

Blockchain Technology in Traditional Finance

ERIC BUDISH AND ADI SUNDERAM¹

March 2024

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 *without* the novel trust model; that is, with trust instead anchored in traditional sources such as rule of law, reputations, relationships, and collateral. We divide our analysis into two parts. First, we analyze potential improvements for transactions that are already taking place in the absence of the technology. We identify three categories of improvement: (i) reducing real resource costs, (ii) improving balance sheet efficiency, and (iii) reducing intermediation rents. Second, we analyze the potential of 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. While the value of such improvements is hard to quantify, it could potentially cumulate to hundreds of billions per year.

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

1. Introduction

Much of finance amounts to record keeping. In order to invest and trade financial claims, all parties must agree on who owns what claims and when. Historically, a relatively small number of government and private sector entities have been responsible for record keeping, with reputation and the rule of law sustaining agreement across parties that the records are accurate.

The emergence of Blockchain technology opens new possibilities for record keeping. As embodied in cryptocurrencies like Bitcoin and Ethereum, this technology combines two novel features, a novel data structure and a novel trust mechanism that allows record keeping to be fully decentralized. Budish (2023) shows that manufacturing trust in a fully decentralized, anonymous way is very expensive, casting doubt on the usefulness of the trust mechanism for traditional finance. In this paper, we ask whether an idealized version of the Blockchain data structure, if supplemented with traditional forms of trust like reputation and rule of law, could generate value for end users of the traditional financial system.

We start by defining the scope of the exercise, clarifying exactly what we will assume about the data structure and which types of financial intermediation we will focus on. With these boundaries in place, we then divide our analysis into two parts, depicted in Figure 1 below. We first discuss potential gains in the set of transactions that are already taking place. We highlight three possibilities. First, an idealized data structure could simply reduce real resource costs of transactions, for instance by decreasing the number of people at financial institutions tasked with transactions clearing. Second, by speeding up clearing it could reduce the amount of balance sheet financial institutions have to dedicate to transactions in the clearing process. We note that this only generates aggregate gains if there are significant violations of the Modigliani-Miller (1958) theorem. Third, we discuss the potential for new technology to lower total markups incurred by end users in financial intermediation, noting that reducing marginal costs for incumbents does not necessarily affect markups.

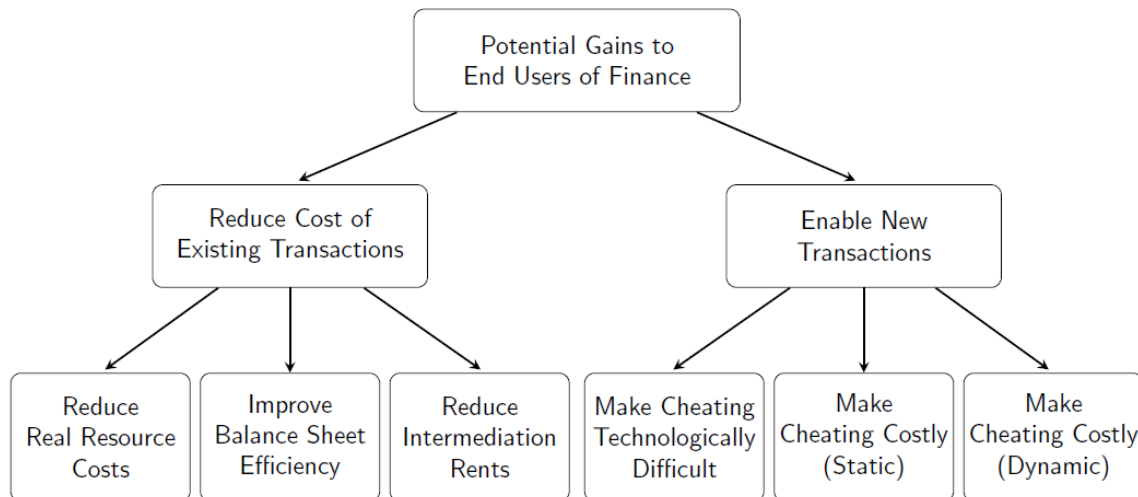


Figure 1. Framework for Analyzing Gains from Idealized Data Structure for Traditional Finance

We then turn to the potential for an idealized data structure to enable new transactions. A key insight is that if the transaction is not taking place today, the surplus it generates must be small compared to the potential loss if one of the parties cheats. Applying folk theorem logic from the theory of repeated games, this must be true on a forward-looking basis. In other words, the parties cannot expect to transact frequently in the future. An idealized data structure could help sustain such transactions in three ways. First, by simply making it more difficult to cheat. Second, by making it easier to punish a cheating counterparty in a static sense. Third, by making it easier to punish a cheating counterparty in a dynamic sense.

We take these ideas to a particular application of Blockchain technology: stablecoins, which are digital currencies recorded on distributed ledgers that are pegged to a reference value. Our discussion of stablecoins' impact on existing transactions draws on our general discussion of market power and markups. To reduce markups in traditional cash transfers, stablecoins must facilitate new entry rather than simply being adopted by incumbents. In addition, we note that in other parts of traditional finance, the cash leg of a transaction is typically not the limiting factor, potentially curbing the impact of stablecoins.

In our discussion of stablecoins' potential to enable new transactions, we focus on the potential for Blockchain technology and smart contracts to expand the set of assets used to back safe stores of value. To date, such efforts have been spectacularly unsuccessful. However, those failures are at least partially due to the fact that the only assets currently transacting on Blockchains are cryptocurrencies and other tokens, which are particularly poor backing for safe assets. A broader adoption of Blockchain technology may facilitate more interesting applications.

We close by noting that our analysis is essentially static in nature. That is, it asks what benefits the idealized data structure would create in the traditional financial system as constituted today. Drawing a historical analogy to the introduction of the relational database, we underscore that the long-run benefits of technology are very hard to forecast.

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

As discussed in Budish (2023), Blockchain technology as embodied in cryptocurrencies like Bitcoin and Ethereum is comprised of both a novel data structure and a novel trust model. The data structure is essentially a distributed database that only allows new data (transactions) to be appended and does not allow old data to be deleted or modified. The trust model (Nakamoto, 2008) is intuitively like a voting system, where the right to update the database depends on how much computational power the updater expends (in proof-of-work systems like Bitcoin) or on the amount the updater has invested in tokens governing the database (in proof-of-stake systems like Ethereum). As Budish (2023) points out, manufacturing trust in these fully decentralized, anonymous ways is expensive. The cost of maintaining the integrity of the system against outside attacks scales linearly with the economic value of the system. Indeed, a conservative calibration suggests

that securing the system against a \$10 billion attack relying on Nakamoto (2008) trust alone costs \$500 billion per year. This makes it unlikely that, as currently designed, the decentralized trust models embodied in cryptocurrencies can play a large role in traditional finance.

Moreover, the economic usefulness of anonymous blockchains has also been limited. To date, the majority of volume appears speculative, with the only other widely-documented use case being the facilitation of black market transactions (Makarov and Schoar, 2021; Foley et al., 2019; Cox, 2021; Gensler, 2021). Ironically, much of this transaction volume flows through cryptocurrency exchanges, exactly the kind of centralized financial intermediaries Bitcoin was designed to eliminate.

The Blockchain data structure itself, however, may prove useful. There are many potential uses for append-only distributed databases with well-defined permissions that adopt some of the features of existing Blockchain technology. For instance, one could imagine a distributed database of transactions that uses cryptographic signatures as identifiers, requires distributed consensus to add transactions, and enables smart contracts—algorithms that can automatically execute transactions on a distributed database.

Trust in such a database could be backed by traditional sources such as existing legal infrastructure, reputations, relationships, and collateral. For instance, though they might transact anonymously, the participants in this hypothetical Blockchain could be restricted to be regulated financial institutions and access to the database could require fulfillment of know-your-customer (KYC) and anti-money-laundering (AML) regulations. The goal in some sense would be to find a middle ground between existing transaction services, which are inaccessible to some parties and based on older technologies, and modern cryptocurrencies, which are too open and anonymous and thus expensive because they do not take advantage of the existing infrastructure for generating trust.

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

Our focus on the data structure aspects of Blockchain technology naturally circumscribes the types of financial intermediation we will discuss. In particular, we focus on the kinds of financial intermediation that are “primarily” about the data structure. 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. For instance, smart

contracts could potentially allow the automation of complex decisions that are currently made by human traders, brokers, and market makers.

The quantity of financial intermediation in the categories we consider is large. According to data from the Securities Industry and Financial Markets Association, total transaction volume in fixed income securities markets averaged nearly \$1 trillion per day in 2023. US equity market volume averaged about \$515 billion per day. In sum, this totals nearly \$375 trillion of volume per year. If we include trade in other financial assets such as real estate and commodities, there is easily \$500 trillion of trading volume per year of assets that shows up directly on financial institutions' balance sheets (this figure will be relevant for some of our calculations below). If we include assets such as futures and derivatives that do not appear directly on balance sheets, at least in their notional amounts, US financial system volume easily exceeds \$1 quadrillion per year.

3. Potential Gains on Existing Transactions

In this section, we consider a first broad source of gains to end users from an idealized data structure in traditional finance: reductions in costs for transactions that would have taken place in the absence of the technology. Such reductions can arise because the real resource costs of executing transactions fall or because the total markups charged by intermediaries for a transaction fall. Total markups may fall because the average markup charged by each intermediary in an intermediation chain falls, or because the number of links in the intermediation chain falls.

3.1 Reducing Resource Costs

Blockchain technology may help reduce the real resource costs of transactions in several ways.

3.1.1 Backoffice Expenses

The simplest one is by reducing the number of people involved in backoffice functions at financial intermediaries. Every large intermediary employs a significant

number of people in backoffice 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 4,300 people. 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 backoffice costs. The number of employees involved in such activities is difficult to determine, but to give a sense of the magnitudes involved consider that the financial services industry employed nearly 53,000 bookkeeping, accounting, and auditing clerks in 2021 according to the Bureau of Labor Statistics. Another 10,000 people were employed as information clerks. There are numerous other similar job categories. If technology reduced the number of such jobs by 100,000 and the average total compensation of employees in such jobs was \$200,000 per year, this cumulates to \$20 billion per year in cost savings.

3.1.2 Netting and Collateral Efficiency

Another potential source of gains is from more efficient netting. In the U.S., the typical equity transaction clears after two days; Treasury transactions take two days; corporate bond transactions can take up to five days, and loan transactions can take up to 21 days.

Discussions of gains from better clearing often focus on the gross gains and losses that arise because the party purchasing a security often sends cash to the party selling the security several days before it receives the security. For instance, suppose Citadel sends cash to Goldman Sachs, and Goldman Sachs sends the purchased bond to Citadel two days later, but does not pay interest on the cash to Citadel for those two days. This is a gain for Goldman and a loss for Citadel. The gross gains and losses are potentially large, especially given that interest rates have risen significantly in the past two years. The gross gains and losses, which equal the total interest on cash in limbo, are on the order of:

$$\text{Riskfree Rate} \times \frac{\text{Days to Clear}}{365} \times \text{Total Transaction Volume.}$$

This is potentially a very large number, given that total transaction volume per year is in the trillions of dollars. For instance, at 5% interest rates, if days to clear averages 2.0 on dollar-weighted terms, then this comes out to 3 basis points of all transaction dollars. Given our discussion of transaction volumes above, total transaction volume is easily \$500 trillion-\$1 quadrillion per year. Using \$500 trillion to be conservative, this totals nearly \$150 billion of interest payments per year.

Note, however, that this is a gross cost, not a net cost. In the example above, Citadel forgoes interest, but Goldman Sachs gains the interest. In other words, interest payments are transferred from one party to another, without any direct aggregate loss. However, \$150 billion is a reasonable quantification of the incentives to avoid gross losses. For instance, the desire to avoid forgone interest may lead Citadel to send cash to Goldman Sachs later than it otherwise would. In equilibrium, if all parties sending cash act this way, it may put operational strain on the system. Due to the circular flow of funds, some parties must wait to receive cash before sending it on to others. Thus, the desire to avoid gross losses may increase the risk of payment failures.

Inefficiencies in transactions clearing may also directly generate net aggregate losses, but these losses are harder to quantify. For instance, if there are violations of the Modigliani-Miller (1958) theorem, then Citadel's cost of financing for the cash it sends to Goldman may exceed the risk-free rate, so Goldman's gain does not fully offset Citadel's loss. These Modigliani-Miller violations can be socially costly, for instance because real resources are spent on trying to monitor financial intermediaries to mitigate agency problems. In this case, there is real deadweight loss associated with transactions clearing, which totals roughly

$$\text{MM cost} \times \frac{\text{Days to Clear}}{365} \times \text{Total Transaction Volume.}$$

Continue to assume that transactions clear in two days on average in dollar-weighted terms. Further suppose that Modigliani-Miller violations add a full percentage point to the cost of capital for financial intermediaries.¹ This is an aggressive (i.e., high) assumption relative to estimates in the literature for the largest global banks (e.g., Basel Committee on Banking Supervision (2010); Kashyap, Stein, and Hanson (2010); Admati et al (2013); Baker and Wurgler (2015); Sarin and Summers (2016); Federal Reserve Bank of Minneapolis (2016); and Firestone, Lorenc, and Ranish (2017)). However, it may be less aggressive for smaller intermediaries like individual hedge funds. Given the assumption, the total cost of these Modigliani-Miller violations per year is roughly 0.5-1 basis point. This is small in percentage terms, but it adds up to nearly \$75 billion per year when applied to the \$500 trillion of annual transaction volume discussed above.

3.2 Market Power and Markups

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.

A technological innovation that simply reduces the marginal costs of executing transactions need not have a big effect on total transactions costs if these costs are dominated by markups. Furthermore, in many settings, we think that monopoly or

¹ It is typical to think of the costs of Modigliani-Miller violations as a percentage of total balance sheet size as opposed to transaction volume. The assumption we are making here is that incremental balance sheet is tied up while transactions clear. For instance, suppose a hedge fund with a \$100 balance sheet trades \$1 every 2 days and every transaction takes 2 days to clear. Then \$1 of the fund's balance sheet is always consumed by transactions and subject to the cost of Modigliani-Miller violations. In this case, the fund's total transaction volume is $\$365/2$, so the formula above says that MM costs should be applied to 2 days to clear/ $365 * \$365/2 = \1 of balance sheet.

oligopoly potentially accelerates technological innovation and adoption because individual firms capture more of the gain from the innovation in such markets than they would in a competitive market. Thus, the likelihood of large untapped reductions in marginal costs seems low, and even if present and realized they might not translate to large reductions in transactions costs absent new entry.

As is always the case, the passthrough of marginal costs to prices depends on the competitive landscape and the elasticity of demand for the services provided. It is quite possible that in highly centralized parts of financial intermediation, lower marginal costs simply increase the profits of intermediaries. For instance, if clearing derivatives transactions has properties that make it a natural monopoly (either due to regulation or its intrinsic nature), a new technology that lowers the cost of clearing does not automatically make the market more competitive. There is a good chance that market participants simply face the same monopolist, which now has lower marginal costs.

For technology to have a significant impact on markups then, it must impact the competitive landscape, for instance by lowering barriers to entry. This would not be the case if, for example, individual intermediaries clear transactions on their own individual Blockchains or a consortium of intermediaries jointly manages a single Blockchain, the way the Depository Trust and Clearing Corporation is currently run. By contrast, if new technology enables a standardized Blockchain that intermediaries can freely enter and compete to clear transactions on, then mark ups could fall significantly. The idea is that concerns about trust or data integrity are a significant impediment to competition in the current world. As a result, markups remain high and accrue to a few large, trusted financial intermediaries. If a new technology can solve this trust problem, then it can significantly increase the degree of competition in financial intermediation, pushing markups closer to zero.

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

and Scharfstein (2013)). These numbers include the costs of recordkeeping, which are our focus, but also the costs of financial decision making discussed in Section 2.2. Thus, new recordkeeping technologies would not reduce the costs of intermediation to zero, but could still generate hundreds of billions of dollars of value per year.

At present, it is hard to see the path to gains of this magnitude for several reasons. First, aggregate historical evidence suggests markups are hard to reduce. As Philippon (2015) notes, the cost of intermediation has been relatively constant over time, despite large changes in the technology and market structure of financial intermediation. Second, microeconomic evidence from modern markets suggests that markups remain high in some areas of finance that are relatively commoditized and competitive, like retail deposit taking. The evidence suggests these markups are high, not because depositors lack options, but because they are quite inelastic, possibly due to high switching costs for moving deposits from one bank to another. After all, there are already many alternatives available to depositors who wish to obtain higher interest rates, and depositors remain at banks that pay low rates despite these alternatives. For instance, Egan, Hortacsu, and Matvos (2017) and Egan, Lewellen, and Sunderam (2021) find a 1 percentage point increase in deposit rates increases a bank's market share only 1-2 percentage points. A third issue is whether innovations in data structure are enough to significantly change the competitive landscape if not paired with an innovative trust model. As Budish (2023) suggests, the trust model associated with Bitcoin is too expensive to accommodate traditional finance. But existing trust models (e.g., regulation, rule of law, reputation) paired with novel data structures may not significantly change the competitive landscape.

Put differently, the impact of a new technology on markups depends on the microeconomic drivers of high concentration in traditional finance. For instance, if concentration is a result of high fixed costs to start an intermediary, then a technology that lowers those fixed costs would potentially have a big impact. On the other hand, if concentration is driven by trust—e.g., the reason that JP Morgan has a large share in many markets is that many parties trust JP Morgan—then the key question is whether Blockchain

technology augmented with traditional trust mechanisms changes the economies of scale associated with being a trusted counterparty.

One way in which technology could alter the competitive landscape is by making more feasible the provision of government services or “public options.” For instance, many advocates of central bank digital currencies (CBDCs) point to the fact that retail depositors typically earn very low interest rates with rents accruing to bank shareholders instead. One solution to this problem, practiced in many parts of the world, is postal banking: the government operates a system of banks, often co-located with post offices, that pays higher rates on deposits. In other words, the government sets a competitive floor on the market by entering itself. One drawback of such systems is operational complexity. It may be challenging for a government entity to effectively manage the large number of properties, employees, and other details associated with running a large national system of postal bank branches.

Technology may substantially reduce these operational challenges. Rather than running thousands of bank branches, the government essentially needs to run a website that provides online banking services, a less operationally complex endeavor.²

4. Enabling New Transactions

In this section, we consider a second broad source of gains from an idealized data structure in traditional finance: the enabling of new transactions that would not have taken place in the absence of the technology. We ask whether new technology can enable cooperation between parties who do not transact frequently. In such cases, long-run considerations are not sufficient to incentivize the parties not to cheat.

4.1 Facilitating Cooperation

² Of course, other challenges that are not linked to technology remain if governments are to offer retail banking services. One important issue is determining what assets should back deposits in the government bank—if the bank grows large, it may need to purchase assets beyond Treasury securities. As discussed above, another issue is whether the government bank would attract enough depositor interest to place competitive pressure on private banks in the first place, given that depositors appear to be quite inelastic.

A simple way to understand the cooperation problem is presented in the payoff matrix in Figure 2 below, which portrays an augmented Prisoners' Dilemma. The two players in the game are financial intermediaries considering a transaction. There are three options. Each intermediary can choose to engage in the transaction or not. If either chooses not to engage, they both get zero payoffs. If both choose to engage, they then each have the option to cooperate or cheat.³ 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 , and the player

		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

Figure 2. Payoffs in Augmented Prisoners' Dilemma.

Note: $f > 0$ represents the gains from trade in an honest transaction. $V > 0$ represents the gain to a cheating party from cheating in the transaction. In traditional finance the relevant case is 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.

³ While we represent the game as a static game with three actions, the game will be equivalent for our purposes to an extensive form game in which in a first stage the players choose engage or not, and then, if both choose engage, there is a second stage in which the players play the standard two-action prisoners' dilemma.

that got cheated gets a large negative payoff, $-V$. (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 $-\varepsilon$.)

While the game is a prisoners' dilemma for any parameters $V > f > 0 > -\varepsilon$, we have in mind the case where f is much smaller than V . Intuitively, the net benefit of a successful transaction, captured by f , is much smaller than the benefit of cheating one's counterparty in the same transaction, captured by V . We can think of f as on the order of the transaction fee the parties would be willing to pay to engage safely in the transaction, while we can think of V as the size of the transaction itself.

For instance, consider a transaction for a \$100 bond. The net surplus from the transaction might be $f = \$1$ because the transaction moves the bond to an owner that values it more. However, suppose the party that is supposed to sell the bond cheats. Then it keeps the bond worth \$100 and keeps the cash sent by the other party, also worth \$100. Thus $V = \$100$.

Note that we take cheating here to be a broad category, including everything from backoffice mistakes and defects in the legal system (e.g., which jurisdiction does the transaction take place in) to willful acts of theft.

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 a socially valuable transaction does not take place because of the risk of cheating. Cooperate-cooperate is not an equilibrium of the one-shot version of the game because Cheat is a best-response to Cooperate, just like in the standard one-shot prisoners' dilemma.

Now consider repeated play of this game. By standard folk-theorem logic, 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

$$\frac{\delta}{1-\delta} Nf > V$$

This observation might explain why parties that repeatedly transact, Goldman Sachs and JP Morgan for instance, do not have trouble sustaining cooperation. In contrast, parties that transact only infrequently, for instance a bank in Poland and a bank in Brazil, might not be able to.

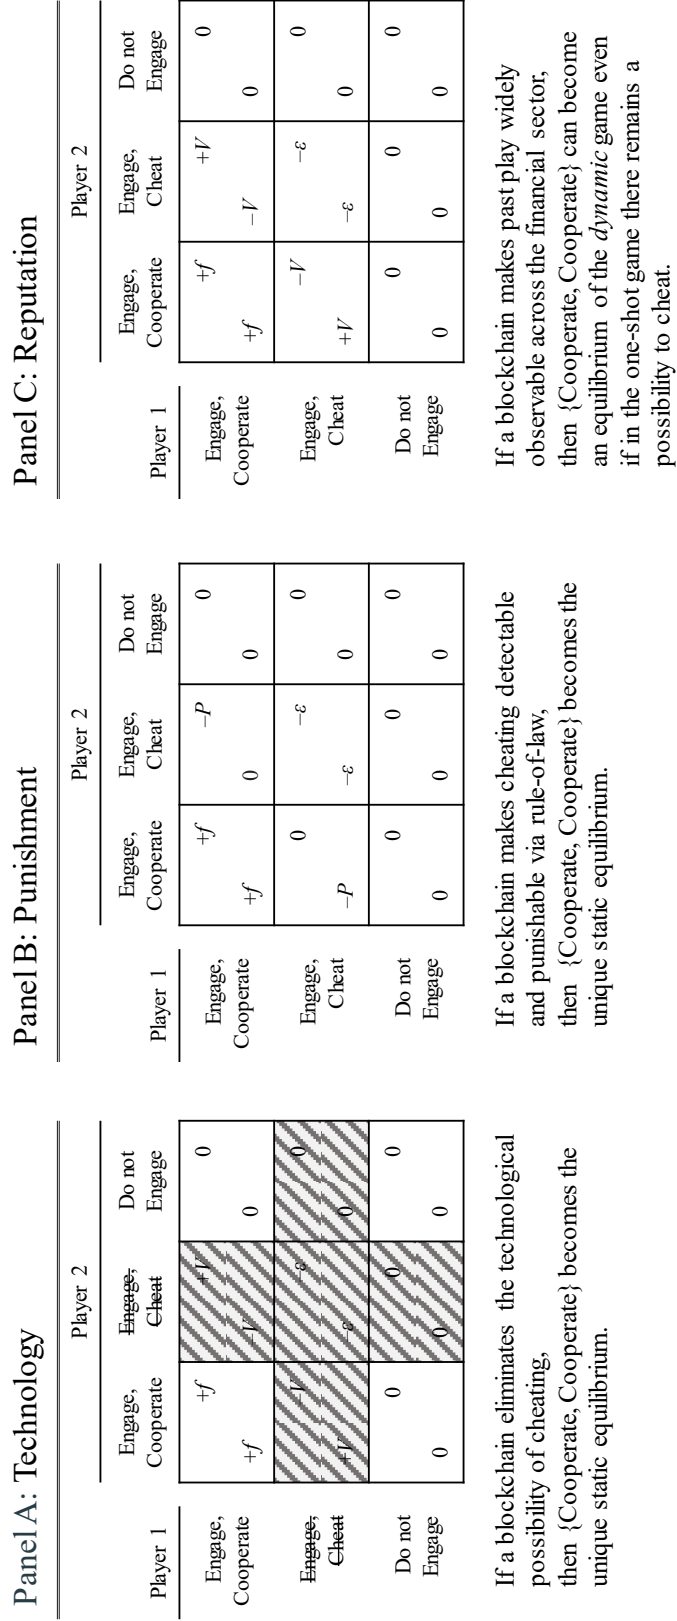
How could an idealized data structure help coordinate on the cooperate-cooperate outcome? We describe three ways, depicted in Figure 3 below.

First, it can change the possible actions in the static game, eliminating the scope for cheating (Figure 3, Panel A). For instance, it could require that the transaction that transfers the bond from the Polish bank to the Brazilian bank be paired and simultaneously cleared with the transaction that transfers cash from the Brazilian bank to the Polish 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.

Second, an 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). We can think of the cheat-cooperate payoff as the payoff net of any enforcement actions and penalties that can be imposed on the cheating bank. Having a public, indisputable data structure may make it easier to impose such penalties and lower the net payoff to cheating, even driving it negative if detection and enforcement are good enough. In other words, an idealized data structure could change the payoffs in the cheat-cooperate cell from $(V, -V)$ to $(-P, 0)$ where P is a legal penalty. This again makes {Engage, Cooperate} - {Engage, Cooperate} the unique equilibrium even in one-shot play.

A third possibility is that, by making intermediary behavior more easily observable and verifiable, the data structure alters the possibilities for cooperation in the repeated game (Figure 3, Panel C). Even if payoffs do not change in the static game, having a public, indisputable data structure may make it easier to impose dynamic penalties that make cooperation easier to sustain. For instance, it could become easier for other financial institutions not party to the transaction to exclude the cheating bank from participating in

Figure 3. Three Ways a Blockchain Can Facilitate Cooperation



other markets. Mathematically, add a parameter M that represents the number of financial institutions using the 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 is

$$\frac{\delta}{1-\delta} MNf > V$$

This is like enlarging the frequency-of-interaction parameter N from the number of times a party interacts with one other counterparty per year, to the number of times a party interacts with *all* other counterparties using the data structure per year. The cost of cheating becomes the cost of exclusion from the entire system, making cooperation more sustainable.

Discussion: Repeated-Game Analysis

A subtle lesson that emerges from this analysis is that the transactions that are not taking place in the current system cannot be transactions that are either high-surplus (high f) or between counterparties that frequently interact (high N). In other words, the missing transactions must be relatively infrequent and relatively low surplus. This does not mean that better technology cannot generate significant value. It could be the case that there are a very large number of such transactions, so that in total the value lost from existing data structures is large.

Discussion: Relation to Verification Problems

A related problem that the idealized data structure could help with is verification, i.e., the problem one party faces in establishing information about the counterparty it is transacting with. For instance, Goldman Sachs may be willing to lend to a small hedge fund with a bond serving as collateral. To do so, however, it must verify that the hedge fund actually has the bond and that it has not pledged the bond as collateral to any other lenders. This is a particular kind of cheating in a transaction.

More broadly, it is potentially valuable for financial institutions to be able to credibly disclose their net worth to one another, thereby establishing a broad notion of

creditworthiness, without disclosing their entire portfolios. An idealized data structure, coupled with so-called zero knowledge proofs, could facilitate such disclosure. Creating a new technology for verifiable disclosure is not the same thing as the ledger but it is a close cousin and seems potentially valuable.

4.2 Eliminating Excess Intermediation

There is an intermediate case between the extensive margin of transactions discussed here and the intensive margin discussed in Section 3. It could be the case that transactions are taking place but indirectly through a trusted intermediary. For instance, in the example above, the Polish bank and Brazilian bank may not be willing to transact directly, but they can still interact through JP Morgan. In this case, an improved data structure could help eliminate excess layers of intermediation. To the extent there are real resource costs of excess intermediation—that is JP Morgan expends real resources on serving as a go-between, rather than simply charging each involved party, then aggregate gains are possible. If there are not real costs, eliminating excess intermediation could still facilitate new transactions by reducing total markups charged along the intermediation chain.

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 US 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, they have grown rapidly, capturing the attention of both cryptocurrency market participants and regulators. This growth reflects a belief that stablecoins could significantly increase the efficiency of the payments system, in addition to serving a critical role in decentralized finance and cryptocurrency markets (e.g., President’s Working Group on Financial Markets et. al. 2021; Adachi et. al. 2022). Second, stablecoins are a natural

application for our framework as, by virtue of their stability, the primary concern with stablecoins is recordkeeping. The third reason stablecoins are a useful setting is that, as we detail below, our framework suggests that they are unlikely to generate much value for traditional finance. This illustrates our framework has some value, as it suggests large gains in some places but not in others.

We start by more carefully elucidating similarities and differences between stablecoins and stable value claims in traditional finance. We discuss potential gains from stablecoin adoption for transactions already occurring and in enabling new transactions.

5.1 Stablecoins and Stable Value Claims in Traditional Finance

There are a variety of stable value claims in the traditional financial system ranging from central bank reserves and Treasury securities to bank deposits and shares in money market mutual funds. From the perspective of their owners, these claims serve purposes that can broadly be categorized as serving either as a store of value or a medium of exchange. Claims frequently mix these purposes. Central bank reserves serve as the formal medium of exchange, but bank deposits for instance serve depositors both as a store of value and as a medium of exchange, with banks handling the transfer of the formal medium of exchange (reserves) on behalf of depositors. Similarly, while Treasury securities typically directly serve only as a store of value, the fact that they are liquid—i.e., easily convertible into the formal medium of exchange—contributes to their value.

Stablecoins are also broadly intended to fill the roles of serving as stores of value and mediums of exchange. To date, there are three main types of stablecoins. First, there are stablecoins that are backed by fiat currency (e.g., dollar) reserves. These include Tether and USDC and vary in the quality of and disclosure about their reserves. For instance, USDC is fully backed by deposits and Treasuries, while Tether is backed by a broader, and not fully disclosed, set of collateral.⁴ Second, there are stablecoins backed by other

⁴https://www.circle.com/hubfs/USDCAttestationReports/2023/2023%20USDC_Circle%20Examination%20Report%20May%202023.pdf;

cryptocurrencies. For instance, Dai is backed by Ethereum and other crypto assets. Because these assets are volatile, such stablecoins are typically overcollateralized—holding substantially more assets than coins issued. Third, there are algorithmic stablecoins that are partially or fully backed by their own investment tokens. This category includes the infamous Terra-Luna collapse.

5.2 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 gains: (i) reductions in real resource costs like back office costs, (ii) improvements in collateral and netting efficiency, and (iii) reductions in intermediation rents.

Stablecoins could offer meaningful reductions in real resource costs and improvements in netting efficiency. One crucial distinction between stablecoins and stable value claims in traditional finance is transferability and fungibility. In traditional finance, stable value “money-like” claims are transferred in what is essentially a two-step process: first they are converted back into the formal medium of exchange and then they are transferred. For instance, suppose Adi banks at Bank of America and wishes to send \$100 to Eric, who banks at Charles Schwab. There are three separate ledgers involved in this transaction. The Federal Reserve’s ledger keeps track of the reserve balances of Bank of America and Charles Schwab. Bank of America’s ledger keeps track of its reserves and Adi’s deposit balance, and Charles Schwab’s ledger keeps track of its reserves and Eric’s deposit balance. To complete the transaction, Bank of America sends \$100 of reserves to Charles Schwab and deducts \$100 from Adi’s deposit account on its ledger. The Federal Reserve shifts \$100 of reserves from Bank of America to Charles Schwab on its ledger. And Charles Schwab records its increase in reserves and a commensurate increase in Eric’s deposit account on its ledger.

https://assets.ctfassets.net/vyse88cgwfb/24G4DuQ0HE7h7EQE6vGy4J/8a8a170edf687ea07b3f86048af8b87b/ESO.03.01_Std_ISAE_3000R_Opinion_31-03-2023_BDO_Tether_CRR.pdf

In contrast, stablecoins do not require conversion back to the underlying medium of exchange in a transaction. In the example above, Adi would simply transfer his stablecoins to Eric, with the transaction being recorded on the relevant Blockchain ledger. No other ledgers would be involved. This reduction in ledgers could reduce both real resource costs—for instance, fewer ledgers require fewer back office employees—and improve netting efficiency by increasing clearing speeds. These are real gains, but they are likely to be small for stable value claims compared to the potential gains from more efficient clearing for complex financial securities.

It is also possible that stablecoins could reduce intermediation rents. A standardized Blockchain could facilitate entry and competition among stablecoin providers, causing mark ups on cash transactions already taking place to fall significantly. A difficulty that arises with this argument is that it is unclear whether new technology will facilitate competition. With current technology, there are large economies of scale in transactions processing, which arise both from physical technological costs and from thick markets externalities—i.e., a transaction medium becomes more attractive when more people are using it. New technology will likely only facilitate competition if it does not feature these kinds of economies of scale. Indeed, one could describe the traditional banking system with the same terms that stablecoin proponents use to describe their technological vision. The government operates a standardized protocol (e.g., Fedwire) to help transfer funds. It does not give individuals access to this protocol, but instead it allows private firms (i.e., banks) to compete to offer individuals services that use the protocol. Entry is not completely unfettered—there are regulatory costs of entry, which help protect the trustworthiness of the system, but it is relatively free, as evidenced by the large number of banks in the world. Despite this competition, markups remain high in traditional banking, in part because consumers are quite inelastic as discussed above.

Potential Limits

There are two main limits to potential gains from stablecoins for transactions already taking place. First, stablecoins essentially only offer efficiencies on the cash side

of a transaction. Transactions typically involve one party sending another cash, while the second party sends a security or other financial asset to the first. In many parts of traditional finance, the inefficiencies in clearing are bigger on the security side of the transaction than the cash side. Thus, the improvements stablecoins offer may not much change the aggregate number of people or aggregate quantity of balance sheet tied up in transactions clearing. Furthermore, there are many forthcoming or potentially forthcoming institutional changes that will make the cash side of the transaction more efficient in the absence of stablecoins. For instance, central banks around the world have been adopting real-time payment systems. Such systems are substitutes for stablecoins in many respects, though they leave a layer of intermediation (banks) between final customers and the central banks, potentially leaving unchanged large markups on cash clearing.⁵

Second, as highlighted above, a key potential source of efficiencies from stablecoins is that their transferability may reduce the number of ledgers involved in a transaction. To the extent that Blockchain and other data technologies reduce the costs of changing all ledgers, then the relative value of this transferability actually falls. In other words, stablecoins may primarily be valuable in the interim period between our current world and a world where new data technologies are fully adopted. The same is true of the stable value claims that currently exist in traditional finance to “fill gaps” in transaction clearing processes. If transactions clearing becomes much faster and more efficient, then the use of stopgap solutions like intraday credit will likely shrink.

5.3 Stablecoins Enabling New Transactions

How might stablecoins and Blockchain technologies more broadly enable new transactions? Our framework in Section 4 suggests three possibilities: (i) eliminating the scope for cheating by changing the action set, (ii) making it easier to detect cheating and

⁵ See, for example, the section on Central Bank Digital Currency in <https://www.federalreserve.gov/publications/files/money-and-payments-20220120.pdf>. The document explicitly describes an “Intermediated” model of CBDC. “An intermediated model would facilitate the use of the private sector’s existing privacy and identity-management frameworks; leverage the private sector’s ability to innovate; and reduce the prospects for destabilizing disruption to the well-functioning U.S. financial system.”

impose penalties, lowering the static payoff to cheating, and (iii) making it easier to impose dynamic penalties that make cooperation easier to sustain.

Stablecoins and blockchain technology could enable new transactions through all three channels. For instance, one could imagine a smart contract that serves the purpose of a bank while using a broader set of assets to back deposits than a traditional bank. Consider a bank that as its assets owns \$100 of S&P 500 equity and puts on \$100 of S&P 500 with a strike price of \$90. This bank can potentially back \$90 of deposits with these assets. A slightly more sophisticated version of the bank would not own the puts directly but instead dynamically trade the S&P 500 to replicate the puts. Such a bank is discouraged by regulations today—it would face very high capital requirements—in part due to asymmetric information: it is hard to verify the bank manager is properly following the strategy and thus the risk that the bank manager may cheat discourages such transactions in the first place. Indeed, the convertibility of deposits back into cash has historically been thought of as important mechanism to discipline bank managers and prevent them from mishandling the funds entrusted to them (e.g., Calomiris and Kahn 1991, Diamond and Rajan 2001). With a public smart contract and public ledger, however, it may be possible to relax these incentive constraints and expand the activities of banks. In other words, a smart contract can change the action set of bank managers in a way that expands the supply of stores of value.

Even without a smart contract, a public ledger may help enable new transactions in banking. For instance, it could help both depositors and regulators monitor the bank's actions and impose static or dynamic penalties for cheating.

Potential Limits

To date of course, the history of stablecoins has not been kind to the argument that new technology can expand the set of assets used to back bank deposits. Many of the most spectacular depegging events, including the Terra/Luna episode, featured stablecoins backed by nonstandard collateral. In these cases, the main asset backing the stablecoin was essentially the firm's own equity (Terra coins) instead of the S&P 500. Even if the equity

has significant fundamental value, this leaves the firm particularly open to the risk of self-fulfilling runs, as significant redemptions reduce the value of the equity, which in turn triggers further redemptions. A traditional bank could also hold its own equity as an asset backing deposits, but it is well understood that doing so would have poor financial stability properties. The hypothetical S&P 500 bank need not face such issues as the value of its assets is not directly tied to the quantity of redemptions. However, Blockchain technology and smart contracts must be more broadly adopted for the hypothetical bank to be feasible—a smart contract must be able to execute trades in S&P 500 stocks.

5.4 Summary of Stable Coin Analysis

Overall, our analysis of stablecoins yields two conclusions. With the current penetration of Blockchain technology into the broader financial system, stablecoins likely must be fully reserved by truly safe claims (i.e., central bank reserves or short-term Treasuries) to maintain stability. In this case, their main benefit stems from their transferability, which reduces the number of ledgers involved in a transaction. As Blockchain and other technological innovations are adopted, this benefit is likely to decline relative to other available alternatives, reducing the utility of stablecoins backed by completely riskless claims. But at the same time, the scope increases for more fundamental innovation in financial intermediation, for instance the creation of a smart-contract bank that issues stablecoins (deposits) backed by risky assets. Such a bank would of course still be subject to traditional bank regulation, but technology essentially expands the production possibilities frontier for safe assets.

6 Static versus Dynamic Considerations

It is worth noting that the analysis here has primarily been static, essentially asking what net benefits Blockchain technology might have in the current financial system. In a more dynamic sense, the transparency and credibility of an idealized data structure might enable innovation and competition that does create a lot of value and efficiency, even if that is hard to forecast today.

For instance, the relational database, which today is a ubiquitous data storage structure, 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.

An additional lesson emerges from this example: the adoption of the relational database was largely driven by the private sector. Coordination was necessary—for instance, the details of the SQL language needed to be standardized for the industry to grow. However, this coordination was largely achieved by the private sector in conjunction with nongovernmental organizations like the International Organization for Standardization and American National Standards Institute. Regulators presumably needed to approve (or at least not disapprove) of the deployment of the technology within large, regulated financial institutions. But these steps all happened incrementally as benefits of the technology and its applications became clear.

7 Conclusion

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

Our analysis of real improvements in existing transactions identifies several potential sources of value (Section 3.1). 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. This is arguably 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 \$0.5 - 1 trillion with significant estimation uncertainty in both directions. This is meaningful in comparison to cryptocurrencies' peak market capitalization of \$3

trillion, reached in 2021. 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 3.2), in facilitating useful new transactions that are currently not taking place because of trust problems (Section 4.1), or in reducing excess intermediation for existing transactions that currently take place through a long intermediation chain (Section 4.2). These discussions are necessarily more speculative.

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.

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 in traditional finance. This light is likely to spur innovation and modernization among traditional financial institutions and central banks.

References

- Abadi, Joseph, and Markus K. Brunnermeier. 2022. "Blockchain Economics." Federal Reserve Bank of Philadelphia Working Paper.
- Adachi, Mitsu, Pedro Bento Pereira Da Silva, Alexandra Born, Massimo Cappucio, Stephanie Czak-Ludwig, Isabella Gschossmann, Georg Paula, Antonella Pellicani, Sarah-Maria Philipps, Mirjam Plooij, Ines Rossteuscher and Pierfrancesco Zeoli, 2022. "Stablecoins' role in crypto and beyond: functions, risk and policy". ECB Macprudential Bulletin 18.
- 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.
- Baker, Malcolm, and Jeffrey Wurgler. 2015. "Do Strict Capital Requirements Raise the Cost of Capital? Bank Regulation, Capital Structure, and the Low-Risk Anomaly." *American Economic Review*, 105 (5): 315-20.
- Basel Committee on Banking Supervision. 2010. *Basel III: A Global Regulatory Framework for More Resilient Banks and Banking Systems*. Basel: Bank for International Settlements. <http://www.bis.org/publ/bcbs189.pdf>
- Budish, Eric. 2023. "Trust at Scale: The Economic Limits of Cryptocurrencies and Blockchains." Becker Friedman Institute Working Paper.
- Calomiris, Charles W., and Charles M. Kahn. 1991. "The Role of Demandable Debt in Structuring Optimal Banking Arrangements". *The American Economic Review*, 81(3), 497–513.
- Cox, Jeff. 2021. "Yellen Sounds Warning About 'Extremely Inefficient' Bitcoin." CNBC. Last modified February 22, 2021. Retrieved May 1, 2022 from <https://www.cnbc.com/2021/02/22/yellen-sounds-warning-about-extremely-inefficient-bitcoin.html>.
- Diamond, Douglas W., and Raghuram G. Rajan. 2001. Liquidity Risk, Liquidity Creation, and Financial Fragility: A Theory of Banking. *Journal of Political Economy*, 109(2), 287–327.
- 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.
- Egan, Mark, Stefan Lewellen and Adi Sunderam. 2021. "The Cross-Section of Bank Value," *The Review of Financial Studies*, vol 35(5), pages 2101-2143.

- Federal Reserve Bank of Minneapolis. 2016. “The Minneapolis Plan to End Too Big to Fail.” <https://www.minneapolisfed.org/publications/special-studies/endingtbtft/the-minneapolis-plan-to-end-too-big-to-fail>
- Firestone, Simon, Amy Lorenc and Ben Ranish. 2017. “An Empirical Economic Assessment of the Costs and Benefits of Bank Capital in the US,” Finance and Economics Discussion Series 2017-034. Washington: Board of Governors of the Federal Reserve System, <https://doi.org/10.17016/FEDS.2017.034>.
- Foley, Sean, Jonathan R. Karlsen, and Talis J. Putnin, 2019. “Sex, Drugs, and Bitcoin: How Much Illegal Activity is Financed through Cryptocurrencies?” *The Review of Financial Studies*, 32(5): 1798–1853.
- Gensler, Gary. 2021. “Remarks Before the Aspen Security Forum.” Last modified August 3, 2021. Retrieved May 1, 2022 from <https://www.sec.gov/news/public-statement/gensler-aspen-security-forum-2021-08-03>.
- Greenwood, Robin, and David Scharfstein. 2013. "The Growth of Finance." *Journal of Economic Perspectives*, 27(2): 3–28.
- Kashyap, Anil K., Jeremy C. Stein, and Samuel Hanson. 2010. “An Analysis of the Impact of ‘Substantially Heightened’ Capital Requirements on Large Financial Institutions.”
- Makarov, Igor, and Antoinette Schoar. 2021. “Blockchain Analysis of the Bitcoin Market.” NBER Working Paper No. 29396.
- Modigliani, Franco, and Merton H. Miller. 1958. “The Cost of Capital, Corporation Finance, and the Theory of Investment.” *American Economic Review* 48, no. 3: 261–97.
- Nakamoto, Satoshi. 2008. “Bitcoin: A Peer-to-Peer Electronic Cash System”. Available at <https://bitcoin.org/bitcoin.pdf>
- Philippon, Thomas. 2015. “Has the US Finance Industry Become Less Efficient?” *American Economic Review*, 105(4): 1408-38.
- President’s Working Group on Financial Markets, the Federal Deposit Insurance Corporation, and the Office of the Comptroller of the Currency, Report on Stablecoins, 2021, https://home.treasury.gov/system/files/136/StableCoinReport_Nov1_508.pdf.
- Sarin, Natasha, and Lawrence H. Summers. 2016. “Understanding Bank Risk through Market Measures.” *Brookings Papers on Economic Activity*, 57–109.
- Stein, Jeremy C. “Monetary Policy as Financial Stability Regulation.” *The Quarterly Journal of Economics*, 127(1): 57–95
- Sunderam, Adi. 2015. “Money Creation and the Shadow Banking System.” *Review of Financial Studies* 28, no. 4: 939–77.