

This PDF includes a chapter from the following book:

Distributed Ledgers

Design and Regulation of Financial Infrastructure and Payment Systems

© 2020 Massachusetts Institute of Technology

License Terms:

Made available under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International Public License

<https://creativecommons.org/licenses/by-nc-nd/4.0/>

OA Funding Provided By:

The open access edition of this book was made possible by generous funding from Arcadia—a charitable fund of Lisbet Rausing and Peter Baldwin.

The title-level DOI for this work is:

[doi:10.7551/mitpress/13382.001.0001](https://doi.org/10.7551/mitpress/13382.001.0001)

3

Ledgers as Financial Accounts

Here we focus on the most obvious component of distributed ledger technology, namely, the ledgers themselves. The unique advantage of DLT as a ledger is that it can be held in common and shared. As a corollary, DLT provides an additional accounting check beyond double-entry bookkeeping on the reliability of recorded transactions. Furthermore, links of DLT ledgers to financial accounts open up a vision for future innovations that could have great power.

3.1 Statement of Cash Flow and Balance Sheet as a Ledger: From Paper Currency to Distributed Ledgers in a Few Steps

We make the link immediately to standard accounting concepts. Cash and paper currency transactions can be recorded on ledgers. For village economies measured in the Townsend Thai project, this is done as in Samphantharak and Townsend (2009), where the statement of cash flow as a standard corporate account is created, along with the stocks recorded in the balance sheet. Tables 3.1, 3.2, and 3.3 provide examples, including the income statement as well.

More specifically, a transaction log operates on the Townsend Thai monthly survey data (<http://townsend-thai.mit.edu>) and records cash transactions that each household i has with any

Table 3.1
Comprehensive financial accounts: balance sheet of household A.

Assets, liabilities, and net worth for a featured household over months 5 through 9					
Month	5	6	7	8	9
Cash in hand	1,966,139	1,862,121	1,701,863	1,663,257	1,593,938
Account receivables	688,971	805,259	952,359	1,059,382	1,126,773
Deposits at financial institutions	167,271	167,969	168,094	156,799	157,474
ROSCA (net position)	33,000	37,000	41,000	11,500	16,050
Other lending	153,136	153,136	153,136	153,136	153,136
Inventories	1,346,939	1,440,729	1,576,481	1,697,413	1,842,527
Livestock	326,280	323,018	319,787	316,590	313,424
Fixed assets	967,342	973,759	970,949	968,151	965,365
Household assets	598,758	596,261	593,775	591,299	588,833
Agricultural assets	66,104	65,829	65,554	65,281	65,009
Business assets	2,479	11,669	11,620	11,572	11,523
Land and other fixed assets	300,000	300,000	300,000	300,000	300,000
Total assets	5,649,079	5,762,991	5,883,669	6,026,228	6,168,687
Total liabilities	1,132,310	1,280,270	1,425,465	1,570,660	1,715,855
Account payables	1,078,505	1,228,465	1,375,660	1,522,855	1,670,050
Other borrowing	53,805	51,805	49,805	47,805	45,805

(continued)

Table 3.1 (continued)

Month	5	6	7	8	9
Total wealth	4,516,769	4,482,721	4,458,204	4,455,568	4,452,832
Initial wealth	3,439,250	3,439,250	3,439,250	3,439,250	3,439,250
Cumulative net gifts received	-6,664	-6,046	-6,357	-6,319	-7,576
Cumulative savings	1,084,182	1,049,517	1,025,311	1,022,637	1,021,158
Total liabilities and wealth	5,649,079	5,762,991	5,883,669	6,026,228	6,168,687

The unit of currency is THB. Month 5 is corresponding to January 1999.

Source: Samphantharak and Townsend (2009).

Table 3.2

Comprehensive financial accounts: income statement of household A.

Revenue, expenses, net profits, and disposition into consumption and savings					
Month	5	6	7	8	9
Revenue from cultivation					
Revenue from livestock	30,485	27,753	26,180	21,780	26,730
Livestock produce	28,985	27,753	26,180	21,780	26,730
Capital gains	1,500				
Revenue from fish and shrimp					
Revenue from business	184,360	145,360	183,875	152,890	160,455
Revenue from labor provision	11,440	11,440	11,440	11,440	11,440
Other revenues	6,000	3,000	6,000	6,000	6,000
Total revenues	232,285	187,553	227,495	192,110	204,625

(continued)

Table 3.2 (continued)

Month	5	6	7	8	9
Cost of cultivation					
Cost of livestock	31,944	30,281	27,642	22,813	21,715
Capital losses					
Depreciation (aging)	3,281	3,263	3,230	3,198	3,166
Other expenses	28,663	27,018	24,412	19,615	18,549
Cost of fish and shrimp					
Cost of business	220,176	167,323	199,933	150,300	159,472
Cost of labor provision					
Cost of other production activities					
Total cost of production	252,120	197,604	227,575	173,112	181,187
Interest revenue					
Interest expense	55	55	55	75	55
Other expenses	2,794	2,783	2,810	2,798	2,786
Depreciation of fixed assets	2,794	2,783	2,810	2,798	2,786
Insurance premium					
Extraordinary items					
Capital gains					
Capital losses					
Net income	-22,684	-12,889	-2,945	16,125	20,597
Consumption	9,035	9,362	8,145	10,849	8,566
Savings	-31,719	-22,251	-11,090	5,276	12,031

Source: Samphantharak and Townsend (2009).

Table 3.3
Comprehensive financial accounts: statement of cash flow of household A.

Adjustments to get to cash flow as distinct from accrued income, then cash flow from production, consumption, and financing activities

Month	5	6	7	8
Net income (+)	-22,684	-12,889	-2,945	16,125
Adjustments				
Depreciation (+)	6,075	6,046	6,040	5,996
Change in account receivable (-)	-147,488	-116,288	-147,100	-107,023
Change in account payable (+)	149,960	149,960	147,195	147,195
Change in inventory (-)	-126,465	-106,205	-148,866	-128,883
Change in other current assets (-)	1,781	3,263	3,230	3,198
Consumption of household-produced outputs (-)	-350	-314	-383	-373
Cash flow from production	-139,171	-76,427	-142,830	-63,765
Consumption expenditure (-)	-8,685	-9,048	-7,762	-10,476
Capital expenditure (-)	-3,281	-12,463	-3,230	-3,198
Cash flow from consumption and investment	-11,966	-21,511	-10,992	-13,674
Change in deposit at financial institution (-)	-8,895	-698	-125	11,295
Change in ROSCA position (-)	-4,000	-4,000	-4,000	29,500
Lending (-)	0	0	0	0
Borrowing (+)	-2,000	-2,000	-2,000	-2,000
Net gifts received (+)	-710	618	-311	38
Cash flow from financing	-15,605	-6,080	-6,436	38,833
Change in cash holding (from statement of cash flows)	-166,742	-104,019	-160,258	-38,606
Change in cash holding (from balance sheet)	-166,742	-104,019	-160,258	-38,606

Source: Samphantharak and Townsend (2009).

other household j . As with Bitcoin, there is an initial state, that is, who holds coins, as in the balance sheet, a state that is modified by a transaction in the cash flow statement to deliver a new state. A difference with paper currency, though, is that currency is held by the household as part of its balance sheet and is not public. Currency is an actual portable physical token but not an electronic entry. It is very much a decentralized way of keeping track of histories, a point we return to subsequently. But the accounting concepts underlying the use of currency and the use of coins are exactly the same: cash flow and balance sheet.

A formal statement of cash flow goes a bit further. It distinguishes the purpose of the cash outflow (or inflow) and thus records cash used for consumption and investment, for production, and for financing as in borrowing and lending. The households in the Thai villages do not keep these cash flow accounts. However, we constructed one for each surveyed household from the Townsend Thai data. Such cash flow statements are essential for the study of liquidity, hinting at a data use for DLT that we revisit throughout this book. Recording liquidity is an essential feature of many distributed ledger systems and a key aspect of mechanism design.

We can now link the statement of cash flow to the new language of distributed ledgers with only a few conceptual steps, as it is not difficult to imagine how the new technology could map onto the current paper currency systems. First, one could imagine in principle that accounts could be kept on a common account or centralized common ledger. To establish a proof of concept, this is being done with the Townsend Thai data. It is “just” a new integrated database that we are creating. A key step here is whether the transactions that would be recorded on the ledger are consistent with each other, which involves rechecking the database to quantify discrepancies. If i transacts with j , is j in the database and, if so, is j also reporting a

transaction with i ? Yet this uncovers discrepancies, and this is one of the main things DLT can remedy. Some of these discrepancies could be innocent measurement errors or honest mistakes in reporting. Unfortunately, when we were first gathering these data two decades years ago, we did not have the conceptualization of the common ledger as a check on the gathered monthly data. These and other types of discrepancies matter, for example, as in the New York financial markets. One purpose of the common ledger component of DLT is to have consensus and avoid subsequent validation.

As an example one can think about Digital Asset's construction for the Australian Stock Exchange (Martin, Lee, and Townsend 2017). First, two parties A and B meet on an outside trading platform and agree to trade. The trade information is sent by the trading platform to the clearing exchange, CSD. The CSD writes an encoded message on the distributed ledgers. The message can be read by A and B and references a contract ID. A and B can read from the distributed ledger the messages that pertain to them and can run Digital Asset's DAML code on the contract to verify that it does what it is expected to do. The contract information makes it into their own personal contract stores. Thus the state of the world, the trade in this instance, is recorded by consensus on a shared distributed ledger. We shall return to the encryption aspects of such transactions in a subsequent section.

In the Thai context, the distributed ledgers could also be implemented in practice—in real time—if transactions were out of e-wallet coins and hence recorded (electronic measurement more generally comes up in the next section). Next, the common ledger could be created and distributed among the households so that each has access to the common account or to its own identical copy (subject to privacy, which we come back to under mechanism design in chapter 6). We would call this new common integrated database a *distributed ledger*.

In summary, the idea of ledgers as a statement of cash flow is not at all new. Yet when put on a common database, discrepancies can be readily spotted and corrected. Approved histories can be thought of as immutable. A limitation is that only subsets of transactions might be recorded, in which case coordinated statements of cash flow could be incomplete. On the other hand, the vision for further utility comes with the creation of complete and integrated accounts.

3.2 Financial Accounts as Ledgers More Generally

A unified, more comprehensive measurement of the financial environment is represented by the entire set of complete financial accounts. Specifically, measured transactions can be used to create formal financial statements, not only the statement of cash flow and balance sheet but also the income statement. More specifically, one can use an initial baseline survey to enumerate financial and real assets held at the beginning point in the time line of the survey.¹ Items on the balance sheet would be the amount of currency held, land, and other assets. Indeed, when on a common ledger, this links to the idea in cryptography of using ledgers as a registry of secure property titles. Likewise, cryptocurrency on the balance sheet would be an asset, hence termed a *digital asset*.

Of course, liabilities can be measured in the same way. Subtraction of liabilities from assets thus determines initial net worth. Then there are transactions over time. A household, for example, surrenders currency to buy another asset. Currency can be used to buy consumption, an expense on the income statement, and income is received as revenue on the income statement. The difference in revenues and expenses is saving, which, along with incoming gifts and remittances, must be equivalent with an increase in net assets. The statement of cash flow is similar to the income statement except that for

the income statement, one typically uses an accrued income concept. Expenses are booked only when there is revenue, as in finance, to measure profits as the return on assets used to operate projects. The distinction between accrual and cash flow methods allows for the distinction between productivity and liquidity and is often essential.

A key point to note is that a given transaction in the data can and typically will enter multiple times across individual statements. Thus, the changes in balance sheet and income statement must be consistent with each other. The books have to balance. This is the idea behind double-entry bookkeeping, done for accuracy at the individual entity level, which was a huge innovation at the time the concept was invented. Luca Pacioli, whose work was published in 1494, is considered to be the “father of accounting,” but the inventor could have been Benedetto Cotrugli, even earlier, in 1458. The use of distributed and common ledgers to reduce discrepancies is another layer on top of the conventional double-entry account and is arguably as important an innovation as double-entry bookkeeping.² This is a more accurate system than accounts individually. All this comes from the log of transactions.

Of course, to create the complete financial accounts from a distributed ledger, certain metadata have to be recorded as part of measured transactions. Again, as an example, the code that creates the accounts for the Townsend Thai data operates on the underlying transactions data, preprogrammed to recognize, from the questions to which the transaction answers are given, where in the accounts particular transactions should be entered. Any entity (e.g., a large firm) is doing this with its own proprietary financial accounts, so the firm at least knows the nature of its own transactions through the lens of financial accounts. In contrast, a distributed ledger that records only transactions without categorization cannot be used to create complete financial accounts. The middle common ground is

perhaps the most interesting: Transacting parties record the categorization, and reconciliation seeks to make the categorization common. This could be an additional advance made possible with common ledgers: *consensus categorization*. The common accounts component of DLT could allow this to be done while maintaining privacy, just as DLT can remove discrepancies in trade.

To sum up, complete financial accounts can have value for the accuracy of measurement, for analysis of data, and hence for the households and businesses themselves. The log of transactions can have value for policy.

We now turn to two example applications that convey the value of enhanced financial accounts for policy in emerging markets and in the United States, to track the impact of tariffs or liberalizations and to measure liquidity flows to build microfounded macro models. A subsequent section on cryptocurrency shows how a log of transactions can be used specifically as a basis for an activist cryptocurrency policy of a digital reserve bank.

3.3 Two Examples of the Use of Village and Community-level Financial Accounts: Tariffs vs. Real and Financial Liberalization and Liquidity Accounts for Multiple Media of Exchange

There is huge interest in the impact of tariffs in the United States, under the Trump administration, and it would be useful for trade and financial flows to be recorded in real time. Likewise, in reverse, one could examine the impact of trade and financial liberalizations in emerging markets.

Paweenawat and Townsend (2012) follow the Bureau of Economic Analysis (2017) guidelines and show first how to reconfigure household and business financial statements and,

second, how to aggregate up to create the set of national income and product accounts (NIPA), with the economy as the village. The income statement is transformed, being careful about value added to the production account, and the balance sheet is transformed by taking time differences to create the savings/investment account. Flow of funds accounts measure net acquisition of financial assets, assets minus net incurrence of liabilities, which is equivalent to gross savings less expenditures on real capital. The balance-of-payments account of the village economy follows, thus explaining how villages, and regions, interact with each other.

3.4 A Counterfactual Policy Analysis

Paweenawat and Townsend (2018) calibrate a model that integrates real and financial sectors, allowing for occupation choice, trade in goods across manufacturing and agricultural sectors, and external borrowing and lending. The model has judiciously chosen obstacles to trade—namely, transactions costs for commodity trade and collateral requirements for credit. After fitting the model-generated village paths to the data, one can examine simultaneously and consistently the activities of the featured case study of sampled households and businesses along with selected aggregates. That is, one can determine what is happening over time at the household level with their own financial accounts and what is happening at the economy-wide, village level with the NIPA accounts. One can also distinguish movements of real capital from movements in paper currency.

It then becomes possible to conduct counterfactual policy analysis: What if trade and capital flows had not been allowed to liberalize, or alternatively, what would happen if there were further innovations stemming from an enhanced financial infrastructure? If, for example, there had been a push

for protection in the past, and somehow trade across regions had been more restricted, then a wedge would move relative prices. Likewise, one can examine counterfactual restrictions on interregional flows of capital if more savings had been targeted to be invested at home. The model predicts what would have happened to interest rates, wages, and prices; to occupation choice, production, profits, and earnings; and, finally, to the trade balance, the current account, and the balancing flows of borrowing/lending.

The impact of altered policy is not homogeneous. Removing an obstacle at the village level, through DLT or other means, is not the same thing as increasing social value. Likewise, imposing obstacles can be welfare improving for some households as a function of balance sheets and income flows.³ The application is, of course, not specific to Thailand. One can imagine examining the impacts of tariffs and flows of funds in the US, if we were to have the requisite data.

3.5 Generalized Statements of Liquidity Accounts in the US

Nothing in these examples is particular to Thailand and the predominant use of paper currency as the medium of exchange there. Samphantharak, Schuh, and Townsend (2016) show how a conventional statement of cash flows can be made for advanced countries such as the United States. Using actual data, the statement of cash flow for households is disaggregated into item-by-item liquidity accounts: the inflows and outflows to and from demand deposits; credit, debit, and prepaid cards; and paper currency. Likewise, though not yet well measured in the surveys of the Federal Reserve Bank of Boston, the statement of liquidity accounts links conceptually to the other financial accounts and thus to the variation in income and long-term financial assets.

Of course, distributed ledger technology is not yet the source of data, though one can envision how we might get there. The Boston Fed survey uses data from survey questionnaires and data from diaries. Likewise, data from the Survey of Consumer Finances (SCF) and Panel Study of Income Dynamics (PSID) are interview-based. Still, there is increasing use of administrative data (which is electronic) as a cross-check on consumer responses. For example: Are households reporting bank transactions consistent with data from corresponding banks? Browning, Crossley, and Winter (2014) seek to integrate the collection of wealth, income, and spending data in the British Household Panel Survey so that for each household the intertemporal budget constraint holds. An Office for National Statistics (ONS) Economic Expert Working Group (EEWG) envisions using web surveys, mobile surveys, and phone apps to scan barcodes and till receipts. There is now also electronic data surrendered voluntarily by customers, as with Mint, and the use of commercial bank information by information aggregators. Use of DLT to create complete financial accounts is not as far as away as it might seem, *a priori*.

How could this data be used? As emphasized here in this section, we could see and understand better the role of liquidity in an economy. Work emphasizing liquidity and payments and links to monetary policy is a bit sparse but increasing. Significant recent contributions include Kaplan and Violante (2014), Piazzesi and Schneider (2018), Adrian and Shin (2009), Doepke and Schneider (2006), Auclert (2019), and Fulford and Schuh (2017). DLT creates the capability of providing enhanced measurement over and above current US surveys. Currently there are discrepancies between the cash flows associated with aggregated income statements and the cash flows associated with changes in aggregated balance sheets (Samphantharak, Schuh, and Townsend 2016).

3.6 DLT vs. Traditional Database: Limitations of Distributed Databases to Be Recognized and Incorporated in Designs

A ledger could be viewed as a traditional database in which a user can create, read, update, or delete (CRUD) (Ray 2017). The risk in having centralized control of a database is that anyone with sufficient access to it can destroy or corrupt data, so users are reliant on the security infrastructure of the database operator and must trust those with write capabilities. The March 2019 episode with Capital One and the Amazon Cloud, hacked by a former employee, is illustrative.

In contrast, distributed ledgers use decentralized data storage, in the sense that the ledgers are distributed among users. With cryptographic rules for change, security is inherent in this structure; there is no single copy. With distributed ledgers a user can read and retrieve data—that is, audit records—and a user can write by adding more data to append only. Newly proposed transactions must be validated in some way, as is discussed in chapter 5. Likewise, past validated histories are immutable. There is no updating of past transactions and no deletion. A key property of blockchains such as Bitcoin is that they do not rely on a single trusted third party as trustee or notary to intermediate transactions. The blockchain network enforces execution, giving this a social aspect. This is what Nakamoto (2008) meant by a system without a trusted third party.

But with the decentralized system of distributed ledgers comes known database problems (Wikipedia 2019a). A theorem in computer science, the CAP theorem, states that it is impossible for a distributed data store to simultaneously provide more than two out of the following three guarantees: (i) Consistency—where every read receives the most recent write or an error; (ii) Availability—where every request receives a

(nonerror) response, without the guarantee that it contains the most recent write as with consistency; (iii) Partition tolerance—in which the system continues to operate despite an arbitrary number of messages being dropped (or delayed) by the network between nodes, so that there is a partition, multiple versions.

To highlight further, in the presence of allowing partitions, one has to choose between consistency and availability. That is, when a network partition occurs, one has to decide to cancel the operation, which decreases availability but ensures consistency, or proceed with the operation but risk inconsistency.

Furthermore, even when the system is running normally, there is a tension between consistency and availability because of latency. Latency is the amount of time a message takes to traverse a system, or how much time it takes for a packet of data to get from one designated point to another (Wikipedia 2018b). High-speed but virtually instantaneously available systems run into latency issues, as this can create inconsistency across multiple versions. With Corda, for example, latency determines the geographic distribution of validators, in some instances, to mitigate delay.

More specifically, we come to the Fischer Consensus Problem of distributed computing (Fischer, Lynch, and Paterson 1985), though we need some definitions first:

In computer science, synchronization refers to one of two distinct but related concepts: synchronization of processes, and synchronization of data. Process synchronization refers to the idea that multiple processes are to join up or handshake at a certain point, in order to reach an agreement or commit to a certain sequence of action. Data synchronization refers to the idea of keeping multiple copies of a dataset in coherence with one another, or to maintain data integrity. Process synchronization primitives are commonly used to implement data synchronization. (Wikipedia 2019b)

Otherwise the process is asynchronous.

Fischer proved that it is impossible to guarantee that any asynchronously connected set of communicating nodes can agree on even a single bit value—a devastating result. On the other hand, the Fischer consensus problem can be resolved simply by synchronizing from a single point. However, doing so introduces a single point of centralization, which is ironic given the decentralized connotation of DLT as emphasized in chapter 1. This centralization in turn can cause scaling problems—that is, as in many DLT consensus algorithms, every node must be connected to every other to achieve consensus, and the costs of messages rise exponentially with the number of nodes. Computer science/distributed systems bounce between these problems: CAP in asynchronous systems, and scaling and fault tolerance in synchronous systems. These features should drive choices as part of a constrained-optimal design (Mallett 2019).

Again, the various distinct consensus protocols for validation cope with these problems in alternative ways and illustrate the trade-off between hyped decentralization, which has high congestion because of the underlying centralized features, and named validators, including single validator systems, which again raises the issue of trusted third parties. The existence of a trusted third party can greatly enhance speed and lower costs. Another example is Digital Asset’s innovation for the Australian stock exchange. Adopting the language of Casey et al. (2018), these systems are known as “permissioned” (or “private”) blockchains, with a limited set of entities, or even a single organization, allowed to write to the blockchain. This can reduce scaling problems.

Another example highlighted in Casey et al. (2018), the Lightning Network, aims to greatly reduce cost and time constraints by shifting small transactions to a cryptographically secure “off-chain” environment so that only large netting transactions need to be directly settled into a resource-constrained

blockchain. With Hyperledger Fabric, a permissioned blockchain, a third-party auditor or regulator can obtain provably correct answers to queries about the system as a whole using zero-knowledge proof concepts. Centralized DLT exchanges for cryptocurrency have “relayers,” application interfaces that allow users to trade in a decentralized manner (Bronstein 2018).

Mallett (2009) compares strictly hierarchical/client servers with fully connected mesh networks and then speaks to the advantages of partial mesh networks. The point is to compare and potentially select among network designs. The US military uses a hierarchical system that suffers from a lack of incorporation of local information but minimizes latency by minimizing communication, as one-way commands from headquarters are obeyed. These hybrid designs need to be integrated further with economic systems, and, indeed, the industrial and management organization may be endogenous with the design selected.

As a suggestive example, Townsend (1978) uses simple, transaction-cost arguments, with fixed costs per node for any given bilateral connection. Optimal risk-sharing arrangements partition agents into segregated subgroups. Despite ever-decreasing per capita costs and ever-increasing gains from having all agents in one mutual fund, due to portfolio diversification and the law of large numbers, marginal costs can exceed marginal benefits from increasing group size. More generally, a related economic issue is whether to have over-the-counter (OTC) markets, centralized platforms, or a hybrid in between. Although it may seem advantageous to have all trade taking place in one spot, the transaction costs integrated with economics may suggest otherwise, potentially.

The main conclusion is that there are trade-offs in design. Which system might dominate is a function of the environment and goals.

