

Blockchain: What It Is, What It Does, and Why You Probably Don't Need One

David Andolfatto

All record-keeping systems (which include monetary systems) must contend with trust issues and methods of organizing historical information. Conventional systems rely on the reputation of central authorities and record-keepers to achieve consensus. Blockchain, which powers Bitcoin, differs from conventional systems by achieving consensus through a community of anonymous (and therefore “trustless”) agents who compete amongst themselves to authenticate transactions. The promise of the blockchain protocol is that it is invulnerable to human foibles. Novel, for sure; but is it worth all the effort? (JEL G23, E50, E59)

Federal Reserve Bank of St. Louis *Review*, Second Quarter 2018, 100(2), pp. 87-95.
<https://doi.org/10.20955/r.2018.87-95>

Public interest in *blockchain* has never been higher. This is in large part due to the eye-popping price dynamics of *Bitcoin*—the original bad-boy cryptocurrency—which everyone hears is powered by blockchain. But what, exactly, is blockchain? The answer, it seems, depends on who you ask.

Things are confusing out there in part because not enough care is taken in defining terms before discussing the subject. And when terms are defined, they sometimes include *desired outcomes* as a part of their definition. For example, blockchain is often described as consisting of (among other things) an *immutable ledger*. To me, this is like defining a *titanic* to be an *unsinkable ship*.

So what do people mean when they bandy about the term blockchain? I recently had a chance to learn how blockchain is viewed from a corporate perspective at a blockchain panel recently hosted by the Olin School of Business at Washington University in St. Louis. My discussion here is based on a presentation I delivered there,¹ and it complements a recent article published in this *Review* (Berentsen and Schär, 2018). My co-panelist at that event, Ed Corno of IBM, provided the following definition:

Blockchain: a shared, replicated, permissioned ledger with consensus, provenance, immutability, and finality.

David Andolfatto is a vice president and economist at the Federal Reserve Bank of St. Louis.

© 2018, Federal Reserve Bank of St. Louis. The views expressed in this article are those of the author(s) and do not necessarily reflect the views of the Federal Reserve System, the Board of Governors, or the regional Federal Reserve Banks. Articles may be reprinted, reproduced, published, distributed, displayed, and transmitted in their entirety if copyright notice, author name(s), and full citation are included. Abstracts, synopses, and other derivative works may be made only with prior written permission of the Federal Reserve Bank of St. Louis.

To put things another way, blockchain is a record-keeping system possessing a specific set of attributes. But which of the attributes listed above distinguish it from other record-keeping systems? Shared, replicated ledgers requiring permissioned access have been around for a long time, so these are not distinguishing characteristics. Provenance simply means a ledger containing all relevant information, starting from the beginning of any relationship. Provenance is nothing new, at least in principle.

The key would seem to lie in the nature of consensus. Consensus means that all relevant parties agree that the information contained in the ledger is true. Depending on how one defines “all relevant parties,” consensus in itself is not new either. What is new is the consensus mechanism—the protocol designed to achieve consensus. Immutability and finality are simply properties that one hopes will be the outcome of the consensus mechanism in place.

Conventional protocols for achieving consensus can be thought of as reputation-based mechanisms. The solution to the record-keeping problem is to delegate the responsibility to a trusted intermediary (or a set of trusted intermediaries). “Trust” in this context means believing that personal and business reputations have too much long-run economic value to be squandered by exploiting an attractive one-time gain.

The promise of blockchain is to replace reputation-based consensus with a “trustless” protocol invulnerable to human foibles. In particular, the reputable delegated record-keeper is replaced by a set of anonymous (hence, untrusted) agents drawn from the broader community who compete amongst themselves by playing a game on a period-by-period basis to update and secure information. In this sense, blockchain can be thought of as a game-based consensus mechanism.

I can’t be entirely sure, but I believe the corporate versions of blockchain are likely to stick to the standard model of reputation-based consensus. If this is correct, then the efficiency gains of “blockchain” boil down to the gains associated with making databases more synchronized across trading partners, more cryptographically secure, more visible, more complete, etc. In short, there is nothing revolutionary or radical going on here—it’s just the usual advancement of the technology and methods associated with the on-going problem of database management. Labeling the endeavor blockchain in this case has more to do with good marketing practices.

On the other hand, game-based blockchains—such as the one that powers Bitcoin—are, in my view, potentially more revolutionary. But before I explain why I think this, I want to step back a bit and describe my bird’s eye view of what’s happening in this space.

A DATABASE OF INDIVIDUAL ACTION HISTORIES

The type of information that concerns us here is not what one might label “knowledge,” say, as in the recipe for chicken cacciatore. The information in question relates more to a set of events that have happened in the past—in particular, events relating to individual actions. Consider, for example, the statement “David washed your car two days ago.” This type of information is intrinsically useless in the sense that it is not usable in any productive manner. In addition to work histories like this, the same is true of customer service histories, delivery/

receipt histories, credit histories, or any performance-related history. And yet, people value such information. It forms the bedrock of reputation and perhaps even of identity. As such, it is frequently used as a form of currency.

Why is intrinsically useless history of this form valued? A monetary theorist may tell you it's because of a lack of commitment or a lack of trust. (See the *MacroMania* blog post “Evil is the root of all money.”) If people could be relied upon to make good on their promises a priori, their track records would largely be irrelevant from an economic perspective. A good reputation is a form of capital. It is valued because it persuades creditors (believers) that more reputable agencies are more likely to make good on their promises. We keep our money in a bank not because we think bankers are angels, but because we believe the long-term franchise value of banking exceeds the short-run benefit a bank would derive from appropriating our funds. (Well, that's the theory, at least. Admittedly, it doesn't work perfectly.)

Note something important here. Because histories are just information, they can be created “out of thin air.” And, indeed, this is the fundamental source of the problem: People have an incentive to fabricate or counterfeit individual histories (their own and perhaps those of others) for a personal gain that comes at the expense of the community. No society can thrive, let alone survive, if its members have to worry excessively about others taking credit for their own personal contributions to the broader community. I'm writing this article in part (well, perhaps mainly) because I'm hoping to get credit for it.

Since humans (like bankers) are not angels, what is wanted is an honest and immutable database of histories (defined over a set of actions that are relevant for the community in question). Its purpose is to eliminate false claims of sociable behavior (acts that are tantamount to counterfeiting currency). Imagine, too, eliminating the frustration of discordant records. How much time is wasted in trying to settle “he said/she said” claims inside and outside of courts of law? The ultimate goal, of course, is to promote fair and efficient outcomes. We may not want something that goes as far as Santa Claus—who, in addition to knowing whether you've been bad or good, also knows your sleep patterns—but something similar defined over a restricted domain for a given application could be useful.

ORGANIZING HISTORY

Let $e(t)$ denote a set of events or actions (relevant to the community in question) performed by an individual at date $t = 1, 2, 3, \dots$. An individual history at date t is denoted

$$h(t-1) = \{e(t-1), e(t-2), \dots, e(0)\}, t = 1, 2, 3, \dots$$

Aggregating over individual events, we can let $E(t)$ denote the set of individual actions at date t and let $H(t-1)$ denote the communal history—that is, the set of individual histories of people belonging to the community in question:

$$H(t-1) = \{E(t-1), E(t-2), \dots, E(0)\}, t = 1, 2, 3, \dots$$

Observe that $E(t)$ can be thought of as a “block” of information (relating to a set of actions taken by members of the community at date t). If this is so, then $H(t-1)$ consists of time-stamped

Andolfatto

blocks of information connected in sequence to form a chain of blocks. In this sense, any database consisting of a complete history of (community-relevant) events can be thought of as a “blockchain.”

Note that there are other ways of organizing history. For example, consider a cash-based economy where people are anonymous and let $e(t)$ denote acquisitions of cash (if positive) or expenditures of cash (if negative). Then an individual’s cash balances at the beginning of date t is given by $h(t-1) = e(t-1) + e(t-2) + \dots + e(0)$. This is the sense in which “money is memory.” Measuring a person’s worth by how much money they have serves as a crude summary statistic of the net contributions they’ve made to society in the past (assuming they did not steal or counterfeit the money, of course). Another way to organize history is to specify $h(t-1) = \{e(t-1)\}$. This is the “what have you done for me lately?” model of remembering favors. The possibilities are endless. But an essential component of blockchain is that it contains a complete history of all community-relevant events. (We could perhaps generalize to truncated histories if data storage is a problem.)

DATABASE MANAGEMENT SYSTEMS (DBMS) AND THE READ/WRITE PRIVILEGE

All right then, suppose that a given community (consisting of people, different divisions within a firm, different firms in a supply chain, etc.) wants to manage a chained-block of histories $H(t-1)$ over time. How is this to be done?


Along with a specification of what is to constitute the relevant information to be contained in the database, any DBMS will have to specify parameters restricting the following:

1. The read privilege (who, what, and how)
2. The write privilege (who, what, and how)

That is, who gets to read and write history? Is the database to be completely open, like a public library? Or will some information be held in locked vaults, accessible only with permission? And if by permission, how is this to be granted? By a trusted person, by algorithm, or by some other manner? Even more important is the question of who gets to write history. As I explained earlier, the possibility for manipulation along this dimension is immense. How does a community guard against attempts to fabricate history?

Historically, in “small” communities (think traditional hunter-gatherer societies) this was accomplished more or less automatically. There are no strangers in a small, isolated village and communal monitoring is relatively easy. Brave deeds and foul acts alike, unobserved by some or even most, rapidly become common knowledge. This is true even of the small communities we belong to today (at work, in clubs, families, friends, etc.). Kocherlakota (1996) labels $H(t-1)$ in this scenario “societal memory.” I like to think of it as a virtual database of individual histories living in a distributed ledger of brains talking to each other in a P2P fashion, with additions to, and maintenance of, the shared history determined through a consensus mechanism. In this primitive DBMS, read and write privileges are largely open, with the latter being subject to consensus. It all sounds so...blockchainy.

Figure 1



SOURCE: Derived from work by Ed Corno of IBM.


While the primitive “blockchain” described above works well enough for small societies, it doesn’t scale very well. Today, the traditional local networks of human brains have been augmented (and to some extent replaced) by local and global networks of computers capable of communicating over the Internet. Achieving rapid consensus in a large heterogeneous community characterized by vast flows of information is a rather daunting task.

The “solution” to this problem has largely taken the form of proprietary databases with highly restricted read privileges managed by trusted entities who are delegated the write privilege. The double-spend problem for digital money, for example, is solved by delegating the record-keeping task to a bank, located within a banking system, performing debit/credit operations on a set of proprietary ledgers connected to a central hub (a clearing agency) typically managed by a central bank.

THE PROBLEM AND THE BLOCKCHAIN SOLUTION


Depending on your perspective, the system that has evolved to date is either a great improvement over how things operated when you were young (if you were born before 1980) or a hopelessly tangled hodgepodge of networks that have trouble communicating with each other and are intolerably vulnerable to data breaches (if you were born after 1980). The figures, derived from work by Ed Corno of IBM, illustrate this general type of network as well as a blockchain-based network and a traditional client-server-based network.

Figure 2



SOURCE: Derived from work by Ed Corno of IBM.

Figure 3



SOURCE: Derived from work by Ed Corno of IBM.

The solution to this present, convoluted state of affairs is shown in Figure 2 as blockchain. And, sure, this looks like a more organized way to keep the books and clear up communication channels, though the details concerning how consensus is achieved in this system remain a little hazy to me. As I mentioned earlier, I'm guessing that it'll be based on some reputation-based mechanism. But if this is the case, then why can't we depict the solution in the following way?

That is, gather all the agents and agencies interacting with each other, forming them into a more organized community, but keep it based on the traditional client-server (or hub-and-spoke) model. In the center, we have the set of trusted "historians" (bankers, accountants, auditors, database managers, etc.) who are granted the write privilege. Communications between members may either be intermediated by historians or take place in a P2P manner with the historians listening in. The database can consist of the chain-blocked sets of information (that is, the blockchain $H(t-1)$ described above). The parameters governing the read privilege can be determined beforehand by the needs of the community. The database could be made completely open, which is equivalent to rendering it shared. And, of course, multiple copies of the database can be made as often as is deemed necessary.

The point I'm making is that, if we're ultimately going to depend on reputation-based consensus mechanisms, then we need no new innovation (like blockchain) to organize a database. While I'm no expert in the field of database management, it seems to me that standard protocols, for example, in the form of SQL Server 2017, can accommodate what is needed technologically and operationally.

EXTENDING THE WRITE PRIVILEGE: GAME-BASED CONSENSUS

As explained above, extending the read privilege is not a problem technologically. We are all free to publish our diaries online, creating a shared/distributed ledger of our innermost thoughts. Extending the write privilege to unknown or untrusted parties, however, is an entirely different matter. Of course, this depends in part on the nature of the information to be stored. Wikipedia seems to work tolerably well. But it's hard to use Wikipedia as currency. This is not the case with personal action histories. You don't want other people writing your diary!

Well, fine, so you don't trust "the Man." What now? One alternative is to game the write privilege. The idea is to replace the trusted historian with a set of delegates drawn from the community (a set potentially consisting of the entire community). Next, have these delegates play a validation/consensus game designed in such a way that the equilibrium (say, Nash or some other solution concept) strategy profile chosen by each delegate at every date $t = 1, 2, 3, \dots$ entails (i) no tampering with recorded history $H(t-1)$ and (ii) only true blocks $E(t)$ are validated and appended to the ledger $H(t-1)$.

What we have done here is replace one type of faith with another. Instead of having faith in mechanisms that rely on personal reputations, we must now trust that the mechanism governing noncooperative play in the validation/consensus game will deliver a unique equilibrium outcome with the desired properties. I think this is in part what people mean when I hear them say "trust the math."

Andolfatto

Well, trusting the math is one thing. Trusting in the outcome of a noncooperative game is quite another matter. The relevant field in economics is called mechanism design. I'm not going to get into details here, but suffice it to say that it's not so straightforward designing mechanisms with surefire beneficial properties. Ironically, mechanisms such as Bitcoin will have to build up trust the old-fashioned way—through positive user experience, much the same way most of us trust our vehicles to function, even if we have little idea how an internal combustion engine works.

Of course, the same holds true for games based on reputational mechanisms. The difference is, I think, that noncooperative consensus games are intrinsically more costly to operate than their reputational counterparts. The proof-of-work game played by Bitcoin miners, for example, is made *intentionally* costly (to prevent DDoS attacks) even though validating the relevant transaction information is virtually costless if left in the hands of a trusted validator. And if a lack of transparency is the problem for trusted systems, this conceptually separate issue can be dealt with by extending the read privilege communally.

Having said this, I think that, depending on the circumstances and the application, the cost associated with a game-based consensus mechanism may be worth incurring. I am inclined to remain agnostic on this matter for now and see how future developments unfold.

BLOCKCHAIN: POWERING DAOs

If blockchain (with noncooperative consensus) has a comparative advantage, where might that advantage be? To me, the clear application is in supporting decentralized autonomous organizations (DAOs). A DAO is basically a set of rules written as a computer program. Because it possesses no central authority or node, it can offer tailor-made “legal” systems unencumbered by prevailing laws and regulations, at least insofar as transactions are limited to virtual fulfillments (e.g., debit/credit operations on a ledger).

Bitcoin is an example of a DAO, though the intermediaries that are associated with Bitcoin obviously are not. Ethereum is a platform that permits the construction of more sophisticated DAOs via the use of smart contracts. The comparative advantages of DAOs are that they permit (i) a higher degree of anonymity, (ii) permissionless access and use, and (iii) commitment to contractual terms (smart contracts).

It's not immediately clear to me what value these comparative advantages have for registered businesses. There may be a role for legally compliant smart contracts (a tricky business for international transactions). But perhaps the potential is much more than I can presently imagine. Time will tell. ■

NOTE

¹ The presentation includes a slide deck (<https://www.slideshare.net/DavidAndolfatto/blockchain-what-it-is-what-it-does-and-why-you-probably-dont-need-one>) and a video (<https://www.youtube.com/watch?v=fWozJbD-QwqQ&feature=youtu.be>). See also David Andolfatto's collected work on Bitcoin and blockchain (<http://andolfatto.blogspot.com/2017/12/my-perspective-on-bitcoin-project.html>).

REFERENCES

Andolfatto, David. "Evil is the root of all money." *MacroMania* blog post, 2012; <http://andolfatto.blogspot.com/2012/09/evil-is-root-of-all-money.html>.

Berentsen, Aleksander, and Schär, Fabian. "A Short Introduction to the World of Cryptocurrencies." Federal Reserve Bank of St. Louis *Review*, First Quarter 2018, pp. 1-16; <https://doi.org/10.20955/r.2018.1-16>.

Kocherlakota, Narayana R. "Money Is Memory." Federal Reserve Bank of Minneapolis Research Department Staff Report No. 218, October 1996; <https://minneapolisfed.org/research/sr/sr218.pdf>.

