Once combined, these two layers—code and data storage—work in tandem to implement systems of control that dictate what people can or cannot do online.

Blockchains are therefore a particularly potent new technology when it comes to algorithmic systems, because they integrate both a storage and a computation layer in a seamless and often indistinguishable manner. Blockchains enable parties to coordinate activity in an automated and decentralized way, and are viewed as a new technology that transforms pillars of industrialized society into entirely or primarily code-based systems. With blockchains, payment systems, financial markets, information systems, and—more generally—the allocation of labor between people and machines can be governed by technical rules.

The maturation and widespread deployment of the technology could therefore accelerate a shift of power from legal rules and regulations to software protocols and other code-based systems.[84] Such a shift would have an important effect on our daily lives: blockchain-based systems and *lex cryptographica* would mold social, economic, legal, and political interactions; they would help us transfer value, protect assets, administer organizations, and validate meaningful life and cultural events. The design of blockchain-based protocols and *lex cryptographica*—and decisions related to their development—would ultimately dictate how these systems work and shape our means of interaction. Existing bureaucratic systems, operated by people and institutions abiding by the rule of law, would be replaced by technocratic systems, operated by technical structures and code-based rules that ultimately constrain human behavior and discretionary choice. Algorithms would define the range of possible actions that individuals may or may not take, to the detriment of potentially valuable alternatives.

The focal point of power in many of these systems, however, would no longer be centralized institutions and hierarchical structures but rather informal systems of (often invisible) rules dictated by programmers deploying code. As a result, the growing reliance on algorithms to shape our interactions with one another and with third-party operators would increasingly subject us to the "rule of code" as opposed to the "rule of law"—eventually placing us in an algocracy.[85]

Today, algorithms are centrally controlled, deployed and stewarded by online intermediaries, which (at least until the development of more emergent artificial intelligence) retain control over these algorithms and the power to tweak them or to shut them off if necessary. Blockchains change this. The "rule of code" established by *lex cryptographica* can be designed to be harder to control and could be used to enable individuals to self-govern and deviate from long-standing legal rules.

The Challenges of Blockchain Technology

At least for the short term, the risks of blockchains and *lex cryptographica* are tempered both by structural problems with blockchain-based networks and by current technical limitations of blockchains (which many are working to surmount). Perhaps the most significant challenge of blockchain-based networks relates to the issues of scalability and security. There are legitimate questions as to whether blockchains are capable of scaling and whether they are secure enough to safely manage the comprehensive and global systems described throughout this book.

Existing blockchains are not as powerful and fast as other data management technologies. These current decentralized networks only handle comparatively few transactions. For instance, the Bitcoin blockchain processes roughly 240,000 transactions per day—far less than the trillions of messages sent across the Internet or the 150 million daily transactions handled by credit card companies such as Visa.[86] What's more, it takes approximately ten minutes for a Bitcoin transaction to be validated by the network and recorded to the shared data set, in contrast to the fraction of a second it typically takes a database to store and record information.[87]

For blockchain technology to achieve mainstream adoption, these emerging technologies will need to handle a seemingly countless number of transactions. The speed and trustworthiness of these networks will likewise need to grow for private and public entities to leverage this technology for the development of novel applications and innovative business models.

Solving scalability issues is no simple task. Because blockchains are append-only databases, each new transaction on the network causes the blockchain to grow. The larger the blockchain, the greater its requirements are in terms of storage, bandwidth, and computational power.[88] If these requirements become too onerous, fewer individuals or entities will be able to invest resources to maintain the shared database, weakening the security of the blockchain by making it easier for a small number of large mining pools to take over the network and potentially compromise its contents.[89]

While there are already a few proposals on how to make blockchains scale in a secure manner—for example, by moving certain transactions off a blockchain, developing faster consensus protocols, or dividing the shared database in ways that would enable a network to process transactions in parallel—these solutions have yet to be implemented in earnest.[90] Whether they materialize will determine the future viability of blockchain-based networks.

Beyond issues related to security and scalability, and despite the autonomous nature of *lex cryptographica,* most blockchain-based networks still remain susceptible to governmental regulations that may either support or hinder the development of blockchain technology. While there are millions of Bitcoin and Ethereum accounts, with thousands of developers worldwide exploring the possible uses of this emergent technology, even the largest blockchain networks are still miles away from gaining the same level of adoption as the World Wide Web, email, or other Internet-based protocols.

Because of the nascent nature of blockchains, governments retain the ability to shape the development of the technology by passing laws and regulations that will either constrain or promote the technology's growth and adoption. Regulations could stymie the development of blockchain technology by making it expensive or difficult to operate digital currencies or deploy autonomous smart contract code. Conversely, governments could implement favorable regulatory frameworks to protect businesses experimenting with blockchains as part of pro-innovation policies.[91]

However, regulation creates its own set of problems. Regulating too soon could provide valuable guidance as to the legitimate uses of blockchain technology but could also stamp out potential benefits.[92] Regulating too late may dissuade the most risk-averse actors from exploring blockchains because of legal uncertainty while simultaneously allowing socially objectionable aspects of the technology to emerge.

Therefore, the first step toward understanding how to regulate blockchain technology requires an analysis of its emerging uses, along with a more detailed examination of the technology's benefits and drawbacks. We will start with the impact of blockchain technology on financial and legal systems and then move on to explore how blockchains shape our interactions with information, organizations, and ultimately machines.