Distributed Ledgers

Design and Regulation of Financial Infrastructure and Payment Systems

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

7

Design Issues: Partitioned Ledgers, the Decision to Decentralize Implementation in Multiparty Contracts, and Incentive-Compatible Token Payment Systems

This chapter considers private information and partitioned ledgers, delegation of authority to a third-party platform, and tokens as communication devices.

7.1 Permissioned Private Ledgers and Gains from Concealment

In many environments, "unique consensus" is not desirable even if it were technologically possible. This is a generic implication of private information for optimal contract design. Though private information is effectively reported via messages or indirectly by choice and display options, this does not mean that such internal message data should be made public on the ledger. Often, the opposite is true: Messages should be kept private. Of course, cryptography with partitioned ledgers makes this possible.

Distributed ledger technology allows partitioning, but it does not have to be done in a mechanical way or, indeed, in the most extreme way often proposed—namely, by keeping all proprietary information entirely private. Likewise, an instance when all messages are made public means that the parties to

the contract would see them, undercutting welfare. Neither of these extremes, all private versus all public, is typically constrained optimal.

Suppose shocks are to preferences and not to endowments. That is, agents are either urgent or patient to consume. This is the standard specification in modeling financial institutions or trade in financial markets. The Diamond and Dybvig (1983) model of banks and runs has versions of this exactly. The Duffie, Gârleanu, and Pedersen (2005) model of over-the-counter (OTC) trade in securities and consumption has stochastically varying security holding costs that motivate the trade. More generally, banks and traders face shocks to their portfolios that arise from the needs of their customers.

As an example, table 7.1 considers an economic environment with two parties to a contract with heterogeneous and random preferences to consume (Townsend 1988). There is positive, but not perfect, correlation in these shocks over time and over agents. Here all agents are risk averse (not just one of them as in the earlier example). However, shocks are private to only one agent at each date, and the identity of that agent alternates over time. That is, only one of the two parties in an initial period announces urgency today in the first period, hopefully compensated by more goods or value today if urgent, and the other, second party, announces tomorrow in the second period. The agents announce their states to the contract node, not to the other party, and then the smart-contract algorithm determines what is made public. This is where the optimal design kicks in.

If the announcement of the first agent in the first date were public, it would undercut insurance possibilities for the second date for the second agent. Typically, with allocations and histories of announcements known, there can be no insurance for the second agent in the second period. More is preferred to less, so it is hard to engineer an incentive-compatible trade-off

Table 7.1
Partitioned private ledgers and the gains from concealment.

Private information solution

θ_0^a	(c_0^a, c_0^b)	$\pi(c_0, \theta_0^a)$	θ_0^a, θ_1^b	(c_1^a, c_1^b)	$\pi(c_1, \theta_1^b)$
0.2	(1.75, 8.25)	1.0	(0.2, 0.2)	(4.75, 5.25)	1.0
			(0.2, 0.9)	(2.0, 8.0)	1.0
0.9	(0.0, 10.0) (1.75, 8.25) (3.25, 6.75)	0.1159346 0.0339681 0.8544384	(0.9, 0.2)	(3.75, 6.25)	1.0
			(0.9, 0.9)	(1.0, 9.0) (10.0, 0.0)	0.86106 0.13839

Agent a has preference shocks, urgency to consume, θ_0^a, at date $t = 0$ of either 0.2 or 0.9. The announcement triggers consumption allocations (c_0^a, c_0^b) for agents a and b, the latter as the second party. These are listed in the second column of the table. Notice that if the announcement is 0.9, then a lottery puts about 3.3% probability on (1.75, 8.25) the deterministic allocation when 0.2 is announced. Probabilities $\pi(c_0, \theta_0^a)$ are in the third column of the table. At $t = 1$, agent b announces 0.2 or 0.9. The mechanism knows the previous history of incentive-compatible announcement of θ_0^a at $t = 0$, but agent b does not. The agent has an incentive to announce θ_1^b truthfully, and there is an insurance transfer along all paths. If, however, agent b were to have seen θ_0^a, there is no insurance over θ_1^b (top row right). Here, without that information, agent b might be tempted to claim urgency, 0.9 always, but doing so risks losing everything with 14% probability if in fact agent a had announced 0.9 (as in the bottom row, last column). In summary, when $(c_0^a, c_0^b) = (1.75, 8.25)$ is observed, agent b remains uncertain of the type agent a reported in $t = 0$, and this is crucial for incentives.
Source: Townsend (1988).

other than a trivial one, no trade-off at all, to transfer a constant amount regardless of the state. However, if agent two in the second period were unsure of what was announced by agent one in the first period—something that a partitioned ledger can keep secret—then an optimized design will cause agent two to weigh the consequences of lying, announcing urgency but with positive probability ending up with very little consumption. Insurance and truth-telling are achieved by having the mechanism itself randomize over the consumption allocations in the

first period as a function of agent one's message, so that actual allocations do not fully reveal either. That is, there are common elements in the support that can happen for any messages. Allocations are seen by all, but the message is not, and allocations do not reveal the whole story. With risk aversion, concave utility, this randomization *per se* comes with a welfare loss, but it is outweighed by the overall insurance benefit made possible by concealment. To summarize, agent two now faces a trade-off: If the agent announces a counterfactual state, succumbing to the temptation to lie, then with positive if small probability agent two would achieve a disastrous outcome.

We shall return to the discussion of randomization and concealment when we later address payments platforms.

7.2 Delegation of Portfolios to a Third Party: Platforms as Custodians

DLT allows commitment to a multiparty smart contract in which awards are allowed to vary over time as a function of shocks. When there are both private, unobserved idiosyncratic shocks and publicly observed aggregate shocks, it can make sense for households to delegate portfolio decisions and commit to a third-party custodian (Townsend 1988). This is a kind of endogenous centralization, with reliance on a third party, which is ironic given that those who promote distributed ledgers say that centralization is bad. This concentration of decision-making can happen in practice as with village funds, cooperatives, or wealth managers, or with exchange-traded funds that can be undone only by a restricted set of designated participants. But the decision of whether or not to do this is a function of the underlying environment, not an exogenous *desideratum*. The gain from a preprogrammed, third-party custodian is that it allows front-loading, as when incentive constraints in future periods bind, limiting the value of having

resources then, so more value is paid out contemporaneously, or back-loading, which strengthens inter-temporal incentives by having more at stake in future periods contemporaneously as a function of what is said today.[1]

DLT can facilitate implementation of a multiparty contract with a commitment to the longer term, including sequestering funds to prevent withdrawal from the arrangement. The third-party custodian could appear as an implementing node, a financial institution, or a reserve bank, for example, engaged in a smart, multiparty contract.

7.3 Private vs. Public and the Role of Tokens

In this section we consider payment systems that are designed to be constrained optimal in support of trade, credit, and insurance.

The essential idea has been presented above several times: Distributed ledgers could keep track of messages as a part of the execution of a multi-period, multi-commodity, multi-agent smart contract and thus optimally allocate underlying risk, while facilitating trade and exchange.[2] Featured here in this section is the use of tokens, both single and multi-colored coins, to achieve this objective. The role of tokens is twofold. First, tokens are one way to interpret ledger entries in a centralized system, as colored entries. Second, tokens as real objects can allow a decentralized implementation of the same allocation, a hybrid system that can mitigate the scaling problems of centralized communication systems. If tokens are held in private, then incentives for voluntary disclosure need to be included, though this is not always a binding constraint. How well these various accounting and token systems function depends on the size of the message space relative to the needs for credit, insurance, and trade coming from the underlying economic environments. Nevertheless, in some instances, information should

be kept private, so more limited message systems are actually preferred as constrained optimal.

7.3.1 Tokens on Ledgers as a Way to Achieve Unique Consensus; an Insurance Example with Voluntary Disclosure

As in Townsend (1987), suppose there are four agents in spatially separate locations and some subset of the agents travel. More specifically, a risk-averse agent a is paired with a risk-neutral agent b initially, in the first period at one of the two locations, such that the risk-neutral agent can insure the agent who is risk averse—and likewise, symmetrically, for agents a' and b' at a second location in the first period. If the pairings do not switch, we are back to the first example in chapter 6 on hybrid borrowing/lending and insurance pairwise.

But now suppose the pairings switch locations in the second period. Agents a and a' make announcements of their urgency to consume to their new partners, b' and b, respectively. To induce truth-telling, if urgent in the first period, the agents receive the good but at the expense of getting less of it in the second period. Likewise, if patient today, the agent receives less of the good today at the benefit of getting more of it in the second period. A centralized public ledger for this four-agent, multiparty arrangement, which records all messages and keeps track of history, is one method of implementation. It reduces, equivalently, to the outcome of two separate pairings.

On the other extreme, if there were no record of announced preference shocks and no record of allocations in the first period, then there could be no link of the first period to the second period, and so, essentially, no insurance can be obtained in either period for agents a and a'.

The introduction of tokens as a hybrid system can solve this problem without the centralization inherent in common ledgers. Announced patient agents in the first period receive more tokens than urgent agents. Tokens could be carried literally

as coins or physical objects. In the second period, agents with more tokens can display them in order to be on the receiving end of goods. Tokens become the communication device and alleviate the burden of keeping track of history on a centralized ledger. Alternatively, consider private, immutable but partitioned ledger entries that are not validated by the entire community. Tokens or DLT entries are equivalent to each other; each conveys the necessary history.

7.3.2 Multiple Colored Tokens and Distinguished Histories: Trade with Insurance

We need not rely on pure insurance examples. In environments with two or more goods, there are exchanges of one good for another at each date, driven by the usual motives for spot trade, except those desires to trade are driven by preference shocks impacting inter-temporal trade-offs, and those are private.

Again, in a hybrid decentralized system, a portable, concealable token system could be used to keep track of trades in the first period. An agent may trade in the first period, give up one good and acquire the other, and be expected to reverse the situation in the future. This is the same patient urgent dichotomy but here for each good separately. That would be fine, as with the earlier examples; portable tokens can handle this situation. But now we introduce additional shocks to inter-temporal discount rates, with different shocks for different goods. After these shocks, agents may have ex post regrets and wish to reverse the trade in the first period in order to get the good they now most prefer in the second. One possible solution is to have multiple colored tokens, so as to have more dimensions in which to keep track of more detailed histories or, equivalently, multiple digital assets on a private ledger (see table 7.2).

These ideas in economics have tight links to cryptography and the idea of colored coins. The discussion here draws on Narayanan et al. (2016). As an example, ordinary Federal

Table 7.2
Colored coins and partitioned private ledgers.

Multi-period private and full information solution, two goods					
Values for $(\theta^a_{1x}, \theta^a_{1y})$	(c^a_x, c^a_y)	Values for (δ_x, δ_y)	Values for $\theta^a_{2x}, \theta^a_{2y}$	(c^a_x, c^a_y)	
(.4, .6)	(2, 8)	(1, 1)	(.6, .4)	8.01	2.0
		(.5, 1.5)	(3, .6)	1.0	8.0
		(1.5, .5)	(.9, .2)	10.0	0.82
(.6, .4)	(8, 2)	(1, 1)	(.4, .6)	2.0	8.0
		(.5, 1.5)	(.2, .9)	0.82	10.0
		(1.5, .5)	(.6, .3)	8.0	1.0

Agent *a* has utility function $U^a(c_x, c_y, \theta^a_1)$ at date $t = 1$ over consumptions c_x and c_y with preference shocks θ^a_{1x} and θ^a_{1y}, respectively. At date $t = 2$, agent *a* has utility $U^a(c_x, c_y, \theta^a_2)$. Note in particular the discount rate δ is random. Specifically, the utility function of agent *a* at date 1 is of the form

$$U^a(c_x, c_y, \theta^a_1) = (c_x)^{\theta^a_{1x}} + (c_y)^{\theta^a_{1y}},$$

with $(\theta^a_{1x}, \theta^a_{1y}) \in \{(.4, .6), (.6, .4)\}$, each with equal probability, and at date 2 of the form

$$U^a(c_x, c_y, \theta^a_2) = (c_x)^{(1-\theta^a_{1x})\delta x} + (c_y)^{(1-\theta^a_{1y})\delta_y}$$

$$= (c_x)^{\theta^a_{2x}} + (c_y)^{\theta^a_{2y}},$$

with

$$(\delta_x, \delta_y) = \begin{cases} (1, 1) & \textit{with Prob } .96 \\ (.5, 1.5) & \textit{with Prob } .02 \\ (1.5, .5) & \textit{with Prob } .02 \end{cases}$$

Agent *b* has linear preferences.

In the table, there are two goods, and agent *a* can be a "borrower" or a "lender" in either good. Still, "preference reversal" shocks at date 2 can cause agent *a* to want to pretend to have been a lender in the commodity the agent did not lend.

Source: Townsend (1987).

Reserve bank notes can be given bar codes so they can be used as purchased tickets to Yankee baseball games. The team signs a message that includes a specific game date, seat number, and serial number of the bill, with the signature of the issuer, all stamped on the bill. The advantage of using preexisting bank notes is that they cannot be easily counterfeited, so there would be no need to print new tickets. And such a system is decentralized. The message is written on the token, so to speak. Alternatively, the stadium could check a central database for information when a "ticketholder" enters the gate with a note having a certain serial number. The serial number links back to the primitive transaction, the purchase of the ticket. Either way, metadata is being attached to the note.

The point is that coins have publicly verified histories to trace ownership. This history can be made meaningful and put to other uses. Coins that "originate" in certain transactions can have associated extra metadata that behave like a color. The colors can also be considered as a metaphor, of course, as they could simply be bit strings. It is important, of course, that all participants understand the rules of this payment system so they know how to interpret the colors.[3]

To summarize simply, and to link back to the economics, distributed ledgers could keep track of messages as a part of the execution of a multi-period, multi-commodity, multi-agent smart contract, and thus optimally allocate underlying risk while facilitating trade and exchange. The "money" here, or more generally the payment system, is not separated from the motives for trade, intertemporal exchange, and insurance.

7.4 Permissioned Private Ledgers, When Consensus Is Not Unique Due to Optimally Kept Secrets

While multi-colored tokens convey more history, one should not jump to the conclusion that more information is preferred to less. We gave an example of private information and

randomization earlier, with two agents, two dates, and one good. Shocks are private to only one agent at each date, and the identity of that agent alternates over time. If the announcement of the first agent in the first date were public, it would undercut insurance possibilities for the future.

If we generalize this example to four agents and imagine that agents cannot know what happened at the first date in a different location, then tokens can be allocated and tied to the random consumption allocation. Tokens can be carried into the second date, so the public part remains public. Yet tokens need not reveal entire histories of messages. That is, if tokens are colored data entries on private systems, there should not be unique consensus, in the words of Corda, in such environments.

The insurance example may seem a bit counterintuitive, but the same idea shows up in applied work in finance and with the same elements: risk aversion and information asymmetry. Lyon (1996) analyzes the optimal transparency of order-flow information as in foreign exchange markets, arguing that slower revelation of information—information that could reveal market-wide order flow—improves risk-sharing among dealers facing unavoidable position disturbances. Garratt et al. (2018) show in a similar but distinct context that some post-trade information disclosure can improve liquidity, but revelation of information (for sale) by a self-interested platform is a worse outcome than no information at all.

There are other examples of optimally limited shared information that make itineraries and validation endogenous. In Prescott and Townsend (2006b), auditors (or one might say validators) make incentive-compatible announcements of underlying states and then depart, making way for an incoming and relatively uninformed agent, assigned to be there as a solution to the mechanism design problem. The role of this auditor would be akin to verifying underlying states or objects on a

ledger. For the incoming uninformed agents, they do not know what path they are on and face trade-offs, both in making announcements or in taking actions. Ironically, in this context, it is an advantage that past history is not known.

All these optimally designed systems require commitment to the design, including the control of information. Leakage is a potential problem in practice, and privacy remains a concern.