

Blockchain Technology for Traditional Finance

ERIC BUDISH AND ADI SUNDERAM¹

March 2026

Abstract

Blockchain technology as embodied in cryptocurrencies like Bitcoin and Ethereum is comprised of both (a) a novel data structure and (b) a novel trust model. This paper analyzes the potential gains for traditional finance from an idealized version of the novel data structure on its own, with trust instead anchored in traditional sources such as rule of law, reputations, relationships, and collateral. Our framework has two parts. First, we analyze potential improvements for financial transactions that are already taking place. We identify three categories of improvement: (i) reducing real resource costs, (ii) improving balance sheet efficiency, and (iii) reducing intermediation rents. While the value of such improvements is hard to quantify precisely, we estimate that potential gains could be significant, especially the reduction in rents. Second, we analyze the potential for the technology to facilitate new transactions. We identify three channels: (i) making it more technologically difficult to cheat, (ii) making it easier to punish a cheating counterparty in a static sense, and (iii) making it easier to punish a cheating counterparty in a dynamic sense. Our key insight is that the potential gains are large if and only if there is a long tail of relatively low surplus, relatively infrequent transactions for which traditional forms of trust are insufficient. Last, we apply our framework to stablecoins. We conclude that if there are large gains from stablecoins for legal actors they are most likely to come from stablecoins putting pressure on intermediation rents or inefficient regulation, or from the programmability of stablecoins facilitating a large number of small transactions that otherwise would not have been trustworthy.

¹ Chicago Booth, eric.budish@chicagobooth.edu; Harvard Business School, asunderam@hbs.edu. This paper was originally prepared for the 2023 Macroeprudential Conference hosted by the Sveriges Riksbank. We thank our two discussants, Markus Brunnermeier and Steven Maijor, our colleagues Eugene Fama, Sam Hanson, Anil Kashyap, Jacob Leshno, Shengwu Li and Jeremy Stein, several financial-sector executives, and conference participants for their helpful comments. We thank George Rao for outstanding research assistance. The authors declare that they have no conflicts of interest that relate to this research.

1. Introduction

We think it is fair to say that technology in the modern financial system is uneven. On the one hand, the largest financial institutions collectively spend hundreds of billions of dollars per year on technology, the financial system processes literally quadrillions of dollars of transactions each year, and trading times are commonly measured in millionths and even billionths of seconds. On the other hand ... have you ever sent a wire? In many markets, settlement times are still days or even weeks. In 2021 Citibank famously sent \$900 million to Revlon creditors *by accident*. Financial columnist Matt Levine said of the case “I was treated to a gothic horror story about software design. I had nightmares all night about checking the wrong boxes on the computer.” Financial exchange computer glitches caused flash crashes in the U.S. stock market in May 2010 and in the U.S. Treasury market in October 2014. Lehman Brothers, on the brink of bankruptcy in early September 2008, literally could not calculate its exposure to U.S. real estate prices. Larry Fink, in his April 2025 Chairman’s letter to Blackrock shareholders, writes “The world’s money moves through plumbing built when trading floors still shouted orders and fax machines felt revolutionary.”¹

Costs are also uneven in the modern financial system. The fee to trade stock on the largest U.S. stock exchanges is about \$0.0001 per share per side—that is just 0.01 basis points on a \$100 stock.² Fees for large currency and bond transactions can be smaller still. But other costs in finance are much higher, such as credit card fees and investment management. The total cost of financial intermediation per dollar of assets has remained roughly constant at 2% per year for over one hundred years. The financial sector’s share of GDP rose from 4.8% in 1980 to 7.3% in 2023. Revenues from intermediating global payments are over \$2.5 trillion per year and forecast to continue growing at 4% per year.³

¹ Sources for this paragraph: Forrester Research (2026); Section 2.3 of this paper; Aquilina, Budish and O’Neill (2022); MacKenzie (2021); Levine (2021); CFTC and SEC (2010); U.S. Department of the Treasury et al. (2015); Valukas (2010); Geithner (2014); Fink (2025).

² For comparison, Uber’s 2024 take rate was 2700 basis points; or over 100,000x higher than stock exchange take rates on a percentage basis. (Calculations based on figures in the Uber 2025 10-K.)

³ Sources for this paragraph: Budish, Lee and Shim (2024); Greenwood and Scharfstein (2013); Greenwood, Ialenti and Scharfstein (2025); Philippon (2015); McKinsey (2025).

Could blockchain improve the financial system’s technology and lower its costs? That is, could blockchain technology generate value for end users of the financial system? There certainly is a lot of hype around this possibility. One can get a sense of the current public discussion from documents such as the President’s Working Group’s July 2025 report “Strengthening American Leadership in Digital Financial Technology,” SEC Chair Paul Atkins’ July 2025 speech “American Leadership in the Digital Finance Revolution,” Blackrock executives Larry Fink and Rob Goldstein’s December 2025 invited article in *The Economist* “On How Tokenisation Could Transform Finance,” and the section of Fink’s 2025 annual letter to stockholders called “Tokenization is Democratization.”⁴

We start with the following crucial distinction: blockchain technology is both (a) a novel data structure, and (b) a novel trust model. The data structure is essentially an append-only distributed database, which uses cryptography to enforce rules about which parties can perform which actions and uses smart contracts to automate certain sequences of actions. The trust model is an elaborate anonymous, decentralized majority voting system that maintains the database and adjudicates conflicts. Scientifically, the trust model was the significant breakthrough in Nakamoto (2008). Computer scientists use the phrase “permissionless consensus” for Nakamoto’s innovation, in contrast with “permissioned consensus” which had been well understood by computer science since the 1980s (e.g., Lamport, Shostak and Pease 1982; Dwork, Lynch and Stockmeyer 1988; see Lewis-Pye and Roughgarden 2023).

Excitement about blockchain technology often conflates these two features. However, as shown by Budish (2025), Nakamoto’s trust model is a non-starter for serious financial applications. The issue is how its costs scale with the economic stakes. For example, if the financial system relied on Nakamoto-style trust alone—without any support

⁴ Representative quotes: The President’s Working Group “endorses the notion that digital assets and blockchain technologies can revolutionize not just America’s financial system, but systems of ownership and governance economy-wide.” SEC Chair Atkins writes “The winds of innovation have always swept through our capital markets, often at gale force,” describes a future of “Big Beautiful On-Chain Software Systems” for U.S. securities markets, and writes “we will ensure that the next chapter of financial innovation is written right here in America.” Fink and Goldstein describe “the next major evolution in market infrastructure” and write “Tokenisation can modernise the infrastructure that still makes parts of the financial system slow and costly, bringing more people into the world’s most powerful engine of wealth creation: the markets.”

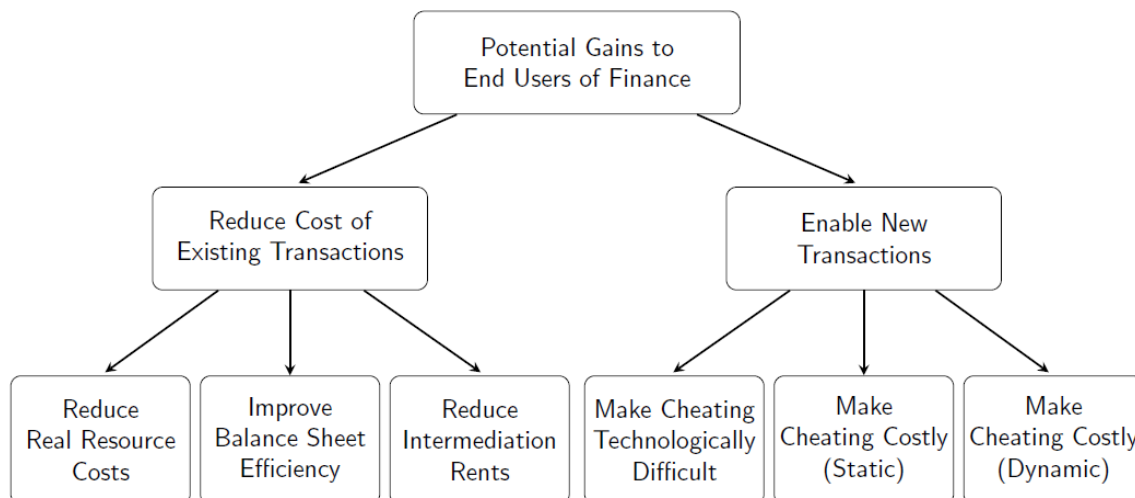
from rule of law, reputations, financial collateral, or other traditional sources of trust—Budish calculates that it would cost all of global GDP to prevent thefts of just \$40 billion. \$40 billion may sound like a lot but it is less than 10 minutes of financial system volume in the United States alone. As expensive as traditional finance is, it is a bargain relative to the cost of Nakamoto-style trust at scale. The core problem is that Nakamoto trust relies on decentralized, anonymous majority voting, and the votes can be bought by bad actors who can then record and confirm false transactions.⁵ The incentives to corrupt the majority vote in this way grow with the economic stakes ... and traditional finance, to quote the famous bank robber Willie Sutton, is “where the money is.”

The blockchain data structure, on the other hand—while prosaic relative to the radical trust model—has several features that seem useful for the traditional financial system. The ideal we have in mind in this paper is a perfectly user-friendly distributed database, with easy-to-use smart contracts (i.e., easy programmability of transactions), and know-your-customer and anti-money-laundering identity verification built in. Cryptographic tools such as signatures and zero knowledge proofs could enable anonymity where both useful and legal. We ask whether such an idealized version of the blockchain data structure could generate value for end users of the traditional financial system, if *trust* in the data is ultimately anchored in traditional sources including rule of law, reputations and financial collateral.

We divide our analysis framework into two parts (see Figure 1 for a schematic). First, we analyze potential gains from the idealized data structure for transactions that are already taking place in the traditional financial system. We highlight three categories of potential gains. First, better data structure could simply reduce the real resource costs of transactions, for instance by decreasing the number of people at financial institutions tasked with transactions clearing or reducing the cost of regulatory compliance. Second, by speeding up clearing it could reduce the amount of balance sheet capital that financial institutions need to dedicate to transactions in the clearing process. We emphasize in our

⁵ This is true of permissionless consensus in all its forms, whether proof-of-work, proof-of-stake, etc. See Budish (2025), Lewis-Pye and Roughgarden (2023), Budish, Lewis-Pye and Roughgarden (2025), and the sources cited therein.

**Figure 1. Analysis Framework:
Blockchain Technology for Traditional Finance.**



calculations that the extent to which this generates net efficiency gains depends directly on the extent to which there are violations of the Modigliani-Miller (1958) theorem. Third, we discuss the potential for new technology to lower total markups incurred by end users of financial intermediaries, i.e., to lower intermediation rents. We emphasize that technologies that improve intermediaries' efficiency are unlikely to reduce intermediaries' rents unless they also affect market structure, e.g., by enabling entry.

Second, we analyze potential gains from the idealized data structure for transactions that are newly enabled. We use a simple game theoretic framework that builds off of game forms used in Dixit (2004). In the game, there are two parties who can engage in mutually-beneficial transactions, but fear that the other party will cheat and steal their assets and that rule of law will fail to protect them. For example, the parties can engage in a \$100 transaction that generates \$5 of surplus for each party, but each party fears the other will just abscond with the \$100 and not follow through, and that it will cost more than \$100 to recover the stolen funds. A key insight from applying the folk theorem of repeated games is that if the parties are not already transacting today, the net present value of the surplus from their relationship (the NPV of the \$5 gains) must be small relative to the potential harm from being cheated (the \$100). The idealized data structure could help enhance trust

in such transactions in three ways. First, by simply making it more difficult to cheat, for instance with programmable escrow. Second, by making it easier for rule of law to punish a cheating counterparty in a static sense, for instance because of indisputable proof of malfeasance. Third, by making it easier for potential trading partners to punish a cheating counterparty in a dynamic sense, for instance, with new abilities to build, and also lose, reputations for trustworthiness.

The last part of our paper applies this framework to a particular application of blockchain technology: stablecoins. Stablecoins are digital currencies recorded on distributed ledgers that are pegged to a reference value, such as the U.S. Dollar. The largest stablecoins at present, Tether and USDC, run on permissionless blockchains that are vulnerable to the Budish (2025) critique. These stablecoins are likely to continue to be popular with users who need the anonymity (illegal activity, tax avoidance, regulatory avoidance, etc.),⁶ but are dangerous for the traditional financial system if used at large scale because of the anonymous majority voting. We ask whether stablecoins that run on an idealized permissioned blockchain protected by rule of law can create value.

We go through each element of our framework in turn. We find that stablecoins are unlikely to reduce real resource costs or improve balance sheet efficiency.⁷ We find that where stablecoins might create significant value is by reducing intermediation rents or facilitating new transactions. As an example of reducing intermediation rents, we note the interest from Amazon and Walmart in creating their own stablecoins, primarily to circumvent credit card fees, or the possibility that a Central Bank Digital Currency (CBDC) could provide services similar to a checking account but with higher interest payments to consumers. Both kinds of potential innovation should be interpreted as entry that reduces

⁶ We note that while the 2025 GENIUS act that legalized stablecoins requires account verification for know-your-customer and anti-money-laundering (KYC and AML) at the creation-redemption level (e.g., turning USDT into USD or vice versa via Tether's creation-redemption mechanism), KYC and AML are intrinsically impossible at the transaction level on a permissionless blockchain. Just as a criminal transacting using a suitcase of \$100 bills can evade KYC and AML detection, so can a criminal sending a payment in USDT or USDC on a permissionless, and hence anonymous, blockchain.

⁷ Throughout the paper, we take regulation as given. We note, however, that if stablecoins put pressure on inefficient aspects of cross-border payments regulation to improve, that will reduce real resource costs for all kinds of cross-border payments, not just stablecoins. See further discussion in Sections 3.1 and 5.1.

incumbent rents (credit card companies, consumer banks), rather than entry that reduces the economic costs of cash transfers per se. As an example of the scope for new transactions, we discuss online micropayments managed by artificially intelligent agents.

The potential gains we discuss throughout the paper are hard to quantify precisely. Where we can, we try to give a credible sense of the order of magnitude of potential gains. Table 1 summarizes. For reducing real resource costs and improving balance sheet efficiency we estimate potential gains on the order of tens of billions of dollars per year. For reducing intermediation rents the potential gains may be on the order of a hundred billion dollars per year or more. We emphasize, however, that gains that come from reducing intermediary rents are intrinsically harder to realize than gains that come from efficiency improvements per se. One central reason is political economy—incumbents have significant incentive to lobby for regulations that preserve their rents (Olson 1971). As one recent example, the 2025 GENIUS act that legalized stablecoins explicitly prohibits stablecoins from paying interest to customers (see Section 4 Item (a)(11) “Prohibition on Interest”). This feature of the bill prevents beneficial competition on interest rates, and in doing so protects economic rents of both permissionless stablecoin issuers like Tether and Circle, and traditional banks with large deposit franchises.⁸

For facilitating new transactions, we are unable even to give a credible order of magnitude. What our applied game theoretic discussion shows is that it is only transactions that are relatively low value and among relatively low-frequency trading partners that are not already happening given current technology. It seems to us intrinsically impossible to estimate how much potential activity there is in this category. In many other contexts the so-called “long tail” of small transactions has proved to be of significant economic value (Anderson 2006), and there could be significant value here too.

⁸ A related bill currently under consideration (as of this writing, it has passed the House but not the Senate), called The Anti-CBDC Surveillance State Act (H.R. 1919), would explicitly prohibit the Federal Reserve from issuing a central bank digital currency (CBDC). This is another way to restrict competition against both permissionless stablecoin issuers and traditional banks with large deposit franchises.

Table 1. Potential Value of Idealized Data Structure for Traditional Finance

Source of Gains	Method	Estimate
Transactions already taking place		
Reduced real resource costs	# back-office employees saved × \$ per employee	\$10 billion
Improved balance sheet efficiency	\$ balance sheet saved × Modigliani–Miller violation costs per \$	\$7.5 billion – \$15 billion
Reduced intermediation rents	\$ intermediated assets × recordkeeping costs per \$	\$100 billion – \$200 billion
New transactions enabled	# new transactions enabled × surplus per transaction	Potentially very large. Very hard to estimate: long tail of small surplus transactions.

Note: See Sections 3 and 4 for details behind estimates.

The remainder of this paper is organized as follows. Section 2 defines the scope of the exercise. Section 3 analyzes the potential gains for existing transactions. Section 4 analyzes the potential gains from facilitating new transactions. Section 5 applies the framework to stablecoins. Section 6 concludes.

2. Scope

This paper aims to understand the potential economic benefits of blockchain technology in traditional financial intermediation. In this section, we define our usage of these terms to draw a boundary around the set of issues we will engage with.

2.1 Blockchain Technology

The Novel Trust Model

As discussed in Budish (2025), blockchain technology as embodied in cryptocurrencies like Bitcoin and Ethereum is comprised of both a novel data structure and a novel trust model. The trust model, known in computer science as “permissionless consensus” (Nakamoto 2008; Lewis-Pye and Roughgarden 2023), is intuitively like an anonymous majority voting system. Votes most commonly take the form of either computational power (in proof-of-work protocols like Bitcoin) or cryptographically locked-up protocol tokens (in proof-of-stake protocols like Ethereum). As Budish (2025) shows, manufacturing trust in this way is very expensive at scale, because of the incentive to obtain a majority of the votes (which are fully anonymous by design) and attack the system. Budish (2025) shows that the cost of maintaining the integrity of the system against such attacks scales linearly with the economic value of attacking the system. The cost of trust thus quickly grows absurd when contemplating the use of this trust model for traditional finance. In some attack scenarios, the cost of the novel blockchain trust model would have to exceed all of global GDP.⁹

Moreover, the economic usefulness of permissionless blockchains has also been limited. To date, the large majority of volume is financial speculation, and the other widely-documented use case is illegal activity (Makarov and Schoar 2025; Foley et al. 2019; Cox 2021; Gensler 2021; Buterin 2023; Faux 2023; Griffin and Mei 2024). Ironically, much of

⁹ The specific number in Budish (2025) is that securing the Bitcoin proof-of-work protocol against a \$40 billion attack, using Nakamoto trust alone (i.e., without any support from rule of law), would cost all of global GDP in Budish’s base case based on a 6-block escrow period. We show below that the U.S. financial system has transaction volume that easily exceeds \$1 quadrillion per year, which is \$4 trillion per trading day (250 trading days), or \$250 billion per trading hour (16 hour trading day), or \$42 billion per ten minutes.

the financial speculation volume flows through cryptocurrency exchanges, such as Coinbase and Binance, and even traditional financial exchanges via products like Blackrock's iShares Bitcoin Trust—exactly the kind of centralized financial intermediaries Nakamoto (2008) was trying to eliminate.

The Novel Data Structure

The blockchain data structure is not as intellectually exciting as the blockchain trust model, but it has several core features that are compelling for some practical applications. These core features are:

- Append-only distributed database.
- Well-defined permissions that define what parties can add what kinds of transactions, enforced using cryptographically secured signatures.
- Smart contracts: programmable instructions that execute transactions based on past data and the current state.

Additionally, in the idealized version we envision that is supported by rule of law, it would also have:

- Legally verified identities for all parties, including for know-your-customer (KYC) and anti-money-laundering (AML) requirements.
- Cryptographic privacy when appropriate that is decryptable by rule of law when appropriate.
- Frictionless user interface.

Again, the blockchain data structure is not particularly exciting from a computer science perspective: append-only databases with well-defined permissions have been around for a long time, and the phrase “smart contracts” is in a sense just a fancy term for ordinary computer code. But, in comparison to some of the data structures used in the modern financial system, a well-architected distributed database sounds like Nirvana. This claim will be obvious to anyone who has sent a wire, and even large financial institutions with strong incentives to invest in technology use antiquated systems that are vulnerable to significant errors. The Citigroup error cited in the introduction is one famous recent

example and there are many others. The Larry Fink quote in the introduction about using technology from the era of fax machines also conveys the issue.

In this paper, we will not engage with the design details of such a data structure. We will instead posit that the ideal data structure exists and explore what economic gains it might provide.

2.2 Financial Intermediation

Given the goal of our paper, we focus on the kinds of financial intermediation where data structure plays a central role. The key distinction we draw is between financial intermediaries that make decisions and financial intermediaries that simply record and execute decisions. Blockchain technology is most directly relevant for the latter. This includes intermediaries involved in payments and transactions clearing.

The following example is illustrative. Consider a transaction in which party A issues a bond to party B. A financial intermediary, for instance the Depository Trust and Clearing Corporation, keeps track of the bond and verifies that the transaction took place. The clearing organization essentially maintains a database and the validity of the database is supported by traditional sources of trust like reputation and the rule of law. For instance, the bond may be registered with the Securities and Exchange Commission.

Now consider a second transaction in which party B deposits money in a financial intermediary, for instance Bank of America, which the intermediary then lends to party A. In many respects, this is similar to the first transaction. Cash flows from B to A and a financial claim flows in the other direction. Bank of America maintains a database keeping track of these flows and that database is supported by traditional sources of trust. For instance, Bank of America is regulated and audited.

The key difference between these transactions is that in the latter case the intermediary made a decision—to lend money to party A. The power to make that decision and the responsibility for assessing A's credit-worthiness lies with the intermediary. In contrast, in the first transaction the decision-making authority is party B's. The

intermediary simply executes and records the decision. This is the kind of intermediation our analysis focuses on.

Of course, technological improvement in recording and executing decisions may ultimately lead to improvement in decision-making technology as well. Our analysis will not attempt to quantify these indirect effects, but they could be very important. The Lehman example cited in the introduction is a case in point.

2.3 Financial Market Quantities

The quantity of financial intermediation in the categories we consider is large, summarized by the key statistics in Table 2 below. According to the Federal Reserve's Financial Accounts of the United States, the total quantity of intermediated assets stands at over \$100 trillion. According to data from the Securities Industry and Financial Markets Association, total transaction volume in fixed income securities markets was \$368 trillion in 2025 (about \$1.5 trillion per trading day), and U.S. equity market volume was about \$206 trillion (about \$800 billion per trading day). In sum, this totals about \$575 trillion of volume per year.

If we include other "on-balance sheet assets" that show up directly on financial institutions' balance sheets, such as currency, real estate and loans, a credible order of magnitude for U.S. financial system volume is \$750 trillion of volume per year. Derivatives volume is larger still. The CME Group alone has nearly \$3 quadrillion of annual trading volume (notional). Thus, a reasonable order of magnitude for all financial system volume, inclusive of both on-balance-sheet assets and derivatives, is \$5 quadrillion per year.

Table 2. U.S. Financial Market Volumes.

Quantity	Source	Estimate
<i>Stocks of Assets</i>		
Stock of intermediated assets	Federal Reserve	\$141 trillion
<i>On-Balance-Sheet Assets Trading Volume</i>		
Annual trading volume for Treasuries	Securities Industry and Financial Markets Association	\$264 trillion
Annual trading volume for all fixed income (incl. Treasuries)	Securities Industry and Financial Markets Association	\$368 trillion
Annual trading volume for equities	Securities Industry and Financial Markets Association	\$206 trillion
Annual spot FX trading volume	NY Fed (Foreign Exchange Committee)	\$122 trillion
<i>Payments Volume</i>		
Annual Fedwire transfers	Federal Reserve	\$1.1 quadrillion
Annual global payments industry volume	McKinsey	\$2 quadrillion
<i>Derivatives Volume</i>		
Annual trading volume for CME Group products (notional)	CME Group	\$2.7 quadrillion
Annual trading volume for FX derivatives (notional)	NY Fed (Foreign Exchange Committee)	\$204 trillion
<i>Overall Volume Magnitudes</i>		
Annual trading volume for on-balance-sheet assets (order of magnitude)	Above sources and authors' estimates	\$750 trillion
Annual trading volume for on-balance-sheet assets and derivatives (order of magnitude)	Above sources and authors' estimates	\$5 quadrillion

Sources: CME Group (2026), Federal Reserve Banks (2026), Federal Reserve Board (2025, 2026), Foreign Exchange Committee (2026), McKinsey (2025), SIFMA (2026a, 2026b, 2026c). The stock of intermediated assets is as of the end of Q3 2025. FX spot and derivatives volumes are annualized based on volume during October 2025 and include volume in all of North America. All other figures are for calendar year 2025.

3. Potential Gains for Existing Transactions

In this section, we consider a first broad source of gains to end users from the idealized data structure in traditional finance: reductions in costs for transactions that are already taking place. Such reductions can arise because the real resource costs of executing transactions fall or because the total markups charged by intermediaries fall.

3.1 Reducing Real Resource Costs

The idealized data structure would help reduce the real resource costs of transactions in several ways. The simplest one is by reducing the number of people involved in back-office functions at financial intermediaries. Every large intermediary employs a significant number of people in back-office and treasury functions, verifying transactions, ensuring that funding is properly managed, and fulfilling legal and compliance obligations. For instance, the Depository Trust and Clearing Corporation (DTCC) employs over 5,000 people (DTCC 2021). Any intermediary transacting through the DTCC will also employ people in similar functions. By simplifying and potentially enabling the automation of transactions clearing, blockchain technology could significantly lower back-office costs.

We conservatively estimate that the idealized data structure could reduce the required number of back-office positions devoted to recordkeeping by 100,000. A reasonable conservative figure for average total compensation of such employees is \$100,000 per year. Thus, this dimension of cost savings cumulates to \$10 billion per year.¹⁰

We make a clarifying remark about regulatory compliance costs. In principle there are two distinct ways that the idealized data structure could reduce the real resource costs of regulatory compliance. First, it could reduce the costs of complying with existing rules

¹⁰ According to the Bureau of Labor Statistics (https://www.bls.gov/oes/2023/may/naics2_52.htm), in 2023 the finance and insurance sectors employed approximately 90,000 book-keeping, accounting, and auditing clerks. In addition, it employed 65,000 computer and information systems managers, 120,000 accountants and auditors, 41,000 bill and account collectors, 20,000 billing and posting clerks, and 45,000 brokerage clerks. The mean wage for book-keeping, accounting, and auditing clerks in all industries was approximately \$50,000 (<https://www.bls.gov/oes/2023/may/oes433031.htm>). Accounting for the 30–50% finance wage premium (Philippon and Reshef 2012) and the fact that wages are only part of total compensation, we use \$100,000 per year as our estimate of total compensation costs for such employees.

and regulations. Second, it could put pressure on regulators to improve these rules and regulations, for example by shining a light on inefficient aspects of existing regulation. Call these a direct effect and an indirect effect. Throughout our analysis, we take regulation as given, capturing the direct effect, not the indirect effect. In the long run, however, the indirect effect could prove more important. We come back to this topic in our discussion of stablecoins and cross-border payments regulations below.

3.2 Improving Balance Sheet Efficiency

Another potential source of gains is from more efficient netting. In the U.S., the typical equity transaction now takes a day to clear, and took two days to clear as recently as 2024. Treasury transactions and corporate bond transactions also take a day to clear and took longer until recently. Many loan transactions can take several weeks to clear.

This clearing process locks up money and assets in the financial system. To get a sense of the magnitudes involved, suppose that across all physical assets, clearing takes one day on average. In Table 2, the annual trading volume of on-balance-sheet assets is on the order of \$750 trillion per year. That means that approximately \$3 trillion ($\$750 \text{ trillion} / 252 \text{ trading days}$) of assets are locked up each trading day.

Note that the securities that are going through the clearing process do not themselves lose any value. They are still accruing dividends and interest. Moreover, transaction prices properly account for the fact that one party in the transaction receives those payments and the other does not. For instance, in the U.S. stock market, the so-called “ex dividend date,” the last date on which a buyer of a stock can purchase it and receive the next dividend, is now one trading day before the “record date,” the date on which ownership is verified for dividend payments, because clearing now takes one day. Stock prices fall immediately after the ex dividend date because after that date buyers receive one less dividend than buyers prior to that date.

Rather, the true net cost of clearing is that the \$3 trillion of assets that are tied up each day must be funded by financial intermediaries. Under the idealized conditions of the Modigliani-Miller (1958) theorem, the net cost of that funding would be zero: the returns

on the assets would be exactly determined by their risk and equal to the cost of funding. However, a large literature has found that Modigliani-Miller does not hold for real world financial intermediaries. The net cost formula is thus given by:

$$Net\ Cost = MM\ violation\ cost\ (\%) \times \frac{\$ Balance\ Sheet\ Tied\ Up}{(Days\ to\ Clear/250) \times Total\ Transaction\ Volume} \quad (1)$$

The literature finds MM violation costs on the order of 0.25–0.5% of financial intermediary balance sheet capital.¹¹ Thus, equation (1) tells us that the net cost of funding \$3 trillion of assets is approximately \$7.5–\$15 billion per year. The idealized data structure could reduce clearing time to approximately zero and save this amount.

3.3 Reducing Intermediation Rents

Transaction costs faced by end users are the sum of the marginal cost of executing the transaction and the markup charged by the intermediary or set of intermediaries involved in the transaction. There are often significant economies of scale in this kind of intermediation. These economies of scale stem both from the nature of the activity—there are often significant market thickness externalities and high fixed costs involved in setting up the infrastructure to execute transactions—as well as regulation, which adds to fixed costs and creates barriers to entry. As a result, many types of financial intermediation are dominated by a few large players, raising concerns that markups are excessively high. Huberman, Leshno and Moallemi (2021) point out that global payments revenues are about \$2 trillion per year.

The gains from reducing markups are potentially large. As documented by Philippon (2015), the cost (including markups) of all financial intermediation in the U.S. has been roughly constant for the past 100 years at 2% per year of intermediated assets. Given the enormous growth we have seen in the quantity of intermediated assets, this means that finance accounts for 7–8% of GDP (Greenwood and Scharfstein 2013). These

¹¹ See, e.g., Basel Committee on Banking Supervision (2010); Kashyap, Stein, and Hanson (2010); Admati et al. (2013); Basel Committee on Banking Supervision (2019); Fleckenstein and Longstaff (2020).

numbers include both the costs of recordkeeping, which are our focus, as well as the costs of financial decision making discussed in Section 2.2. Thus, the idealized data structure would not reduce all of these intermediation costs, but could still generate hundreds of billions of dollars of value per year. A simple back of the envelope calculation would be to suppose that 0.10–0.20% of Philippon (2015)’s total 2.00% cost of financial intermediation is related to recordkeeping. In other words, suppose that intermediation costs related to recordkeeping are 5–10% of the total cost of intermediation. U.S. financial assets are over \$100 trillion, so the costs of recordkeeping could easily be \$100–\$200 billion per year.

An important caveat is that even the hypothetical idealized data structure might not eliminate all intermediary rents associated with recordkeeping. There are two separate reasons which may reinforce one another. First is the “money doctors” role of financial intermediaries emphasized by Gennaioli, Shleifer and Vishny (2015). Even if technology makes high-quality recordkeeping trivial and costless, investors may still value the peace of mind that comes from knowing their money is in the hands of institutions with a long history of competence.

Indeed, empirical evidence suggests that markups remain high in some areas of finance that in principle seem like they should be relatively commoditized and competitive, like retail deposit taking. The issue is not that depositors lack options, but that they are quite inelastic, perhaps because they value the peace of mind they get from trusted names in banking with a long track record (see, e.g., Egan, Hortacsu, and Matvos (2017) and Egan, Lewellen, and Sunderam (2021)). This makes entry more difficult than one might guess given the simplicity of the technology per se.

Second, incumbent firms earning rents from recordkeeping will have significant incentive to lobby to preserve their rents, which in turn may affect the regulation of entry (Olson 1971). The GENIUS Act and Anti-CBDC Surveillance State Act discussed in the introduction are salient recent examples.

4. Enabling New Transactions

In this section, we study a second potential source of gains from the idealized data structure: enabling new transactions that would not have otherwise taken place. We consider a game-theoretic environment where two parties can engage in mutually beneficial transactions but can cheat. Rule of law is imperfect in the sense that recovering the stolen funds may cost more than the stakes on a one-shot basis (i.e., ignoring reputational considerations). A simple insight from this approach is that transactions that generate relatively large surplus between parties that transact relatively frequently are likely already taking place. What technology can unlock are transactions that are either relatively lower surplus or that are between parties that transact relatively infrequently. In other words, better data structure can enable traditional trust to extend to a potentially much larger set of transactions.

4.1 Setup

We formalize our approach with the payoff matrix in Figure 2 below. The figure portrays an augmented Prisoners' Dilemma similar to game forms studied in Dixit (2004). The two players in the game are considering a transaction. There are three options. Each player can choose to engage in the transaction or not. If either player chooses not to engage, they both get zero payoffs. If both players choose to engage, they then each have the option to cooperate or cheat. This part of the game then has the standard payoffs to a Prisoners' Dilemma. If both players cooperate, they both get a small positive payoff: the surplus from completing the transaction, f . However, if one player cheats and the other player cooperates, then the player that cheats gets a large positive payoff, $V > f$, and the player that got cheated gets a large negative payoff, $-V$. We think of V as representing the nominal size of the transaction (e.g., a \$100 bond) in contrast to f which represents the net benefit of the transaction (e.g., moving the \$100 bond from the party who values it less to the party who values it more creates \$5 of surplus). For completeness, if both players cheat, for instance, by refusing to send the cash or securities they were supposed to send to the other party, then they both get a small negative payoff, denoted $-\varepsilon$.

Figure 2. Payoffs in Augmented Prisoners' Dilemma.

		Player 2		
		Engage, Cooperate	Engage, Cheat	Do not Engage
Player 1				
Engage, Cooperate		$+f$	$+V$	0
	$+f$		$-V$	0
Engage, Cheat		$-V$	$-\varepsilon$	0
	$+V$		$-\varepsilon$	0
Do not Engage		0	0	0
	0		0	0

Note: $f > 0$ represents the gains from trade in an honest transaction. $V > f$ represents the gain to a cheating party from cheating in the transaction. In traditional finance the relevant case is that f much smaller than V . f is on the order of the willingness to pay to engage in the transaction, whereas V is on the order of the size of the transaction.

We need to make two points of clarification about this game. First, we take cheating here to be a broad category including everything from back-office mistakes to willful acts of theft. We assume in this part of the analysis that rule of law is incomplete in the sense that the cost to the cheated party of recovering the V via rule of law is larger than V on a one-shot basis (i.e., ignoring any reputational benefits of committing to always try to recover). We will make the role of rule of law more explicit below.

Second, while we represent the game as a static game with three actions, this static game is equivalent for our purposes to the following extensive form game: in a first stage, the players decide whether or not to engage; in a second stage, the players decide whether or not to cheat; last, if one player cheats the other, there is a third stage in which the cheated party decides whether or not to try to recover their funds. In the current form of the game

we assume that the cheated party's cost of trying to recover their funds exceeds their expected recovery, so they choose not to try and their payoff is $-V$.

4.2 Initial Analysis

Static Equilibrium

The game depicted in Figure 2 has a unique static equilibrium: both parties choose not to engage. We interpret the static equilibrium as capturing that some socially valuable transactions do not currently take place because of the risk of cheating.

Repeated Play and Reputation

Now consider repeated play of this game. By standard folk-theorem logic (Schelling 1956; Aumann 1959), repeated interaction can facilitate cooperation. Formally, let N be the frequency with which the parties interact (e.g., number of transactions per year), and let δ be the discount factor. Then the condition for {Engage, Cooperate; Engage, Cooperate} to be a Nash equilibrium of the repeated game is

$$\frac{\delta}{1-\delta} Nf > V \quad (2)$$

Equation (2) says that even if rule of law is imperfect in the sense that it costs more than V for a cheated party to recover their funds, parties that transact frequently (N), are long-lived and patient (δ), and generate surplus from transacting (f) are able to transact safely. Thus, folk-theorem logic alone can help explain why financial institutions that interact frequently, for instance Goldman Sachs and JPMorgan, can sustain cooperation.

In contrast, parties that have the potential to transact only infrequently will have more fear of cheating and hence less trust. Here are two examples to fix ideas. First, consider small banks in two different countries that do not transact much with each other. The banks might reasonably prefer not to transact directly but instead through a mutually trusted large intermediary with a presence in both countries (like a JPMorgan). In effect, the small banks would pay some portion of the gains from trade f to the mutually trusted

intermediary as a fee for creating trust. If the gains from trade are not sufficiently large relative to intermediary fees, the transaction will not occur at all.

Second, consider a small online payment between two anonymous parties, e.g., a payment by one AI agent to another, or a transaction between humans for a pair of concert tickets. The parties might prefer to transact through a trusted intermediary (e.g., an escrow service, a secondary market platform), paying some of the gains from trade to the intermediary. Again, if the gains from trade are not sufficiently large relative to intermediary fees, the parties might not transact at all.

4.3 Three Mechanisms for Enabling Cooperation

How could the idealized data structure help coordinate on the cooperate-cooperate outcome? We describe three mechanisms, depicted in Figure 3 below.

First, it can change the possible actions in the static game, reducing the scope for cheating (Figure 3, Panel A). For instance, suppose that two small banks in different countries want to engage in a bond trade but worry that the other party will not follow through. A smart contract could encode that the transaction that transfers the bond from the selling bank to the buying bank is paired and simultaneously cleared with the transaction that transfers cash from the buying bank to the selling bank. This effectively eliminates {Engage, Cheat} from the game between the two banks. If {Engage, Cheat} is not an option, then {Engage, Cooperate} - {Engage, Cooperate} becomes the unique equilibrium even in one-shot play. This eliminates any benefit from transacting through a larger trusted intermediary like a JPMorgan.¹²

Second, the idealized data structure could make cheating in the static game less attractive, e.g., by making it harder to cover one's tracks (Figure 3, Panel B). Having a public, indisputable data structure may make it easier and cheaper for the cheated party to

¹² Of course, this requires that the parties can verify that the relevant {Engage, Cooperate} transactions are programmed correctly. There have been numerous examples of sophisticated parties duped into executing malicious transactions. A recent case in point is the March 2025 \$1.5 billion Bybit hack, where the cryptocurrency exchange Bybit's CEO was duped into signing a transaction that transferred \$1.5 billion worth of Ethereum to North Korean hackers (Rajic and Brock 2025).

complain to law enforcement about the cheating party and try to recover their funds. If detection and enforcement are good enough, this could change the cheater's payoff from the $V > 0$ of stolen funds to a $-P < 0$ legal penalty.¹³ This again makes {Engage, Cooperate} - {Engage, Cooperate} the unique equilibrium even in one-shot play.

A third possibility is that, by making players' behavior more easily observable and verifiable by others, the data structure enhances the possibilities for reputation building and hence cooperation (Figure 3, Panel C). For instance, if one of the small banks in the bond trade described above cheats, it could become easier for other financial institutions not party to the transaction to observe the cheating (e.g., that they received the bond but did not follow through with the cash) and stop trusting the cheater. Mathematically, add a parameter M that represents the number of financial institutions using the idealized data structure, assume that each pair of institutions interacts N times per year as before, and assume that any one act of cheating can be detected by all parties. Then the condition for cooperation to be an equilibrium becomes:

$$\frac{\delta}{1-\delta} MNf > V \quad (3)$$

Equation (3) captures that the idealized public data structure could enlarge the frequency of interaction parameter from the number of times a party interacts with one counterparty per year (N) to the number of times a party interacts with *all* other counterparties using the data structure per year (MN). The cost of cheating becomes the cost of exclusion from the entire system, making cooperation more sustainable.

We emphasize that these three mechanisms are logically complementary. All three mechanisms could work together to enable transactions that are not currently possible (at least directly between the parties) because of gaps in trust.

¹³ The 0 for the cheated party in the in the cheat-cooperate cell assumes that the cheated party recovers their stolen funds and any legal costs. Alternatively, the payoff in the cheat-cooperate cell could be modified to $(-P, -H)$ where the H term represents the hassle cost to the cheated party of recovering their funds. As long as $H < V$ they will address the cheating (in the suitably augmented extensive form game discussed above in 4.1) and the analysis is the same as in the $(-P, 0)$ case. If the cost of addressing the cheating is too high, then we are back to the original scenario depicted in Figure 2.

Figure 3. Three Ways a Blockchain Can Facilitate Cooperation

Panel A: Technology		Panel B: Punishment		Panel C: Reputation	
Player 1		Player 1		Player 1	
Engage, Cooperate	Engage, Cheat	Engage, Cooperate	Engage, Cheat	Engage, Cooperate	Engage, Cheat
+f	+V	+f	-P	+f	+V
+f	-V	+f	0	+f	-V
Engage, Cheat	-ε	Engage, Cheat	-ε	Engage, Cheat	-ε
+V	0	-P	0	+V	0
0	0	0	0	0	0
0	0	0	0	0	0
Do not Engage	Do not Engage	Do not Engage	Do not Engage	Do not Engage	Do not Engage
Player 2		Player 2		Player 2	
Engage, Cheat		Engage, Cheat		Engage, Cheat	
+V	-ε	-P	-ε	+V	-ε
-V	0	0	0	-V	0
-ε	0	-ε	0	-ε	0
0	0	0	0	0	0
0	0	0	0	0	0

If a blockchain eliminates the technological possibility of cheating, then {Cooperate, Cooperate} becomes the unique static equilibrium.

If a blockchain makes cheating more easily detectable and punishable via rule-of-law, then {Cooperate, Cooperate} becomes the unique static equilibrium.

If a blockchain makes past play widely observable across the financial sector, then {Cooperate, Cooperate} can become an equilibrium of the *dynamic* game even if in the one-shot game there remains a possibility to cheat.

4.4 Discussion

While stylized, this analysis based on gains from trade and fear of cheating has two important messages.

First, parties that have the opportunity to transact relatively frequently (high N), and/or have high gains from trade (high f), are likely already able to trust each other under the current system (equation (2)). The transactions that are missing must be either relatively low surplus (low f) or relatively infrequent (low N).

Second, better data structure can indeed facilitate these missing transactions. We described three complementary channels: technologically eliminating the possibility of cheating, providing better records for traditional law enforcement, and better records for forming public reputations (Figure 3).

How much value is there in such missing transactions? Define Q to be the number of pairs of potential transacting parties. Then the total value that could be generated by facilitating these missing transactions is:

$$\text{Value of Enabled Transactions} = \underbrace{Q \cdot \Pr \left[Nf \leq V \frac{1-\delta}{\delta} \right]}_{\# \text{ of pairs}} \cdot \underbrace{E \left[Nf \mid Nf \leq V \frac{1-\delta}{\delta} \right]}_{\text{Value per pair}} \quad (4)$$

We know from our discussion above that the last term in equation (4), the value per pair, cannot be very large; else the pair would be transacting already. However, the number of potential transacting pairs could be very large. Potential examples that seem credible to us include transactions between AI agents and other forms of online micropayments.

If Q is large, then the aggregate value of enabled transactions could be large as well, even if each individual transaction pair is modest in size.

5. Application: Stablecoins

In this section, we apply the framework built up in the prior two sections to stablecoins. Stablecoins are digital currencies recorded on distributed ledgers that are pegged to a reference value. The reference value is typically a fiat currency such as the U.S. dollar, but can also be a basket of currencies, commodities, or other assets.

Stablecoins are a useful setting for illustrating the value of our framework for three reasons. First, as underscored by the recent passage of the GENIUS Act in the U.S., they have captured the attention of both market participants and regulators. Second, stablecoins are a natural application for our framework as, given their intended stable value, their use case is payments and recordkeeping. Third, as we detail below, our framework suggests that they are likely to generate value for traditional finance in certain areas but not others. This illustrates that our framework has analytical bite.

Following the key distinction we make throughout the paper, we focus on the potential for stablecoins that use the blockchain data structure but rely on traditional sources of trust including rule of law and reputation. In other words, our focus is on the potential for stablecoins operating on permissioned blockchains. In contrast, the largest existing stablecoins like Tether and Circle operate on permissionless blockchains, relying on the decentralized and anonymous trust model inherent in those systems which makes them dangerous for traditional finance. Permissionless blockchains are also intrinsically anonymous at the account level, making them attractive to illegal actors. Indeed, documented use cases of stablecoins include the online scam economy, financing of terrorist groups, and money laundering (Faux 2023; Griffin and Mei 2024; Foldy 2023; United Nations 2024).¹⁴ That said, by far the largest use case of stablecoins to date is

¹⁴ There are also some more sympathetic “illegal” use cases for permissionless stablecoins. These include their use as a payments medium by individuals in countries with volatile fiat currencies, and their use for international remittances to circumvent high fees. For example, see the discussion in Jones and Tabarrok (2025). Evidence in surveys by Castle Island and Brevan Howard (2024) and Ante (2025) suggests that both of these use cases are common in the relevant populations, e.g., 34% of cryptocurrency users in five developing countries reported “buying or selling goods or services” as one of their primary goals when using stablecoins, and 26% of U.S. adults actively engaged in remittances reported that they sent a remittance in stablecoins in the past year.

cryptocurrency trading (McKinsey 2026; Visa and Allium Labs 2026).¹⁵

Throughout our analysis, we compare stablecoins to three alternatives in the traditional financial system. First, credit card providers like Visa. These providers charge high fees, on the order of 2–3% of transaction volume. However, they provide a bundle of services beyond basic transactions clearing, including short-term lending, fraud protection, and credit card points and other benefits. Second, existing services like Zelle and Venmo that clear transactions at zero cost to consumers within the existing regulatory and legal framework. These services offer relatively little fraud protection and other features to users. However, they tend to limit transaction size, not due to technological limitations, but as a form of consumer protection. Third, services like the Automated Clearing House (ACH) that allow users to transfer large sums of money with safeguards, but that are slower and based on legacy technology.

5.1 Stablecoins in Transactions Already Taking Place

We now apply our framework developed in Section 3 to evaluate possible gains from stablecoins for transactions already taking place. Recall from Section 3 that there were three potential sources of value: (i) reductions in real resource costs like back-office costs, (ii) improvements in balance sheet efficiency, and (iii) reductions in intermediation rents. We consider each source of value in turn.

5.1.1 Do Stablecoins Reduce Real Resource Costs?

Stablecoins seem likely to offer at most modest reductions in real resource costs for transactions already taking place in the traditional financial system under existing laws and regulation. The \$10 billion of real resource cost savings identified in Section 3.1 largely

¹⁵ Visa and Allium report \$12.4 trillion of annual adjusted transaction volume for stablecoins (December 2025 annualized). This calculation strips out blockchain noise as in Makarov and Schoar (2025); e.g., if user A sends user B \$1 from an address with \$10 in it, this might get recorded as a transaction of \$1 and a transaction of \$9, where the latter represents A sending money to themselves as a remainder. McKinsey in conjunction with Artemis computes that annual stablecoin payments volume is about \$390 billion (December 2025 annualized), or about 3.2% of Visa’s total. They conclude that stablecoin volume “consists mainly of trading, internal shuffling of funds, and automated blockchain activity,” echoing similar findings for Bitcoin volume from Makarov and Schoar (2025).

stems from recordkeeping for more complex securities like stocks, bonds, and loans. Updating ledgers keeping track of cash could offer some gains, but they are likely to be a small part of the total.

The simplest evidence for this view is that the cost to users of services like Zelle, Venmo and ACH is already zero. Current stablecoin transactions on permissionless blockchains actually have a slightly positive transaction cost because they take up space on a permissionless blockchain ledger and those intrinsically have to charge fees for security (Huberman, Leshno and Moallemi 2021; Halaburda et al. 2022; Budish 2025).

We note in contrast that current stablecoins are noticeably cheaper than debit and credit cards. However, those are different products. Both debit and credit cards bundle in fraud protection, whereas a stablecoin transaction on a permissionless blockchain is irreversible other than via majority attack.¹⁶ Credit cards also bundle in float, short-term lending, and incentives like cashback or frequent flyer points.

Thus, for assessing the ability of stablecoins to reduce real resource costs, the relevant comparison is to services like Zelle, Venmo and ACH, and the savings look to be small at best, at least for domestic transactions.

For cross-border transactions, some of the current appeal of stablecoins on permissionless blockchains like Circle and Tether is that they evade payments regulations. Users of stablecoins on the permissioned blockchains we analyze, however, would still have to comply with these regulations and thus pay the costs of regulatory compliance. In other words, taking regulation as given, as we do throughout our analysis, the benefits of stablecoins for cross-border transactions may be limited when those stablecoins are on permissioned blockchains. On such blockchains, better recordkeeping may reduce regulatory costs, but it will not eliminate them in the way that permissionless blockchains can. Where we see the largest potential for real cost savings is what we described in Section 3.1 as the indirect effect, if the existence of permissionless stablecoins puts pressure on

¹⁶ This irreversibility is the source of so-called “wrench attacks,” in which a criminal forces a crypto holder at gunpoint to irreversibly send their funds. See Carr (2026).

regulators to improve the regulatory efficiency of cross-border payments on permissioned financial rails.

5.1.2 Do Stablecoins Improve Balance Sheet Efficiency?

Improved balance sheet efficiency is also unlikely to be a source of significant value for stablecoins. When financial intermediaries transact, the cash part of the transaction already clears very quickly. Thus, in equation (1) days to clear is approximately zero for cash transactions, making the net gains approximately zero.

Sometimes it takes longer for consumer or small-business cash transactions to clear, for example for sending a large wire for a home purchase, but this is usually because of fraud checks and other regulatory requirements. As discussed above, it is certainly possible that some of these regulations are inefficient but if so the potential for gains is not the blockchain technology per se but the potential to improve inefficient regulation.

5.1.3 Do Stablecoins Reduce Intermediation Rents?

Where stablecoins could add significant value for transactions already taking place is by reducing intermediation rents, particularly when compared to credit card providers like Visa, Mastercard and American Express. Large retailers like Amazon and Walmart have expressed interest in issuing their own stablecoins to avoid paying transaction fees to these providers (Heeb et al. 2025). To get a rough sense of potential gains, note that Visa, Mastercard, and American Express together have a combined market capitalization of nearly \$1.5 trillion. If stablecoins eliminated half of these companies' intermediation rents, this could have a net present value on the order of 750 billion, or roughly \$75 billion per year at a 10% weighted average cost of capital.

However, it is unclear whether stablecoins per se will be the technology to unlock such gains, or whether existing services like Zelle and Venmo will continue to grow and displace more of the payments volume currently paying 2–3% fees to credit card providers.

5.2 Stablecoins Enabling New Transactions

Another way stablecoins could add significant value is by enabling new transactions. Here, stablecoins have a meaningful advantage relative to services like Zelle and Venmo because of their programmability via smart contracts (Garratt and Lee 2025; Jones and Tabarrok 2025).

In Section 4, we argued that such new transactions would have to be relatively infrequent (low N) and relatively low surplus (low f). The programmability of stablecoins may be particularly suited to facilitating such transactions. For instance, stablecoins on the idealized data structure may make it substantially easier than app-based services like Zelle and Venmo for micropayments to flow between users and websites or between AI agents.

6 Conclusion

We have developed a conceptual framework for analyzing the potential gains for traditional finance from an idealized version of the blockchain data structure (Figure 1 in the introduction).

Our analysis of real improvements in existing transactions identifies two potential sources of value: reducing real resource costs and improving balance sheet efficiency (Sections 3.1–3.2). Whether the numbers discussed sound big or small depends on one’s perspective. The figures are large in dollar terms, easily tens of billions of dollars annually, while modest relative to the scale of modern finance, especially given that the hypothetical thought experiment is an idealized data structure across all of finance. On the other hand, the capitalized value of the savings we discuss is on the order of \$175 billion to \$500 billion (with significant estimation uncertainty in both directions). This is meaningful in comparison to cryptocurrencies’ peak market capitalization of \$4 trillion. It is also meaningful in comparison to the global banking sector’s current market capitalization of \$8 trillion.

If there is a larger payoff lurking, it is either in reducing the ability of firms to exercise market power (Section 3.3), or in facilitating useful new transactions that are currently not taking place because of trust problems (Section 4), or that are currently taking place but are unnecessarily intermediated (e.g., two small banks transacting via a larger bank rather than directly). Reducing intermediation rents could easily be in the hundreds of billions of dollars per year and trillions of present value. However, these savings are the hardest to realize because of incumbents’ incentives to protect their rents. The widespread political opposition in the United States to CBDC—which would reduce the rents of *both* permissionless stablecoin providers like Circle, Tether, and now the Trump family, *and* traditional banks with deposit franchises—is a case in point.

One interpretation of our analysis, which seems intuitively appealing to us, is that blockchain-like technologies will be valuable for finance in an analogous way to how other computational technologies have turned out to be valuable in finance. It is likely to create

small efficiencies in the near term, but if there are large efficiencies they will develop slowly over time.¹⁷

One additional prediction that we are relatively confident about: for all of the hype about decentralized trust, disrupting centralized financial institutions, etc., one of the biggest lasting benefits of the past decade of excitement about blockchains will be simply that it has shined a light on excess rents and outdated technologies and regulations in traditional finance. It will be ironic if the lasting impact of Nakamoto (2008) is to improve the efficiency of the traditional financial system.

¹⁷ By analogy, consider the relational database. The relational database, which today is ubiquitous, was first proposed in 1970 (Mishra and Beaulieu 2002). The SQL query language, which implements the idea, first appeared in the mid 1970s, with the first commercial versions released in 1979 by IBM and Oracle. If asked in the early 1980s to quantify the dollar benefits of the relational database to traditional finance, it likely would have been hard to point to huge amounts of value creation. Today, however, it would be difficult to pay any amount of money to get traditional financial firms to relinquish their relational databases. This suggests that significant value has amassed over time.

References

- Admati, Anat R., Peter M. DeMarzo, Martin F. Hellwig, and Paul Pfleiderer. 2013. “Fallacies, Irrelevant Facts, and Myths in the Discussion of Capital Regulation: Why Bank Equity Is Not Socially Expensive.” Working Paper No. 161. Stanford: Stanford University, Rock Center for Corporate Governance, November 4, 2013. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2349739.
- Anderson, Chris. 2006. *The Long Tail: Why the Future of Business Is Selling Less of More*. New York, NY: Hyperion, July 11, 2006.
- Ante, Lennart. 2025. “From Adoption to Continuance: Stablecoins in Cross-Border Remittances and the Role of Digital and Financial Literacy.” *Telematics and Informatics* 97 (February 2025): 102230. <https://doi.org/10.1016/j.tele.2024.102230>.
- Anti-CBDC Surveillance State Act, H.R.1919, U.S. House of Representatives. <https://www.congress.gov/bill/119th-congress/house-bill/1919>.
- Aquilina, Matteo, Eric Budish, and Peter O’Neill. 2022. “Quantifying the High-Frequency Trading ‘Arms Race.’” *The Quarterly Journal of Economics* 137 (1): 493–564. <https://doi.org/10.1093/qje/qjab032>.
- Atkins, Paul S. 2025. “American Leadership in the Digital Finance Revolution.” U.S. Securities and Exchange Commission, July 31, 2025. <https://www.sec.gov/newsroom/speeches-statements/atkins-digital-finance-revolution-073125>.
- Aumann, Robert J. 1959. “Acceptable Points in General Cooperative N-Person Games.” In *Contributions to the Theory of Games (AM-40), Volume IV*, edited by Albert William Tucker and Robert Duncan Luce, 287–324. Princeton University Press, May 1, 1959. <https://doi.org/10.1515/9781400882168-018>.
- Basel Committee on Banking Supervision. 2010. *Basel III: A Global Regulatory Framework for More Resilient Banks and Banking Systems*. Basel: Bank for International Settlements, December 2010. <http://www.bis.org/publ/bcbs189.pdf>.
- Basel Committee on Banking Supervision. 2019. “The Costs and Benefits of Bank Capital - a Review of the Literature.” Working Paper No. 37. Basel: Bank for International Settlements, June 24. <https://www.bis.org/bcbs/publ/wp37.htm>.
- Budish, Eric. 2025. “Trust at Scale: The Economic Limits of Cryptocurrencies and Blockchains.” *The Quarterly Journal of Economics* 140 (1): 1–62. <https://doi.org/10.1093/qje/qjae033>.

- Budish, Eric, Robin S. Lee, and John J. Shim. 2024. “A Theory of Stock Exchange Competition and Innovation: Will the Market Fix the Market?” *Journal of Political Economy* 132 (4): 1209–46. <https://doi.org/10.1086/727284>.
- Budish, Eric, Andrew Lewis-Pye, and Tim Roughgarden. 2025. “The Economic Limits of Permissionless Consensus.” Preprint, 2025.
- Buterin, Vitalik. 2023. *Make Ethereum Cypherpunk Again*. December 28, 2023. <https://vitalik.eth.limo/general/2023/12/28/cypherpunk.html>.
- Carr, Austin. 2026. “Small-Time Crypto Investors Are Facing Violent Attacks.” Bloomberg Businessweek, January 2, 2026. <https://www.bloomberg.com/features/2026-crypto-thieves-kidnappers/>.
- Castle Island Ventures and Brevan Howard Digital. 2024. *Stablecoins: The Emerging Market Story*. Castle Island Ventures, September 12, 2024. <https://castleisland.vc/writing/stablecoins-the-emerging-market-story/>.
- CME Group. 2026. “Access Global Markets.” 2026. https://www.cmegroup.com/#tab_wCfNRwi=access-global-markets.
- Cox, Jeff. 2021. “Yellen Sounds Warning About ‘Extremely Inefficient’ Bitcoin.” CNBC, February 22, 2021. <https://www.cnbc.com/2021/02/22/yellen-sounds-warning-about-extremely-inefficient-bitcoin.html>.
- Depository Trust & Clearing Corporation. 2021. “DTCC 2021 Annual Report.” 2021. <https://www.dtcc.com/annuals/2021/purpose/people>.
- Dixit, Avinash K. 2004. *Lawlessness and Economics: Alternative Modes of Governance*. The Gorman Lectures in Economics. Princeton, NJ: Princeton University Press, 2004.
- Dwork, Cynthia, Nancy Lynch, and Larry Stockmeyer. 1988. “Consensus in the Presence of Partial Synchrony.” *Journal of the ACM* 35 (2): 288–323. <https://doi.org/10.1145/42282.42283>.
- Egan, Mark, Ali Hortaçsu, and Gregor Matvos. 2017. “Deposit Competition and Financial Fragility: Evidence from the US Banking Sector.” *American Economic Review* 107 (1): 169–216. <https://doi.org/10.1257/aer.20150342>.
- Egan, Mark, Stefan Lewellen, and Adi Sunderam. 2021. “The Cross-Section of Bank Value.” *The Review of Financial Studies* 35 (5): 2101–43. <https://doi.org/10.1093/rfs/hhab089>.
- Faux, Zeke. 2023. *Number Go Up: Inside Crypto’s Wild Rise and Staggering Fall*. First edition. New York: Currency, 2023.

- Federal Reserve Banks. 2026. “Fedwire Funds Service Volume and Value Statistics.” January 26, 2026. <https://www.frbservices.org/resources/financial-services/wires/volume-value-stats>.
- Federal Reserve Board. 2025. “Federal Reserve Payments Study.” March 2025. <https://www.federalreserve.gov/paymentsystems/fr-payments-study.htm>.
- . 2026. *Financial Accounts of the United States*. January 9, 2026. <https://www.federalreserve.gov/releases/z1/>.
- Fink, Larry. 2025. “The Democratization of Investing: Expanding Prosperity in More Places, for More People.” April 3, 2025. <https://www.blackrock.com/corporate/literature/presentation/larry-fink-annual-chairmans-letter.pdf>.
- Fink, Larry, and Rob Goldstein. 2025. “Larry Fink and Rob Goldstein on How Tokenisation Could Transform Finance.” *The Economist*, December 1, 2025. <https://www.blackrock.com/corporate/literature/article-reprint/larry-fink-rob-goldstein-economist-op-ed-tokenization.pdf>.
- Fleckenstein, Matthias, and Francis A. Longstaff. 2020. “Renting Balance Sheet Space: Intermediary Balance Sheet Rental Costs and the Valuation of Derivatives.” *The Review of Financial Studies* 33 (11): 5051–91. <https://doi.org/10.1093/rfs/hhaa033>.
- Foldy, Ben. 2023. “From Hamas to North Korean Nukes, Cryptocurrency Tether Keeps Showing Up.” Markets. *Wall Street Journal*, October 27, 2023. <https://www.wsj.com/finance/currencies/most-popular-cryptocurrency-keeps-showing-up-in-illicit-finance-71d32e5e>.
- Foley, Sean, Jonathan R. Karlsen, and Tālis J. Putniņš. 2019. “Sex, Drugs, and Bitcoin: How Much Illegal Activity Is Financed through Cryptocurrencies?” *The Review of Financial Studies* 32 (5): 1798–853. <https://doi.org/10.1093/rfs/hhz015>.
- Foreign Exchange Committee. 2026. “Survey of North American Foreign Exchange Volume.” February 10, 2026. <https://www.newyorkfed.org/fxc/fx-volume-survey>.
- Forrester Research. 2026. *US Tech Forecast 2026: What It Means For Financial Services*. <https://www.forrester.com/report/us-tech-forecast-2026-what-it-means-for-financial-services/RES191756>.
- Garratt, Rod, and Michael Junho Lee. 2025. “The Future of Payment Infrastructure Could Be Permissionless.” *Liberty Street Economics*, November 25, 2025. <https://doi.org/10.59576/lse.20251125>.

- Geithner, Timothy F. 2014. *Stress Test: Reflections on Financial Crises*. Crown, May 12, 2014.
- Gennaioli, Nicola, Andrei Shleifer, and Robert Vishny. 2015. “Money Doctors.” *The Journal of Finance* 70 (1): 91–114. <https://doi.org/10.1111/jofi.12188>.
- Gensler, Gary. 2021. “Remarks Before the Aspen Security Forum.” August 3, 2021. <https://www.sec.gov/news/public-statement/gensler-aspen-security-forum-2021-08-03>.
- Greenwood, Robin, Robert Ialenti, and David Scharfstein. 2025. “The Evolution of Financial Services in the United States.” *Annual Review of Financial Economics* 17 (November 2025): 189–206. <https://doi.org/10.1146/annurev-financial-082123-105625>.
- Greenwood, Robin, and David Scharfstein. 2013. “The Growth of Finance.” *Journal of Economic Perspectives* 27 (2): 3–28. <https://doi.org/10.1257/jep.27.2.3>.
- Griffin, John M., and Kevin Mei. 2024. “How Do Crypto Flows Finance Slavery? The Economics of Pig Butchering.” Working Paper. SSRN, 2024. <https://doi.org/10.2139/ssrn.4742235>.
- Guiding and Establishing National Innovation for U.S. Stablecoins Act (GENIUS Act), Pub. L. No. 119-27 (2025). <https://www.congress.gov/bill/119th-congress/senate-bill/1582/text>.
- Halaburda, Hanna, Guillaume Haeringer, Joshua Gans, and Neil Gandal. 2022. “The Microeconomics of Cryptocurrencies.” *Journal of Economic Literature* 60 (3): 971–1013. <https://doi.org/10.1257/jel.20201593>.
- Heeb, Gina, AnnaMaria Andriotis, and Josh Dawsey. 2025. “Walmart and Amazon Are Exploring Issuing Their Own Stablecoins.” Markets. *Wall Street Journal*, June 13, 2025. <https://www.wsj.com/finance/banking/walmart-amazon-stablecoin-07de2fdd>.
- Huberman, Gur, Jacob D Leshno, and Ciamac Moallemi. 2021. “Monopoly without a Monopolist: An Economic Analysis of the Bitcoin Payment System.” *The Review of Economic Studies* 88 (6): 3011–40. <https://doi.org/10.1093/restud/rdab014>.
- Jones, Garrett, and Alex Tabarrok. 2025. *President Trump’s Stablecoin EO: What It Means for American Innovation, Crypto Taxation, and Global Dominance*. Mercatus Center at George Mason University. Webinar. <https://www.mercatus.org/economic-insights/event-videos/president-trumps-stablcoin-eo>.
- Kashyap, Anil K., Jeremy C. Stein, and Samuel G. Hanson. 2010. “An Analysis of the Impact of ‘Substantially Heightened’ Capital Requirements on Large Financial

- Institutions.” Mimeo. Harvard Business School, 2010.
<https://www.hbs.edu/faculty/Pages/item.aspx?num=41199>.
- Lampert, Leslie, Robert Shostak, and Marshall Pease. 1982. “The Byzantine Generals Problem.” *ACM Transactions on Programming Languages and Systems* 4 (3): 382–401. <https://doi.org/10.1145/357172.357176>.
- Levine, Matt. 2021. “Citi Can’t Have Its \$900 Million Back.” Bloomberg, February 17, 2021. <https://www.bloomberg.com/opinion/articles/2021-02-17/citi-can-t-have-its-900-million-back>.
- Lewis-Pye, Andrew, and Tim Roughgarden. 2023. “Permissionless Consensus.” arXiv:2304.14701. Preprint, arXiv, April 28, 2023.
<https://doi.org/10.48550/arXiv.2304.14701>.
- MacKenzie, Donald. 2021. *Trading at the Speed of Light: How Ultrafast Algorithms Are Transforming Financial Markets*. Princeton University Press, May 25, 2021.
- Makarov, Igor, and Antoinette Schoar. 2025. “Blockchain Analysis of the Bitcoin Market.” Working Paper. July 24, 2025. <https://doi.org/10.2139/ssrn.3942181>.
- McKinsey & Company. 2025. *The 2025 McKinsey Global Payments Report: Competing Systems, Contested Outcomes*. McKinsey & Company, September 26, 2025.
<https://www.mckinsey.com/industries/financial-services/our-insights/global-payments-report>.
- . 2026. *Stablecoins in Payments: What the Raw Transaction Numbers Miss*. January 23, 2026. <https://www.linkedin.com/pulse/stablecoins-payments-what-raw-transaction-numbers-4qjke/>.
- Mishra, Sanjay, and Alan Beaulieu. 2002. *Mastering Oracle SQL: Putting Oracle SQL to Work*. 1st Edition. O’Reilly, 2002.
- Modigliani, Franco, and Merton H. Miller. 1958. “The Cost of Capital, Corporation Finance, and the Theory of Investment.” *American Economic Review* 48 (3): 261–97.
- Nakamoto, Satoshi. 2008. “Bitcoin: A Peer-to-Peer Electronic Cash System.” 2008.
<https://bitcoin.org/bitcoin.pdf>.
- Olson, Mancur. 1971. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Second Printing. Harvard Economic Studies. Harvard University Press, 1971.
- Philippon, Thomas. 2015. “Has the US Finance Industry Become Less Efficient?” *American Economic Review* 105 (4): 1408–38. <https://doi.org/10.1257/aer.20120578>.

- Philippon, Thomas, and Ariell Reshef. 2012. “Wages and Human Capital in the U.S. Finance Industry: 1909–2006.” *The Quarterly Journal of Economics* 127 (4): 1551–609. <https://doi.org/10.1093/qje/qjs030>.
- President’s Working Group on Digital Asset Markets. 2025. *Strengthening American Leadership in Digital Financial Technology*. July 30, 2025. <https://www.whitehouse.gov/crypto/>.
- Rajic, Taylar, and Julia Brock. 2025. “The ByBit Heist and the Future of U.S. Crypto Regulation.” March 18. <https://www.csis.org/analysis/bybit-heist-and-future-us-crypto-regulation>.
- Schelling, Thomas C. 1956. “An Essay on Bargaining.” *The American Economic Review* 46 (3): 281–306.
- Securities Industry and Financial Markets Association. 2026a. “US Equity and Related Statistics.” January 5, 2026. <https://www.sifma.org/research/statistics/us-equity-and-related-securities-statistics/>.
- . 2026b. “US Fixed Income Securities Statistics.” January 9, 2026. <https://www.sifma.org/research/statistics/us-fixed-income-securities-statistics/>.
- . 2026c. “US Treasury Securities Statistics.” January 9, 2026. <https://www.sifma.org/research/statistics/us-treasury-securities-statistics/>.
- Uber Technologies, Inc. 2025. “Form 10-K.” U.S. Securities and Exchange Commission, February 14, 2025. <https://www.sec.gov/Archives/edgar/data/1543151/000154315125000008/uber-20241231.htm>.
- United Nations. 2024. *Casinos, Money Laundering, Underground Banking, and Transnational Organized Crime in East and Southeast Asia: A Hidden and Accelerating Threat*. January 15, 2024. <https://www.unodc.org/roseap/en/2024/casinos-casinos-cryptocurrency-underground-banking/story.html>.
- U.S. Bureau of Labor Statistics. 2023. “Occupational Employment and Wage Statistics.” 2023. <https://www.bls.gov/oes/home.htm>.
- U.S. Commodity Futures Trading Commission and U.S. Securities and Exchange Commission. 2010. *Findings Regarding the Market Events of May 6, 2010*. <https://www.sec.gov/news/studies/2010/marketevents-report.pdf>.
- U.S. Department of the Treasury, Board of Governors of the Federal Reserve System, Federal Reserve Bank of New York, U.S. Securities and Exchange Commission, and

U.S. Commodity Futures Trading Commission. 2015. *Joint Staff Report: The U.S. Treasury Market on October 15, 2014*. <https://www.sec.gov/files/treasury-market-volatility-10-14-2014-joint-report.pdf>.

Valukas, Anton R. 2010. *Report of the Examiner in the Chapter 11 Proceedings of Lehman Brothers Holdings Inc.* Jenner & Block LLP, March 11, 2010. <https://www.jenner.com/en/news-insights/news/lehman-brothers-holdings-inc-chapter-11-proceedings-examiner-s-report>.

Visa and Allium Labs. 2026. “Visa Onchain Analytics Dashboard.” <https://visaonchainanalytics.com/transactions>.