Medium of Exchange in Blockchain Economy

Agostino Capponi* Zhengge Zhou†

This version: March 2026
Incomplete draft. Please do not circulate!

Abstract

This paper asks: Do we need a medium of exchange, such as money, in a blockchain economy? Due to the nature of blockchain technology, a validator is randomly selected to produce the new block at each block slot. Because block production consumes computational resources, the validator is incentivised to re-arrange transactions within the block for minimal block production costs. We show that this incentive would result in a hub-spoke transaction network. The hub in this optimal network, therefore, intermediates all transactions and acts as a medium of exchange endogenously in the blockchain economy.

JEL Codes: D40, E50, G10

Keywords: Medium of Exchange, Blockchain Economy, Network.

*Department of Industrial Engineering and Operations Research, Columbia University, New York, NY 10027, the United States, ac3827@columbia.edu.

†Accounting and Finance Department, University of Western Australia, Crawley, WA 6009, Australia, zhengge.zhou@uwa.edu.au.

1 Introduction

The role of a medium of exchange in a barter economy has been a central topic in economics at least since Cournot (1838), where the author observed that in an economy with n commodities, direct barter on a point-to-point basis requires up to $\frac{1}{2}n(n-1)$ distinct pairwise trading markets. As argued by Jevons (1875), the fundamental inefficiency of such a barter system is the "double coincidence of wants" problem: a transaction can occur between a pair of agents only if each agent possesses precisely what the other desires. This friction is the genesis behind the emergence of indirect trade and the use of a medium of exchange. A commonly accepted medium, such as money in modern society, allows this n-commodity economy to function efficiently with only $n-1$ (or n) pairwise trading markets, thereby eliminating the need for long chains of intermediaries to pass goods along to ultimate consumers.

Recent financial innovation – blockchain technology – has renewed the view of how a barter system can function efficiently. In a blockchain economy, rather than interacting with any intermediaries or searching for counterparties, agents directly submit their transaction orders, say generic barter orders that exchange one asset for another, to a canonical ledger, where orders are to be verified, included, and settled by a randomly selected computer node at each block slot. The blockchain economy embedded with such a decentralised infrastructure thus allows agents to barter directly in an economic system at a relatively low cost. From this perspective, it may appear that the conventional rationale for the need for a commonly accepted medium of exchange in a barter system no longer applies, particularly in a blockchain-based economy. One core question of this paper is to study whether this conventional wisdom regarding the medium of exchange is still valid for blockchain exchange.

By considering a barter environment setting purely blockchain-based and otherwise intermediary-free, we ask whether, when, and through what mechanisms a medium of

exchange may still emerge endogenously in this environment. To the best of our knowledge, this paper is the first in the literature to provide a formal analysis of the emergence of a medium of exchange in a blockchain economy and to identify the underlying economic forces driving this phenomenon. Beyond this theoretical contribution, note that blockchain stablecoin is pegged to fiat money, with the latter a commonly accepted medium of exchange in the traditional economy. Therefore, results of this paper also have direct implications on stablecoin practices and policies and speaks about the consequence of stablecoin/CBDC issuance on the blockchain ecosystem.

To capture the fundamental role of a medium of exchange in the blockchain economy, our model features two types of agents. The first type, users, generate generic barter demands in the economy. They want to exchange one asset for another. Their barter or transaction orders are executed only when these orders are included and settled by one particular computer node that is randomly selected from the blockchain network. These computer nodes are the second type of agent in our blockchain economy, which we refer to as mediators throughout the paper. Mediators are responsible for the process of block production, such as verifying, settling, and including users' transactions on the canonical chain. Due to the nature of the blockchain protocol, at each block slot, only one mediator is randomly selected to produce the block.[1] Block production consumes computational resources borne by mediators, and the production costs increase with the number of transactions included in a block. To compensate for these costs, the randomly selected mediator receives transaction fees paid by users for the transaction inclusion and settlement.

Once selected, the mediator at the given block slot possesses monopoly power over users' transaction inclusion and settlement. We would expect that the mediator

[1]The specific selection probability depends on the type of consensus mechanism in the blockchain network, such as Proof-of-Work (e.g., Bitcoin) or Proof-of-Stake (e.g., Ethereum).

has a strong incentive to minimise her block production cost by optimally organising users' transaction orders in the block. Such a cost-minimising consideration will be formalised as an optimal route optimisation problem for the monopolistic mediator in this paper. In the language of economic networks, we model each type of tokenised asset as one node and each transaction settlement infrastructure for a direct barter between one pair of tokens as a link.[2] As a result, the block production process by the mediator corresponds to choosing both a transaction network that connects assets and an allocation of transaction flows along paths in this network.

The block space is scarce and expensive. We therefore argue that the cost of the computational resource borne by the mediator in the process of building new blocks includes two additive components. The first component $c_0 > 0$, where we hereafter refer to a link-formation cost, is a positive number representing the fixed cost incurred by the mediator when enabling a new channel for a pair of distinct assets g and h to be directly bartered. The second component is a variable cost $f(\cdot)$ whose exact value critically depends on the size of transaction flows routing on each link. Importantly, no restriction is imposed on the set of feasible networks in the paper, whereby the mediator is free to choose any pattern of connections among assets. This leads to both the cost-efficient network structure and the optimal transaction routes being endogenously determined.

When variable costs exhibit economies of scale, a simple yet economically important condition, we obtain the main result of this paper: the optimal transaction network is either a hub-spoke network or a complete network. More specifically, we show that if the fixed cost is sufficiently low, $c_0 < c^*$, the complete network is the unique optimum, where the mediator has no incentive to reroute users' transactions. By contrast, when

[2] A barter transaction between a pair of assets can be represented as a directed edge from the origin asset to the destination asset, with the weight of the edge equal to the total transaction volume in that direction.

$c_0 > c^*$, the optimal network takes the form of a connected star (hub-and-spoke) structure. In this case, transactions are either between a spoke asset and the hub asset or between two spoke assets that must be routed through the hub asset. The hub asset, therefore, intermediates all transactions in the blockchain economy. Such a hub-token in turn functions as an endogenous medium of exchange in an otherwise intermediary-free blockchain barter environment.

The comparison between a star (i.e., hub-and-spoke) network and a complete (i.e., point-to-point) network is instructive. A complete network requires $\frac{1}{2}n(n-1)$ links, say on the order of n^2, to connect n distinct tokens on a point-to-point basis. In contrast, only $n-1$ links, say on the order of n, are required to connect the same set of tokens in a star network centred around a hub token. Accordingly, the total fixed link-formation cost grows quadratically in a complete network but only linearly in a star network. The star network economises on the number of connections needed to connect a large number of origin-destination barter transactions. Therefore, the benefits of "hubbing" become more pronounced when the size of the economy n is large or when the link-formation cost c_0 is a significant component of total costs, which explains the efficiency of enabling a medium of exchange in the blockchain-based barter system.

The result of this paper speaks directly to the relation between the (cost) efficiency of a blockchain economy and the necessity of well-aligned incentives across mediators on blockchain. In existing blockchain systems, although the selection probabilities vary across consensus mechanisms, only one mediator, such as the miner in a Proof-of-Work blockchain network or the validator in a Proof-of-Stake blockchain network, is selected to include and settle transactions in each block slot. This institutional feature naturally generates a strong cost-minimisation incentive for the selected mediator, which in turn can lead to a cost-efficient transaction-routing structure.

Such cost efficiency is conditional on the arbitrage profit-seeking behaviour of me-

diators being abstracted away. In practice, however, such arbitrage-profit motives can generate sizeable economic rents known as MEV that may potentially dominate the pure cost savings from an efficient network design. In this sense, our results here echo and confirm the recent industrial attempts to separate the monopoly power of the mediator across different functions, with each operating at distinct stages between users and the canonical chain. The builder-proposer separation framework in Ethereum is a leading example, which ensures that the incentive of the proposer is aligned well with the cost efficiency objective of the blockchain. In our model, we take all market participants involved in the transaction settlement as one representative mediator. So, one may want to read the mediator in our model as the proposer in the context of the builder-proposer separation framework.

Remark that including MEV problem in the analysis may affect the optimal network characterisation. Nevertheless, it does not imply that the hub-spoke or star network becomes no longer optimal. On the opposite, the point-to-point complete network can be less attractive than the hub-spoke network, even when c_0 approaches zero, when we account for the computational complexity involved in MEV extraction. The hub-spoke network, as one of the simplest network structures, has a clear advantage in terms of its minimal computational complexity.[3] This provides an additional incentive for the mediator to select the hub-spoke network, thereby increasing the likelihood of the medium of exchange emergence in a blockchain-based barter system.

Lastly, our analysis suggests that the most recent rapid proliferation of stablecoin could potentially deteriorate blockchain market efficiency, as mediators may find it increasingly difficult to coordinate with each others on a single or over a few dominant media of exchange payment nodes. As a result, users' generic barter orders are forced to route through a more complex transaction network in the blockchain economy with

[3]For discussions of computational complexity across different network structures in the context of MEV, see Zhang et al. (2025).

few medium of exchange hub-tokens. Compared to routing orders through a simple star network with only one hub-token, the coordination failure due to a more complex transaction network will eventually come at the cost of the blockchain infrastructure burden.

In terms of literature contributions, our work complements the classic, longstanding literature on the theory of money, but from a particular, new medium of exchange lens. Shubik (1999) and Starr (2012) provide comprehensive surveys regarding the literature on the medium of exchange theory. While a complete review of the vast literature on the role of money in exchange mechanisms is beyond the scope of this paper, two broad strands are particularly relevant. The first builds on search-theoretic models with random bilateral meetings among long-lived agents, following the classic insight of Jevons (1875); see, for example, Jones (1976), Banerjee and Maskin (1996), Kiyotaki and Wright (1989), Kiyotaki and Wright (1993), and Ostroy (1989), among many others. A parallel and equally important strand develops general equilibrium models with alternative trading frictions, such as transaction costs and limited trading opportunities; see, for instance, Wallace (1978), Hahn (1989), Howitt and Clower (2000), and Foley (1989). Our work complements these strands by providing a novel rationale for the medium-of-exchange role of money that is distinct from, yet consistent with, existing explanations, arising in a blockchain-based pure exchange economy.

Our paper also contributes to the rapidly growing literature on the economics of blockchain; see John et al. (2022) for a comprehensive survey. Representative studies examine blockchain adoption (Cong et al., 2021; Hinzen et al., 2019), consensus mechanisms (Biais et al., 2019; Ebrahimi et al., 2025; Saleh, 2021), and strategic behavior and attacks in blockchain systems (Chiu and Koeppl, 2022; Capponi et al., 2024; Garratt and van Oordt, 2023; Budish, 2025; Zhang et al., 2025). Our contribution differs fundamentally from these prior studies by providing, to our knowledge, the first eco-

nomic analysis of the necessity and underlying forces governing the emergence of a medium of exchange within a blockchain protocol.

The remainder of the paper is organised as follows. Section 2 provides the institutional background of the blockchain economy. Section 3 presents the model. Section 4 delivers the main results. Section 5 provides additional discussions. Section 6 concludes the paper. All proofs are relegated to the Appendix Section.

2 Blockchain Infrastructures

Two main types of transactions exist in the blockchain network: payment transactions that involve token transfers between users (or equivalently, wallets), and smart contract transactions that users interact with blockchain-stored programs by sending transactions to them. Leading smart contract applications are decentralised exchanges and decentralised lending platforms, where users can swap tokens and deposit collateral, respectively. Because of the nature of blockchain technology, any transaction order submitted from an individual (wallet) is truly executed only if it is included in a block on the canonical chain. Producing a block thus means the inclusion and settlement of a set of transaction orders submitted from individuals, and the canonical chain is a sequence of such blocks.

At the core of the above block production process are validators, who are responsible for appending blocks of transactions to the canonical chain. For each block slot, however, only one random validator in the network would be selected and assigned to produce the new block. This selected validator receives transaction fees paid by users for its block production task, as this task is at the expense of computational power from the selected validator. The exact value of this blockchain infrastructure expense, which is also known as gas cost, depends critically upon the complexity of

the block content, say the number of transactions and the types of transactions in the block. As for the selection probability, it is set to the validator's share of computation powers under a Proof-of-Work network (in short, PoW, e.g., Bitcoin), but to the share of staked assets under a Proof-of-Stake network (PoS, e.g., Ethereum after September 2022).

Other than selection probability, the above two blockchains also differ in the roles assigned to the selected validator. In the Bitcoin case, the selected validator not only constructs block content via bundling transactions together but also sends (i.e., broadcasts) the block to its peers on the network.[4] The Ethereum, however, transitioned from PoW to PoS and introduced the Proposer-Builder separation (PBS) framework since September of 2020. This framework separates the above dual roles of validators into block builders and block proposers, the former being responsible for creating the block of transactions and offering the block to the proposer, and the latter being responsible for submitting (i.e., broadcasting) the block of transactions to its peers on the network. The block of transactions offered by block builders is an option to block proposers, who can either outsource the task of block construction to a market of block builders or may instead choose to submit a block of transactions to the chain network they have constructed themselves.[5]

Although the role between block building and block proposing is supposed to be separated under the PBS framework, there is no formal restriction preventing proposers from acting as builders.[6] Moreover, given their temporary monopolistic power in determining which transactions are included in a block and in what order, validators under PoW, as well as builders and proposers under PBS, often deploy sophisticated algo-

[4]In the context of the Bitcoin network, validators are often called miners.

[5]Empirical evidence such as Wahrstätter (2024) shows that the block value is significantly higher for blocks sourced from builders, as opposed to self-building blocks from proposers.

[6]For a detailed discussion of the incentives for proposers to become builders, see Capponi et al. (2024).

rithms to extract additional value from block production beyond standard transaction-fee rewards. This additional value was initially referred to as *miner-extractable value* (MEV) and later renamed *maximum-extractable value* following the protocol transition in September 2020.

The allocation of block value among different agents within the mediator sector is an important question for blockchain design. However, it primarily concerns re-distribution within the mediator sector and therefore lies outside the scope of this paper.

Despite differences in institutional roles, proposers and builders in PoS systems, together with validators in PoW systems, share a common incentive to rearrange barter transaction orders within a block to maximise block-production benefits. We therefore refer to these agents collectively as a representative (selected) mediator and model the block-production process as a monopoly problem faced by this representative mediator. The mediator's key decision variable is the optimal arrangement of barter transaction orders within the block content.

3 Model

Consistent with the description of the blockchain production process in the previous section, we now formalise the optimisation problem faced by the monopolistic mediator. The transaction-settlement infrastructure used to arrange block contents is modelled in the language of an economic network in this paper. We next introduce the formal environment in detail.

3.1 Network

Let $N = \{1, 2, \ldots, n\}$ denote a set of $n \geq 3$ nodes. Each node $h \in N$ represents a distinct type of tokenised asset. Accordingly, a transaction in the blockchain economy is interpreted as a swap between a pair of tokens. Throughout the paper, subscripts g, h, g', h' are used to index asset types. It is also possible to interpret each $h \in N$ as a digital wallet address on the blockchain. Under this alternative interpretation, a transaction represents a payment from one wallet to another.

We say that tokens g and h are *directly linked* (or *directly connected*) if direct barter between them, say trading g for h or vice versa, is permitted by the settlement infrastructure. For example, such a direct link may correspond to a payment channel between two wallets in the Bitcoin Lightning Network or to a liquidity pool established on a decentralised exchange.

The transaction-settlement infrastructure is therefore represented by a network $\mathbf{G}$, where $\mathbf{G} : N \times N \to \{0, 1\}$ satisfies

$$\mathbf{G}(g, h) = \mathbf{G}(h, g) = \begin{cases} 1 & \text{if } g \text{ and } h \text{ are directly linked,} \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

For ease of exposition, we often write $gh \in \mathbf{G}$ to denote $\mathbf{G}(g, h) = 1$ and use network and infrastructure structure interchangeably throughout this paper.

A pair of tokens (g, h) is said to be *linked* (either directly or indirectly) in network $\mathbf{G}$ if there exists a sequence of distinct tokens $\{g_1, g_2, \ldots, g_{z+1}\}$ such that (i) $g_1 = g$, $g_{z+1} = h$, and (ii) $g_s g_{s+1} \in \mathbf{G}$ for all $s = 1, \ldots, z$. Such a sequence, denoted by $a := \{g_1, g_2, \ldots, g_{z+1}\}$, is called a *path* from g to h. The length of the path is defined as z, i.e., the number of links along the path.

In general, there may exist multiple paths connecting g and h. Let L_{gh} denote

the set of all such paths. Moreover, we let l_{gh} denote the length of the shortest path connecting g and h. We are particularly interested in the path length of the shortest path as we will show later that they are directly related to cost-efficient routing.

Note that if g and h are directly linked, they are necessarily connected. The converse need not hold, as two tokens may be connected only indirectly via a path of length greater than one. A network $\mathbf{G}$ is called *connected* if every pair of tokens in the network is connected.

The size of a network $\mathbf{G}$, which we denote as $m(\mathbf{G})$, is given by the total number of links in this network, say $m(\mathbf{G}) = \frac{1}{2}\sum_{g=1}^{n}\sum_{h\neq g} G(g,h)$. We often use $m(\mathbf{G})$ and m interchangeably in this paper. Among all feasible networks, m has the lower bound with $m = 0$ that implies an empty network and a upper bound with $m = \frac{1}{2}n(n-1)$ that implies a point-to-point or a complete network. In a point-to-point network, every pair of tokens can be serviced by direct links. We are also interested into another important network, which we refer to as a hub-spoke network. For a hub-spoke network of size m, it means that there exists a hub-token h such that $\sum_{g\in N} G(h,g) = m$.

3.2 Demand

Let q_{gh} denote the size of barter demand in the (g,h) token-pair market, i.e., the number of transaction orders submitted by blockchain users who wish to trade token g for token h. We aggregate these transactions into a (g,h) token-pair market because they share a common origin token g and a common destination token h in the network.

For each transaction from the (g,h) market, in our view, the length of the shortest path connecting g and h in the network directly determines the trading cost taken by the users in the (g,h) market. Accordingly, we let the demand q_{gh} be a function of l_{gh} and impose the following assumption.

Assumption 1. *For every token-pair market (g,h), demand q_{gh} is determined by a*

function $q(\cdot) : \mathbb{Z}_+ \to [0, \infty)$. *Moreover,* $q(\cdot)$ *is decreasing and follows* $q(l_{gh}) = 0$ *whenever* $l_{gh} \geq \bar{z} > 2$.

For a well-defined demand function $q(\cdot)$, we impose $q(l_{gh}) := 0$ and $l_{gh} := 0$ if any token pair (g, h) is disconnected in the network. Collectively, we denote by

$$\mathbf{q}(\mathbf{G}) := \{q(l_{gh})\}_{g,h \in N}$$

the array of barter demands generated in the blockchain economy.

The idea behind the above assumption is that the longer the shortest path connecting g and h, the more intermediate barters along the path are required to complete each generic barter between g and h. As a result, more block space is consumed. This increases the total gas cost paid by users or, equivalently, decreases the size of barter orders submitted from each market. This explains why we let demand in the economy be a function of the shortest-path length in the network.

Our demand specification abstracts from several realistic considerations. In particular, demand in each token-pair market depends solely on the shortest-path length connecting the tokens. The function $q(\cdot)$ is assumed to be homogeneous across token pairs and independent across markets. These restrictions are imposed primarily for tractability, as the central focus of the paper is the cost-efficient design of the optimal network structure. We revisit and partially relax these assumptions in Section 6.

3.3 Order Flow Allocation

Rather than settling generic barter demands between g and h directly within the network, it may be cost-minimising for the mediator to re-route transactions through an intermediate token g'. We now formalise this mediator's order-flow allocation decision. This requires additional notation.

First, for each demand order arising from (g, h) market, the mediator must determine which path(s) within her infrastructure network each order follows. This issue is particularly relevant when set L_{gh} is not a singleton. Let $\Delta_{gh}(\cdot)$ denote the mediator's allocation decision over L_{gh}, where $\Delta_{gh}(a)$ represents the proportion of transaction volume in the (g, h)-market routed through path $a \in L_{gh}$. Clearly, $\Delta_{gh}(a) \geq 0$ and $\sum_{a \in L_{gh}} \Delta_{gh}(a) = 1$.

Next, we let $\mathbf{\Delta} := \{\Delta_{gh}\}$ be the allocation profile across all token-pair markets. This profile uniquely determines the order flow on each link $gh \in \mathbf{G}$, which we can thus denote by $Q_{gh}(\mathbf{\Delta})$. As we will show later in the paper that there exists one unique optimal allocation profile Δ^*. Hence, with slight abuse of notation, we write the induced link-level order flow as $Q_{gh}(\mathbf{G})$ for simplicity.

It is worth noting that the link-wise order flow size $Q_{gh}(\mathbf{G})$ has a weak relation to the generic market-wise demand $q(l_{gh})$, even though both depend on the network structure. The former reflects the mediator's endogenous routing decision in response to aggregate demand $\mathbf{q}(\mathbf{G})$, whereas the latter captures users' primitive demand in the local (g, h)-market.

3.4 Network costs

In the model, the cost for a transaction-settlement infrastructure network $\mathbf{G}$ involves two additive components: a fixed cost and a variable cost. Specifically, the cost incurred by the mediator to operate a $gh \in \mathbf{G}$ link in the network to route a size of order flow $Q_{gh} + Q_{hg}$ is given by

$$c_0 + f(Q_{gh} + Q_{hg}). \tag{2}$$

Above, $c_0 > 0$ is the amount of fixed costs to form a new link. Having such a

new link gh in the network permits direct barter between g and h. Otherwise, barters between g and h can only be settled via a sequence of barters with other tokens. Term $f(Q_{gh} + Q_{hg})$ captures the variable operational cost on each link $gh \in \mathbf{G}$, with its exact value critically depending upon the size of order flow routing this link in both directions. Furthermore, we make one following standard assumption.

Assumption 2. *The variable cost function* $f : [0, \infty] \to [0, \infty]$ *is continuous, strictly increasing, concave, and follows* $f(0) = 0$.

Under this assumption, we can now specify the network cost incurred by the mediator when routing a demand $\mathbf{q}$ under allocation $\mathbf{X}$ in a $\mathbf{G}$, which is given by

$$\sum_{g=1}^{h} \sum_{h>g} f\Big(Q_{gh}(\mathbf{\Delta}) + Q_{hg}(\mathbf{\Delta})\Big) + c_0 m(\mathbf{G}). \tag{3}$$

Our cost structure (2) deserves several remarks.

First, the fixed-cost component captures the one-off cost incurred by the mediator when creating a new barter channel. For example, this cost may be interpreted as the initial infrastructure expenditure required to open a new payment channel between two wallets[7] in the Bitcoin Lightning protocol.

Second, the continuity and monotonicity of the variable-cost function f are economically and technologically plausible, considering the data-driven nature of blockchain systems. A blockchain network is built on a decentralised ledger to enable trustless trading. Adding a bundle of transactions to the canonical ledger requires the selected mediator to verify every transaction and then broadcast them to peer mediators in the blockchain network. Consequently, block production costs depend critically on the number of transactions in each block, as this directly determines the time and computational resources required for storage, downloading, and verification by all decentralised

[7]In this interpretation, the network connecting tokens becomes a network of wallets.

ledgers in the network. We therefore model block production costs as a function of the size of transaction flow, abstracting from transaction-specific characteristics, such as trading price and trading direction.

Third, the concave variable cost f enables a system-wide economies of scale property. Such a property in the block-production cost may arise either from declining marginal computational costs incurred by the mediator when verifying additional transactions for identical token-pairs or from the spreading of the fixed cost over a larger number of transactions.

3.5 An Optimisation Problem

We now formalise the optimisation problem for the monopolistic mediator.

Let $R(q_{gh})$ denote the revenue earned by the mediator from a token-pair market (g, h) with the demand q_{gh}. Due to Assumption 1, under which demands across token-pair markets are independent from each other, the total network revenue in $\mathbf{G}$ across all token-pair markets then follows an additive structure. That is,

$$R(\mathbf{G}) := \sum_{g=1}^{n} \sum_{h \neq g} R(q_{gh}) = \sum_{g=1}^{n} \sum_{h \neq g} R\big(q(l_{gh})\big). \tag{4}$$

Next, the optimisation problem faced by the monopolistic mediator is to choose a network structure $\mathbf{G}$ to maximise

$$\pi(\mathbf{G}) := R(\mathbf{G}) - \sum_{g=1}^{n} \sum_{h>g} f\Big(Q_{gh}(\mathbf{G}) + Q_{hg}(\mathbf{G})\Big) - c_0 m(\mathbf{G}). \tag{5}$$

Before we proceed with the solution characterisationn to the problem, note that if the fixed cost c_0 is sufficiently large, the mediator would optimally choose the null network. To rule out this trivial outcome, let $\pi^*(x)$ denote the maximised value of

$\pi(\mathbf{G})$ for each $c_0 := x > 0$, and let $\bar{c}$ be the smallest solution to $\pi^*(x) = 0$. Throughout the paper, we impose $0 < c_0 < \bar{c}$. Furthermore, we let

$$R(q_{gh}) \geq R(q(2)) > \bar{c} \quad \forall g, h \in N \tag{6}$$

The idea is that the revenue from every token-pair market has to be sufficiently large. Otherwise, the mediator may give up some n' markets less profitable but focus on serving other $n-n'$, more profitable markets. We assume $R(q_{gh})$ has to at least exceed the fixed link-formation cost c_0 (but not necessarily larger than fixed cost plus variable cost) when the network is complete and a hub-spoke. The reason we consider these two special network structures is that they are proven to be the only possible optimal networks in the model. Therefore, under (6), we indeed ensure that the mediator is free to choose any pattern of networks among tokens at the optimum, thereby providing a clean idea of how the interaction between fixed and variable costs shapes the resulting optimal network structure.

4 Main Results

In this section, we characterise the optimal network $\mathbf{G}$. By restricting network size to $n-1$, we characterise that the optimal network of size $n-1$ is a hub-spoke network. Then, we extend the discussion to networks with size $m(\mathbf{G})$ following $n-1 < m(\mathbf{G}) < \frac{1}{2}n(n-1)$. We find that networks with an intermediate size $m \in (n-1, \frac{1}{2}n(n-1))$ are always suboptimal to the mediator. Lastly, by comparing the two corner solutions between the hub-spoke network of size $n-1$ and the complete network of size $\frac{1}{2}n(n-1)$, we characterise the specific cost condition under which one network emerges an the optimal.[8]

[8]It is important to note that Assumption (6) implies that the optimal network must be connected. This fact yields a lower bound of $n-1$ on the network size, so there is no need to discuss networks of

It is worth emphasising that once the optimal network $\mathbf{G}^*$ is shown to display the hub-and-spoke structure, the emergence of a medium of exchange in the blockchain economy is established. In such a network, all transactions between non-hub tokens are intermediated exclusively through the hub. That is, the hub token functions as a unique medium of exchange in the economy.

4.1 Optimal Network without variable costs

We begin by examining an extreme case in which the effect of variable costs, or, equivalently, economies on the network size, is assumed away. That is, we assume the variable cost is negligible and impose a constant network revenue for each token-pair market. Despite its simplicity, this specification is useful for refining the set of optimal networks in a large economy. Formally, we obtain the following theorem.

Theorem 1. *For each token-pair market (g, h) in the network, suppose $R(q_{gh}) = R_0 > \bar{c}$ and $f(q_{gh}) = 0$. Then, a network $\mathbf{G}$ is optimal if and only if it is connected and of size $n - 1$.*

The intuition is straightforward: On the one hand, the network cost in the optimisation problem (5) changes to the total link-formation cost $c_0 m(\mathbf{G})$ under the zero variable cost assumption. On the other hand, the constant revenue assumption implies that the mediator can achieve the maximum network revenue given by $\sum_{g,h \in N} R_0$ as long as the network is connected. Of course, due to (6), it is clearly in the mediator's interest to have all token pairs connected in the network.

Therefore, among all connected networks, the problem faced by the mediator becomes: what is the minimum number of links (i.e., the network size) required for the network to be connected among n tokens? Note that many network structures connect

size less than $n - 1$. This fact is proved in the Appendix. The upper bound on the size of networks among n tokens is $\frac{1}{2}n(n-1)$, and this upper bound is achieved if and only if the network is complete.

all n tokens but with different sizes, including a complete network of size $\frac{1}{2}n(n-1)$, a cycle network of size n, and a chain or a hub-spoke network of size $n-1$. We prove that the minimum network size is $n-1$, thereby establishing Theorem 1.

The set of connected networks of size $n-1$, however, is not a singleton. For instance, they may take the form of a chain, a hub-and-spoke, or a multi-hub structure. Theorem 1 thus provides just one existence result. To establish uniqueness and obtain stronger results regarding the optimal network structure(s) of size $n-1$, we must generalise the cost structure and incorporate economies of scale into the model.

4.2 Optimal Network with a Minimal Size

Theorem 1 establishes that when variable costs are negligible, network revenue is independent of the particular routing of each barter to follow in the blockchain infrastructure network. This property no longer holds when the variable cost function f is positive, that is $f(x) > 0$ for any $x > 0$. Indeed, the economies of scale implied by a positive, concave variable cost function has a direct implication on the mediator's optimal order-flow allocation, thereby the revenue in the network as well as the structure of optimal network. Indeed, we show below that the set of optimal connected networks of size $n-1$ collapses to a singleton when the variable cost is no longer negligible. In particular, the unique optimal network exhibits the hub-and-spoke structure.

Theorem 2. *Suppose* $\mathbf{G}$ *is a connected network of size* $n-1$ *and the variable cost* f *is positive, that is,* $f(x) > 0$ *for all* $x > 0$. *Then* $\mathbf{G}$ *is optimal if and only if it is a hub-and-spoke network.*

In the presence of a positive concave variable cost, the mediator is incentivised to economise a given number of links by concentrating order flows on paths as few as possible. This implies that the mediator would allocate all barter orders from the same

token-pair market to a single path in the network. Also, it implies that the mediator has no incentive to allocate barter orders between these two directly linked tokens in the network to a longer path with any intermediate tokens. Armed with this optimal order flow configuration, we argue that the solution network of size $n-1$ turns out to be the one such that the aggregated length of paths connecting all token pairs is minimised. The above argument is valid because, for a given size of network, both the revenue and cost-saving favour shorter paths in the network. This cost-saving part is trivial. As for the revenue side, note that shorter path between two distinct tokens in the network increases the demand and thus increases the revenue in this token-pair market.

Among all networks of size $n-1$ among n connected tokens, the hub-and-spoke network uniquely minimises the total path lengths. This is due to a special structural feature only implied by the hub-and-spoke network. That is, amongst all networks of size $n-1$ that connect n tokens, only the hub-and-spoke network has all its shortest paths smaller than or equal to two: Every pair of non-hub tokens in the hub-and-spoke network is connected through exactly one intermediary, say a path of length of two, and every pair of hub token and non-hub token is directly connected, say a path of length of one, and there is no other pattern of connection in this network.

To conclude this section, we point out that the theorem here represents one corner solution as we restrict the network size to $m(\mathbf{G}) = n-1$. The other corner is obviously the complete network, which uniquely arises when the fixed link-formation cost c_0 is sufficiently small. Intuitively, as $c_0 \to 0$, adding links increases revenue by shortening paths while imposing negligible additional network costs. Hence, any non-complete network can be strictly improved by adding links at a expense of negligible cost, implying that the optimum is the complete network. One may then wonder under which cost conditions a network of intermediate size may arise at the optimum. We

explore this in the next section.

4.3 Optimal Network with an Intermediate Size

We now turn to networks of intermediate size, $n-1 < m < \frac{1}{2}n(n-1)$. Given that the complete network is optimal as $c_0 \to 0$, while the hub-and-spoke network is optimal when c_0 is sufficiently large,[9] one might then expect that networks of intermediate size become optimal for intermediate values of c_0.

Perhaps surprisingly, this reasoning is shown to be not correct in this section. We find that no interior network size can be optimal and any network of a size between $n-1$ and $\frac{1}{2}n(n-1)$ will be strictly dominated by either a hub-and-spoke network or a complete network.

Due to the involving effects of the general concave variable cost f on the order flow configuration $\mathbf{\Delta}$, we begin the analysis in this section with a special linear variable cost, and then, building upon that, extend it to a general concave variable cost. The idea is that a general concave cost function can be approximated by a collection of piecewise linear cost functions.

Linear variable cost For any variable cost function f with constant marginal cost $c > 0$, say $f(x) = cx$, the economics of scale implies that the network cost of serving each token-pair market (g, h) only depends upon the product between the number of submitted barter orders in this market q_{gh} and the number of links used by each of these barters, which we denote as ℓ_{gh}. Then, we can rewrite the monopolist's objective

[9] Indeed, as $c_0 \to \bar{c}$, the mediator strictly prefers networks of size $n-1$, whereby leading to the hub-and-spoke structure at the optimum.

(5) as below:

$$\pi(\mathbf{G}) = R(\mathbf{G}) - \sum_{g=1}^{n} \sum_{h>g} f\Big(Q_{gh}(\mathbf{G}) + Q_{hg}(\mathbf{G})\Big) - c_0 m(\mathbf{G})$$
$$= \sum_{g=1}^{n} \sum_{h\neq g} \Big[R(q(\ell_{gh})) - c\,\ell_{gh}\,q(\ell_{gh})\Big] - c_0 m(\mathbf{G}) =: \hat{\pi}(m, cx) - c_0 m(\mathbf{G}). \quad (7)$$

Here, we denote by $\hat{\pi}(m, cx)$ the network revenue net of the fixed network costs in a network of size m with a linear variable cost specification $f(x) = cx$. Next, by taking the linear cost as given, we vary the network size m and yield the following result.

Lemma 1. *For a network of size $n-1 < m < \frac{1}{2}n(n-1)$, $\hat{\pi}(m, cx)$ is bounded from above by*

$$2m \times h(c) + [n(n-1) - 2m] \times h(2c), \quad (8)$$

where $h(\ell c) := R(q(\ell)) - c\,\ell\,q(\ell)$ and $\ell = 1, 2$.

This lemma states that there exists an upper bound for $\hat{\pi}(m, cx)$ and this upper bound is attainable. To understand this result, note that for each connected network of size m there are $2m$ directly linked token-pair markets, as well as $2\left[\frac{1}{2}n(n-1) - m\right]$ indirectly connected token-pair markets. In those directly linked markets, barter transactions are executed without routing through intermediate tokens, which follows immediately from economies of scale in the cost structure. For indirectly connected markets, their barter transactions must be routed along some paths with length greater than one in the network. To maximise profits in these indirectly connected markets, the mediator will try to assign their barter orders to the shortest paths in the network. Clearly, the shortest such paths have a length of two, if such networks exist.

Indeed, this network can be constructed by firstly forming a hub-and-spoke network of size $n-1$, and then adding extra spoke-spoke links until the total number of links

reaches $m > n-1$. In such a network, all token pairs are either directly connected or indirectly connected via a shortest path of length two.

We now substitute this upper bound (8) back into the monopolist's objective. This yields

$$\pi(\mathbf{G}) = n(n-1)h(2c) + m\Big[2\big(h(c) - h(2c)\big) - c_0\Big].$$

The extra profit earned by the monopolist from adding one additional link in the network is thus given by $2\big(h(c) - h(2c)\big) - c_0$. Built upon this, we characterise the optimal network structure in the linear variable cost case in the following way:

- The optimal network is a star network of size $n-1$ if $h(c) - h(2c) < \frac{1}{2}c_0$;
- Otherwise, it is a complete network of size $\frac{1}{2}n(n-1)$.

It shall be clear now that any network of size m with $n-1 < m < \frac{1}{2}n(n-1)$ can never be optimal when the variable cost function is linear.

Concave variable cost We now move to the general concave variable cost case. Perhaps surprisingly, we find that the upper bound as defined in the above linear case still holds for the general concave variable cost case. Moreover, this upper bound for the general concave variable cost is still attainable when the network is a hub-spoke network or a complete network.

Let us denote $\hat{\pi}_S(f)$ and $\hat{\pi}_F(f)$ be the respective profit component net of fixed link-formation cost earned by the mediator in a star network and a complete network under the setting with a general concave variable cost function f. Then, as a sufficient condition, what remains to show is the following:

$$\hat{\pi}(m, f) \leq \hat{\pi}_S(f) + \frac{m-(n-1)}{\frac{1}{2}n(n-1)-(n-1)}\Big(\hat{\pi}_F(f) - \hat{\pi}_S(f)\Big) \tag{9}$$

for any $n-1 < m < \frac{1}{2}n(n-1)$. To understand the sufficiency, it is helpful to note

that the right side of (9) is a linear combination between $\hat{\pi}_S(f)$ associated with the hub-spoke network and $\hat{\pi}_F(f)$ associated with the complete network.

Characterisations related to $\hat{\pi}_S(f)$ and $\hat{\pi}_S(f)$ for a general variable cost function f are, however, not trivial. Instead of deriving them directly, we utilise our earlier results in the linear cost case by approximating any given concave function f via a collection of upper piecewise linear functions $\{f_i(x)\}_{i\in I} := \{c_i x + b_i\}_{i\in I}$. Formally, we obtain the following theorem.

Theorem 3. *For any network* $\boldsymbol{G}$ *of size* m *with* $n-1 < m < \frac{1}{2}n(n-1)$*, the inequality (9) always holds. Therefore, the network* G *can never be optimal.*

The result here deserves some attention. Under a concave variable cost, the marginal cost of serving one additional transaction order in the economy decreases as the density of the transaction flow increases. One might, therefore, pre-assume that the monopoly rent from adding a new connection increases as the network expands, and that the optimal network size should depend upon the concavity of the cost function. Our result in this theorem, however, shows that this intuition is incorrect. In particular, if adding one more connection to a star network of size $n-1$ is profit-improving for the monopolist, then it is optimal to keep adding connections until the network becomes complete, a fully point-to-point connected network.

4.4 Hub-spoke Network vs Complete Network

Having established the optimal network for different network sizes, we indeed have obtained a monopoly rent frontier in network size. That is, the optimal network is either a hub-spoke network of size $n-1$ or a complete network of size $\frac{1}{2}n(n-1)$.

As for the conditions under which network structure emerges at the optimum, we obtain the following theorem.

Theorem 4. *For a general concave and strictly increasing variable cost function f, there exists a positive constant c^* such that*

- *a hub-spoke network of size $n-1$ is optimal if and only if $c_0 > c^*$;*
- *a complete network of size $\frac{1}{2}n(n-1)$ is optimal if and only if $c_0 < c^*$.*

The emergence of either optimal network follows directly from the break-even condition under which the monopolist becomes indifferent between operating a complete network and a hub-spoke network. That is,

$$\hat{\pi}_F(f) - \frac{1}{2}n(n-1)c^* = \hat{\pi}_S(f) - (n-1)c^*.$$

Rearranging this break-even condition immediately yields a level of fixed link-formation cost c^* as below:

$$c^* = \frac{2}{(n-1)(n-2)}\big[\hat{\pi}_F(f) - \hat{\pi}_S(f)\big]. \tag{10}$$

Hence, this proves the theorem above.

The takeaway message of this paper is now clear: The conventional wisdom regarding the necessity of a medium of exchange remains valid even in a nominally intermediary-free blockchain economy, as long as on-chain block space remains scarce. However, the underlying economic mechanism here that drives the emergence of a medium of exchange differs fundamentally from the traditional setting. In classical monetary theory, the role of money arises from search frictions associated with the double coincidence of wants. In contrast, in a blockchain-based barter environment, the demand for a medium of exchange is driven by infrastructure cost efficiency considerations, say the scarcity of block space and the associated transaction inclusion costs faced by the (randomly selected) mediator.

5 Further discussion

This section discusses several factors that we abstract away in our model and explains why none of them overturns the emergence of a medium of exchange in a blockchain economy.

Notion of optimality. Our notion of optimality concerns the efficiency of blockchain infrastructure with respect to base gas costs. It is largely orthogonal to alternative fee-design mechanisms commonly observed in practice, practically the upfront-fee mechanism.

Under the upfront-fee mechanism, a user attaches to her transaction a total fee equal to the base gas fee plus a user-chosen priority fee. The base fee is fixed at the block level, whereas the priority fee acts as a tip that determines execution priority. Users with time-sensitive or information-sensitive trades may therefore bid aggressively in order to secure earlier inclusion.

Such fee bidding may generate increasing marginal inclusion costs, potentially contrasting with the concave cost structure assumed in the model studied in this paper. In particular, strategic bidding competition can induce discrete jumps in effective inclusion costs when users outbid one another. We abstract from this dimension by assuming an exogenous demand schedule $q(\cdot)$ because our objective is not to model priority competition per se, but to isolate the infrastructure-driven forces, say the base gas costs, underlying the emergence of a medium of exchange.

Natural monopoly structure. Finally, we argue that a hub-spoke infrastructure naturally arises in an unregulated blockchain economy in the face of competition from entrants.

Consider a star network of size $n-1$ links. Total revenue is $\sum_{g,h\in N} R(p^*_{gh})$, while

total fixed cost equals $(n-1)c_0$. Profitability requires

$$\sum_{g,h \in N} R(p^*_{gh}) > (n-1)c_0.$$

To get a sense of implications, let R^* denote a lower bound on revenue per token-pair market, whereby the profitability condition can be relaxed to

$$n(n-1)R^* > (n-1)c_0 \quad \Longleftrightarrow \quad R^* > \frac{c_0}{n}.$$

By contrast, a start-up entrant operating a single token-pair market earns revenue R and incurs fixed cost c_0, so profitability requires

$$R > c_0.$$

When fixed costs are large relative to marginal costs, operating an integrated network is substantially more profitable—on a per token-pair basis—than operating isolated links. A hub–spoke architecture internalises cross-link externalities: adding a new spoke increases traffic over existing links at low marginal cost while generating sizeable incremental revenue across the network. These scale effects are captured by the incumbent hub but not by a stand-alone entrant.

By contrast, a complete network of size $n(n-1)$ does not enjoy the advantage but achieves a profitability condition that resembles that of an entrant operating isolated links. As a result, when fixed infrastructure costs are the dominant component of total costs, the hub–spoke network strictly dominates the complete network in deterring the start-up entrant We therefore expect the hub–spoke infrastructure to emerge as the natural outcome in an unregulated blockchain economy characterised by substantial fixed costs and strong economies of scale.

Incentive of users to deviate Thus far, our characterisation is built upon the demand pattern of users in response to the shortest path length in the network. However, given the monopolistic power of the mediator in organising block content, it appears plausible to introduce price-discrimination among users according to their specific demands in the model.

This introduces a two-stage model. In stage-1, the mediator has to determine not only the network $\mathbf{G}$ that connects every token pair but also an array of prices $\mathbf{p} := \{p_{gh}\}$ with each p_{gh} denoting the price charged by the mediator for including one barter order submitted from the token-pair market (g, h). Having observed $(\mathbf{G}, \mathbf{p})$, the users in stage 2 determines their demands $\mathbf{q}(\mathbf{G}, \mathbf{p}) := \{q_{gh}(\mathbf{p})\}_{g,h\in N}$. Note that the user in stage 2 now may split her barter demand between g and h via an intermediate token g' if

$$p_{gg'} + p_{g'h} < p_{gh}.$$

Taking the above strategic response into the consideration, we see that if there exists a solution $(\mathbf{G}^*, \mathbf{p}^*)$ to the optimization problem faced by the mediator in stage-1, the demand $\mathbf{q}(\mathbf{p}^*, \mathbf{G}^*)$ induced in stage 2 must coincide with the demand exactly anticipated by the mediator when solving the optimisation problem of stage 1.

Hence, showing that the hub-spoke network under $c > c^*$ and the complete network under $c < c^*$ are still equilibria becomes equivalent to showing that there is no profitable deviation for users from the prescribed routing in each corresponding optimal network.

Let us begin with the case $c_0 < c^*$. In this region, the optimal infrastructure is the complete network. One can easily show that the optimal fee is identical across all token pairs:

$$p^*_{gh} \equiv p^* \quad \text{for all } g, h \in N.$$

Under this pricing rule, a user trading from token g to h would incur cost $2p^*$ if she splits her transaction through an intermediate token g', whereby a higher trading cost $2p^* > p^*$. Hence, no user has an incentive to deviate from the generic barter demand.

Consider next the case $c_0 > c^*$, in which the optimal infrastructure is a hub-spoke network. Consequently, barter demands in all token pairs sharing the same path length would face one identical price. A spoke–spoke barter requires two links while a hub–spoke barter requires only one. Hence, the optimal price array chosen by the mediator in stage 1 features two distinct prices. Denote by p^*_{sh} the price charged by the mediator for each hub-spoke barter transaction, and by p^*_{ss} the price for each spoke-spoke barter. We obtain the following result.

Result 1. *Suppose inverse demand is concave. Then, users have no incentive to deviate from their generic barter demands under the optimal hub-spoke network* $(\mathbf{G}^*, \mathbf{p}^*)$, *as* $p^*_{sh} < p^*_{ss} < 2p^*_{sh}$.

The first inequality is intuitive: Establishing a path between two spokes requires two links, whereas a hub–and–spoke trade requires only one. Consequently, the mediator optimally sets $p^*_{ss} > p^*_{sh}$. This rules out any profitable deviations in which a user wishing to trade a spoke for the hub first trades to another spoke and then re-trades with the hub.

The second inequality, under the mild assumption on the concavity of the inverse demand, states that for the stage-1 solution $(\mathbf{G}^*, \mathbf{p}^*)$, the price charged by the mediator for a spoke-spoke barter demand is less than twice that for a hub-spoke barter demand. Hence, a user wishing to submit a spoke-spoke barter demand does not benefit from the splitting deviation by submitting two hub–spoke barter demands, as such deviation would cost $2p^*_{sh}$, which strictly exceeds p^*_{ss}.

Hence, users optimally adhere to the prescribed routing and submit their generic barter demands in stage 2. The induced demand coincides with the mediator's stage-1

anticipation, thus reconfirming that the optimality of the hub-spoke network even if the mediator does the price discrimination across token pairs.

MEV considerations. One may further ask whether maximal extractable value (MEV) alters the attractiveness of the hub-spoke structure. Existing work, such as Zhang et al. (2025), shows that the computational complexity of solving the MEV problem depends critically on the market structure. When tokens are organised in a hub–spoke structure (a Fisher-market formulation with a commonly accepted numéraire token), the MEV problem can be solved in polynomial time. In contrast, under a fully decentralised Arrow–Debreu structure without a designated hub, the corresponding optimisation becomes NP-hard.

This complexity distinction strengthens the mediator's incentive to organise trade around a hub token. A hub–spoke architecture not only economises on infrastructure costs, as shown in our model, but also reduces the computational costs associated with transaction ordering and MEV extraction. Hence, MEV considerations reinforce rather than undermine the emergence of a monetary hub.

6 Conclusion

This paper challenges the conventional presumption that blockchain economies have to be inherently intermediary-free. Although block producers (mediators) are randomly selected across block slots, the mediator chosen in any given slot holds temporary monopoly power over transaction inclusion and block composition.

We show that when block production exhibits economies of scale in variable costs and block space remains scarce in new payment channel creation, this temporary monopoly power induces a hub-spoke transaction network within each block. The resulting network architecture endogenously generates a hub token that performs the

function of a medium of exchange.

Our analysis extends the classic theory of the emergence of money to a blockchain-based, nominally intermediary-free barter economy. In contrast to traditional environments where money arises from search frictions and the double coincidence of wants, the demand for a medium of exchange here is driven by the consideration of cost-efficient blockchain infrastructure in transaction settlement.

7 Appendix

7.1 Proof of Theorem 1

Theorem 1 is a direct implication of one fact as follows.

Fact 1. *If a network has a size $m \leq n-1$, then the maximum number of connected token pairs in this network is $m(m+1)$.*

Proof. The proof has three steps.

Step 1: It is easy to check that the bound is attained by a star network of size m.

Step 2: In this step, we want to prove that for the bound to be attained, all tokens have to be connected in the network. That is, the bound $m(m+1)$ can be attained only if the network of size m is connected.

We proceed by induction. The above claim is obvious for $m = 1$. Suppose it is true for $m = k-1$. For the bound to be attained when $m = k$, let us now prove that all tokens in the network must be connected, which is to show that if $G(g,h) = 1$ and $G(g',h') = 1$, then g, h, g' and h' are connected.

We can decompose the set of N into a collection of disjoint sets N_s such that for any $g \in N_s$, $G(h,g) = 0$ if $h \notin N_s$. Denote by k_s the number of direct connections in N_s. Suppose that $k_s < k$ for all s. Then we get

$$k^2 = (\sum_s k_s)^2 = \sum_s k_s^2 + 2\sum_s \sum_{s \neq s'} k_s k_{s'}.$$

This implies that

$$k(k+1) = \sum_s k_s^2 + \sum_s k_s + 2\sum_s \sum_{s \neq s'} k_s k_{s'} = \sum_s k_s(k_s+1) + 2\sum_s \sum_{s \neq s'} k_s k_{s'},$$

whereby results in

$$\sum_s k_s(k_s+1) < k(k+1).$$

This yields the contradiction because we have seen that $k(k+1)$ is attainable (see step 1). In other words, all tokens are connected in the network when the bound is attained.

Step 3: As the last step, what remains to show is that the maximum number of connected pairs in the connected network of size m is $m(m+1)$.

This proof is built upon the following fact.

Fact 2. *A connected network with n nodes must have at least $n-1$ edges.*

As long as this fact holds, we can naturally have that for a given m edges, the maximum number of vertices a connected network can have is $n = m+1$. In a single connected network with $m+1$ vertices, every pair of vertices is connected. Therefore, the number of such pairs is given by $2 \times \binom{m+1}{2} = m \times (m+1)$, where the multiplier 2 in front of the combination formula is because transactions are assumed to be single-directional in the paper.

Therefore, what remains here is to prove the above fact and let us now focus on that.

First, consider a connected network with $n+1$ nodes/vertices. If the network contains any cycles, one can remove an edge from a cycle without breaking the network's connectivity. Repeating this process until no cycles remain yields a connected acyclic network, which we often refer to as a tree.

Next, we claim a tree T with $n+1$ nodes has exactly n edges.

To prove by induction. It is trivial when $n = 0$. Assume that the claim is true for the tree with $n = k+1$ nodes, where $k \geq 0$ is some integer.

Now consider a tree network T with $n = k+1+1 = k+2$ nodes. Note that the tree

T has at least two leaves (as the longest path in the network contains two endpoints). Let v be a leaf in T. By removing vertex v and the unique edge connected to it from T, denoted by vv', we obtain a new network T'. Now, T' has $k+2-1=k+1$ nodes and remains connected because removing a leaf would not affect the connectivity of a tree network. Clearly, T' still has no cycles, as removing an edge doesn't create cycles. Therefore, T' is a tree (i.e., by definition, a connected acyclic graph) and it has $k+1$ nodes.

Remark that T was created by adding one vertex (v) and one edge (vv') to T'. Thus, the total number of edges in T is the number of edges in T' plus one: $(k+1-1)-1=k+1$.

Proof completes.

□

7.2 Proof of Theorem 2

Proof. What we have to show is that a connected network $\mathbf{G}$ of size $n-1$ has no path with more than two direct connections if and only if $\mathbf{G}$ is a connected star network.

$\Leftarrow$: It is trivial.

$\Rightarrow$: Proof by contradiction. Suppose $\mathbf{G}$ is not a star network. As $\mathbf{G}$ is connected and is of size $n-1$, the set of paths L_{gh} that connects the token-pair (g,h) must be a singleton. To see this point, note that a connected network with n nodes and $n-1$ direct connections is a tree, whereby there exists a unique path from each origin g to each destination h.

Next, for connected network $\mathbf{G}$, we claim there exists at least one g whose neighbourhood set $I_g=\{h|G(g,h)=1\}$ satisfies $2\leq|I_g|<n-1$. If not, then all nodes in the network have degree 1 or $n-1$. This contradicts the assumption that $\mathbf{G}$ is connected and that there is no path in $\mathbf{G}$ with more than two direct connections.

Therefore, the connectivity of the network $\mathbf{G}$ implies that there exist a node k such that $k \notin I_g$ and $k \neq g$ and a node h such that $h \in I_g$ and $G(k,h) = 1$. As $|I_g| \geq 2$, we can find another node $g_1 \in I_g$ with $g_1 \neq h$. Clearly, $G(g,g_1) = 1$ and $G(g,h) = 1$. All the above arguments together imply

$$G(k,h) = G(h,g) = G(g,g_1) = 1.$$

Clearly, $k \neq g_1$ as $k \notin I_g$, whereby results in a path with three direct connections. This contradicts the pre-assumption that no path in $\mathbf{G}$ has more than two direct connections. Therefore, $\mathbf{G}$ is a star or a hub-spoke network. We complete the proof.

□

7.3 Proof of Theorem 3

Proof. Before proceeding with the proof, let us first specify the validity of approximating a concave f via piecewise linear functions $\{f_i(x)\}_{i\in I}$, which is formally established below.

Indeed, for any network $\mathbf{G}$ and any concave, strictly increasing function f, the separating hyperplane theorem immediately implies the existence of a piecewise-linear function $f_{PL}(\cdot)$ such that

$$f_{PL}(x) \geq f(x) \quad \forall x \geq 0 \quad \textit{and} \quad \hat{\pi}(m,f) = \hat{\pi}(m,f_{PL}),$$

where $f_{PL}(x) := \min\limits_{i\in I}\{c_i x + b_i\}$ with $I := \{1,2,...,m\}$.

For illustration of the validity of this approximation, we depict below an example concave function $f(x) = 2 - \frac{1}{2}(2-x)^2$ and its relation with its five upper piecewise linear envelope segments.

With that, it is now sufficient to construct an upper bound for the $\hat{\pi}(m,f_{PL})$. To

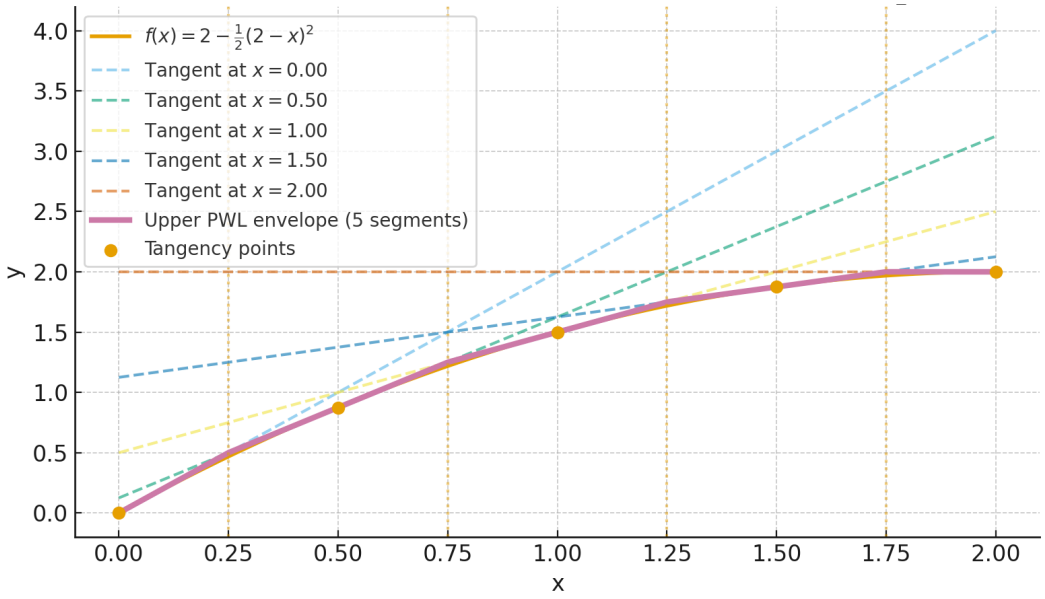


Figure 1: As an illustration, here we approximate a concave function $f(x) = 2 - \frac{1}{2}(2 - x)^2$ through five upper piecewise-linear envelope segments.

proceed, for all indirectly connected markets, let us consider the setting for the cost incurred by the initial and final connection. We define a_{ij} as an indicator function that takes the value of 1 if and only if the path connecting tokens g and h contains $i \in I$ and $j \in I$ as its respective origin connection and destination connection. Notably, this "if and only if" relationship is guaranteed by the economies of scale condition. As a result, for $f_{PL}(Q) = min_{i \in I}\{c_i Q + b_i\}$, we obtain

$$\hat{\pi}(m, f_{PL}) \leq \sum_{i \in I}(2h(c_i) - b_i) + \sum_{i,j \in I} a_{ij} \times h(c_i + c_j).$$

Next, due to the convexity of $h(\cdot)$, we notice that

$$\sum_{i,j \in I} a_{ij} h(c_i + c_j) \leq \frac{1}{2} \sum_{i,j \in I} a_{ij}[h(2c_i) + 2h(2c_j)],$$

whereby implies that

$$\hat{\pi}(m, f_{PL}) \leq \sum_{i \in I}(2h(c_i) - b_i) + \frac{1}{2}\sum_{i,j \in I} a_{ij}[h(2c_i) + 2h(2c_j)], \tag{11}$$

To proceed, inspired by (9), let us characterise $\hat{\pi}_S(f_i)$ and $\hat{\pi}_F(f_i)$, respectively. Recall that they represent the total network profits earned by the monopolist in a star network of size $m = n - 1$ and in a complete network (of size $\frac{1}{2}n(n-1)$), respectively. Therefore, we have

$$\hat{\pi}_S(f_i) = \underbrace{(n-1)[2h(c_i) - b_i]}_{\text{directly connected markets}} + \underbrace{(n-1)(n-2)h(2c_i)}_{\text{indirectly connected markets}} \tag{12}$$

and

$$\hat{\pi}_F(f_i) = \underbrace{\frac{1}{2}n(n-1)[2h(c_i) - b_i]}_{\text{only directly connected markets}} \tag{13}$$

By taking the maximum over both sides and applying the above piecewise linear approximation, it becomes clear that

$$\hat{\pi}_S(f) \geq max_{i \in I}\{\hat{\pi}_S(f_i)\} \quad and \quad \hat{\pi}_S(f) \geq max_{i \in I}\{\hat{\pi}_F(f_i)\}$$

Given that, it becomes evident that establishing (9) is equivalent to verifying the following relation

$$\begin{aligned}\sum_{i \in I}(2h(c_i) - b_i) + \frac{1}{2}\sum_{i,j \in I} a_{ij}[h(2c_i) + 2h(2c_j)] \leq max_{i \in I}\{\hat{\pi}_S(f_i)\} + \\ + \frac{m - (n-1)}{\frac{1}{2}n(n-1) - (n-1)}\Big(max_{i \in I}\{\hat{\pi}_F(f_i)\} - max_{i \in I}\{\hat{\pi}_S(f_i)\}\Big)\end{aligned} \tag{14}$$

We simplify the right side of the above inequality as follows

$$
\begin{aligned}
&max_{i\in I}\{\hat{\pi}_S(f_i)\} + \frac{m-(n-1)}{\frac{1}{2}n(n-1)-(n-1)}\Big(max_{i\in I}\{\hat{\pi}_F(f_i)\} - max_{i\in I}\{\hat{\pi}_S(f_i)\}\Big)\\
=&\frac{n(n-1)-2m}{(n-1)(n-2)}max_{i\in I}\{\hat{\pi}_S(f_i)\} + \frac{2m-2(n-1)}{(n-1)(n-2)}max_{i\in I}\{\hat{\pi}_F(f_i)\}\\
=&\left\{\frac{n(n-1)-2m}{(n-1)(n-2)}(n-1) + \frac{2m-2(n-1)}{(n-1)(n-2)}\frac{1}{2}n(n-1)\right\}max_{i\in I}\{2h(c_i)-b_i\}\\
&+\left\{\frac{n(n-1)-2m}{(n-1)(n-2)}\times(n-1)(n-2)\right\}max_{i\in I}\{h(2c_i)\}\\
=&m\times max_{i\in I}\{2h(c_i)-b_i\} + [n(n-1)-2m]\times max_{i\in I}\{h(2c_i)\},
\end{aligned}
$$

where we apply (12) and (13) to obtain the last second equality in the above.

Next, what remains to show is that

$$
\begin{aligned}
&\sum_{i\in I}(2h(c_i)-b_i) + \frac{1}{2}\sum_{i,j\in I}a_{ij}[h(2c_i)+2h(2c_j)]\\
\leq\ & m\times max_{i\in I}\{2h(c_i)-b_i\} + [n(n-1)-2m]\times max_{i\in I}\{h(2c_i)\}. \qquad (15)
\end{aligned}
$$

The validity of the above relation is trivial if one notices that the symmetry between i and j, as well as

$$
\sum_{i\in I}1 = |I| = m \quad \text{and} \quad \sum_{i,j\in I}a_{ij} \leq 2\times[\frac{1}{2}n(n-1)-m] = n(n-1)-2m,
$$

where the inequality in the above is because each a_{ij} is associated with one unique indirectly connected token-pair market, thereby resulting in a maximum aggregated number of token-pairs equal to $n(n-1)-2m$.

The proof is complete. □

References

Abhijit V Banerjee and Eric S Maskin. A walrasian theory of money and barter. *The Quarterly Journal of Economics*, 111(4):955–1005, 1996.

Bruno Biais, Christophe Bisiere, Matthieu Bouvard, and Catherine Casamatta. The blockchain folk theorem. *The Review of Financial Studies*, 32(5):1662–1715, 2019.

Eric Budish. Trust at scale: The economic limits of cryptocurrencies and blockchains. *The Quarterly Journal of Economics*, 140(1):1–62, 2025.

Agostino Capponi, Ruizhe Jia, and Sveinn Olafsson. Proposer-builder separation, payment for order flows, and centralization in blockchain. *Payment for Order Flows, and Centralization in Blockchain (February 12, 2024)*, 2024.

Jonathan Chiu and Thorsten V Koeppl. The economics of cryptocurrency: Bitcoin and beyond. *Canadian Journal of Economics/Revue canadienne d'économique*, 55 (4):1762–1798, 2022.

Lin William Cong, Ye Li, and Neng Wang. Tokenomics: Dynamic adoption and valuation. *The Review of Financial Studies*, 34(3):1105–1155, 2021.

Augustin Cournot. Researches into the mathematical principles of the theory of wealth. In *Forerunners of Realizable Values Accounting in Financial Reporting*, pages 3–13. 1838.

Zahra Ebrahimi, Maxi Guennewig, Bryan Routledge, and Ariel Zetlin-Jones. Achieving consensus on blockchains. Technical report, University of Bonn and University of Mannheim, Germany, 2025.

Duncan K Foley. Economic equilibrium with costly marketing. In *General Equilibrium Models of Monetary Economies*, pages 183–199. Elsevier, 1989.

Rodney J Garratt and Maarten RC van Oordt. Why fixed costs matter for proof-of-work–based cryptocurrencies. *Management Science*, 69(11):6482–6507, 2023.

Frank H Hahn. Equilibrium with transaction costs. In *General Equilibrium Models of Monetary Economies*, pages 201–225. Elsevier, 1989.

Franz J Hinzen, Kose John, and Fahad Saleh. Bitcoin's fatal flaw: The limited adoption problem. *NYU Stern School of Business*, 2019.

Peter Howitt and Robert Clower. The emergence of economic organization. *Journal of Economic Behavior & Organization*, 41(1):55–84, 2000.

William Stanley Jevons. *Money and the Mechanism of Exchange*. D. Appleton, London., 1875.

Kose John, Maureen O'Hara, and Fahad Saleh. Bitcoin and beyond. *Annual Review of Financial Economics*, 14(1):95–115, 2022.

Robert A Jones. The origin and development of media of exchange. *Journal of Political Economy*, 84(4, Part 1):757–775, 1976.

Nobuhiro Kiyotaki and Randall Wright. On money as a medium of exchange. *Journal of political Economy*, 97(4):927–954, 1989.

Nobuhiro Kiyotaki and Randall Wright. A search-theoretic approach to monetary economics. *The American Economic Review*, pages 63–77, 1993.

Joseph M Ostroy. The informational efficiency of monetary exchange. In *General Equilibrium Models of Monetary Economies*, pages 113–128. Elsevier, 1989.

Fahad Saleh. Blockchain without waste: Proof-of-stake. *The Review of financial studies*, 34(3):1156–1190, 2021.

Martin Shubik. *The theory of money and financial institutions*, volume 1. Mit Press, 1999.

Ross M Starr. *Why is there money?: Walrasian general equilibrium foundations of monetary theory*. Edward Elgar Publishing, 2012.

Toni Wahrstätter. Is it worth using mev-boost? *Ethereum Research, https://ethresear.ch/t/is-it-worth-using-mev-boost/19753*, 2024.

Neil Wallace. The overlapping-generations model of fiat money. *Staff Report*, 1978.

Mengqian Zhang, Yuhao Li, Xinyuan Sun, Elynn Chen, and Xi Chen. Maximal extractable value in batch auctions. In *Proceedings of the 26th ACM Conference on Economics and Computation*, pages 510–510, 2025.