IOTA 2.0 Incentives and Tokenomics Whitepaper

November 2023

IOTA 2.0 Incentives and Tokenomics Whitepaper

Olivia Saa, Andrew Cullen, Luigi Vigneri

November 2023
Version 1.0

Contents

Abstract

This whitepaper introduces IOTA's groundbreaking approach to tokenomics and incentives, challenging prevalent models in the crypto space. Traditional cryptocurrencies often rely on token-based incentives, resulting in a degradation of civic duty and a skewed distribution of wealth. IOTA 2.0 diverges from this trend, offering access to the network as a reward for maintaining the network, thus creating an inclusive and accessible cryptocurrency ecosystem and making digital autonomy a reality for a broader user base.

IOTA 2.0's leaderless consensus eliminates fees for token holders, who instead burn Mana, a resource generated by their tokens, to produce their own blocks. By rewarding participation in the network with Mana and not the base token, IOTA 2.0 tokenomics prevents value extraction and exploitation by profit-motivated validators. Our approach eliminates inflation entirely, ensuring a fixed token supply and preventing wealth concentration. Tying rewards directly to the system's utility also encourages sustained, long-term engagement from early adopters, and accommodates users restricted from receiving cryptocurrency rewards.

1

Introduction

In 1993, economists from the Institute for Empirical Economic Research at the University of Zurich asked the residents of Wolfenschiessen, a Swiss village, whether they would vote to accept a nuclear waste repository in their community. Although the facility was widely viewed as undesirable, 51 percent of residents said they would accept it. Apparently, their sense of civic duty outweighed their concern about the potential drawbacks. The economists then presented the residents with the option of receiving an annual monetary payment as compensation. Surprisingly, the financial incentive cut the acceptance rate in half, from 51 to 25 percent [1].

The example of this Swiss village shows that cash incentives can erode an existing sense of civic duty, a phenomenon that can be observed in many market-leading traditional cryptocurrencies that rely on generous inflation and fees to attract participation. In many cases, this has led to greed obscuring the technology's moral backbone, which we believe should be to enable digital autonomy for everyone. End-users of many such cryptocurrencies are systematically exploited and excluded, leaving them disillusioned with a technology that they once believed to be a force of good. Moreover, the prevailing inflationary models employed in many cryptocurrencies further exacerbate wealth imbalance, consolidating power in the hands of a few while leaving many participants disempowered. Centralization of wealth and power not only goes against the core ideology of any cryptocurrency but undermines the security and basic operation of the system, leaving it open to corruption and tampering.

In IOTA 2.0, we take a different approach: rather than cash incentives, we provide access to public goods as a reward for public service. We understand that incentives remain essential in any system, ensuring fairness among diverse actors and covering necessary costs. However, our approach transcends the limitations of cash rewards by decoupling rewards from the IOTA token and rewarding contributions to the system with access to the system. This access attracts key players who want to utilize the IOTA network to its full potential, rather than those who seek only to exploit the system for profit and nothing else. Alongside our unique DAG-based consensus, IOTA 2.0's access-based incentives model enables an unprecedented set of properties. These properties illustrate that rewarding participation with access aligns incentives with IOTA's long-term vision of digital autonomy for everyone. We empower participants to actively engage with the protocol, forging a mutually beneficial relationship between users and the technology.

No fees for token holders. IOTA's access-based incentive scheme means that no fees are required for token holders, which sets IOTA apart from other distributed ledgers and addresses a significant barrier to adoption present in many real-world scenarios. By removing token fees from the equation, IOTA provides a system designed for everyone, regardless of transaction type or geographical boundaries.

No inflation of the IOTA token. In addition to eliminating fees for token holders, rewarding with access also removes the need for IOTA token inflation to support incentives. Instead, the IOTA token supply remains fixed, preventing dilution of token holders' funds due to inflationary rewards, the likes of which can be seen in many large-scale cryptocurrencies.

Fairer wealth redistribution. By decoupling rewards from the base token, we prevent the typical redistribution of wealth that can be found in systems with token fees and inflation. In these systems, wealth flows from end-users to validators, making the rich richer and the poor poorer.

Empowered end users. Access-based rewards empower users to engage actively with the protocol from the outset by enabling any token holder to issue their own blocks and make use of the system. Because IOTA 2.0 has no token fees and no inflation, this participation is possible without leaching value from their IOTA token holdings.

Leaderless access. IOTA 2.0 consensus is leaderless, meaning that access to write to the ledger is never controlled by a single entity. Instead of relying on centralized entities or validators, IOTA places the power directly in the hands of its users, allowing all token-holders to issue their own blocks and play an active role in consensus. The leaderless DAG-based consensus of IOTA 2.0 increases censorship resistance and minimizes value extraction by powerful validators.

Long-term commitment. Our reward scheme encourages meaningful long-term commitment from early adopters of IOTA 2.0 technology. This is because rewards are not simply tokens that can be cashed out to make an immediate profit, but are tied directly to the utility of the system: access becomes more valuable as the technology and ecosystem matures, increasing utility and demand for use of the system.

Legal/regulatory flexibility. As a final note, due to rewards coming in the form of access, IOTA 2.0 allows active participation from users who are restricted from receiving cryptocurrency rewards for legal or regulatory reasons.

In summary, IOTA 2.0 represents a transformative approach to tokenomics and incentives, challenging traditional models prevalent in the crypto space. In the remainder of this paper, we will explain how our new tokenomics scheme works, and why we believe that it will pave the way to digital autonomy for everyone.

2

The IOTA economy

Only cryptocurrencies that have well-designed tokenomics and incentive schemes will survive in the long term, as it is economic factors that ensure their ultimate success. Poorly thought-out tokenomics can leave projects vulnerable to short-term changes in the economic environment, resulting in selling spirals and hyperinflationary scenarios. (This is not exclusive to cryptocurrencies; it is also a well-known phenomenon in traditional monetary dynamics [2].) Ultimately, the results of poorly thought-out tokenomics could lead to the demise of a DLT, given the correlation between asset price and the security of the system in Proof of Stake (PoS)-based systems. Equally, projects without appropriate incentive schemes may never attract the engagement required to build a robust decentralized infrastructure, which can doom them to failure from the outset.

As discussed above, rewarding validators through token fees and inflation, whilst punishing users who fund this, is the *de facto* tokenomics scheme adopted by most significant players in the cryptocurrency market. This leads to corrupt and centralized systems which leave their users vulnerable and disempowered. IOTA's tokenomics, however, is fundamentally different.

Before delving deeper into IOTA's tokenomics, we must first emphasize some of our protocol characteristics and differences from most other DLT projects. First consider a typical *leader-based* blockchain with which many readers will be familiar, as illustrated in Figure 2.1. In a leader-based DLT, a block is usually a set of data constructed by a block issuer (a miner in Proof of Work or validator in Proof of Stake), including a set of transactions selected from a mempool. The end-users send their transactions to the mempool, and the block issuer decides which transactions in the mempool will be added to the block. This selection is done at the block issuer's will, meaning that the block issuer can select the transactions that pay higher fees, censor transactions as they please, or even order transactions inside a block in a way that maximizes their financial gain [3].

Figure 2.2 represents the *leaderless* approach adopted in IOTA 2.0. In our setting, end-users issue their own blocks to a directed acyclic graph (DAG) structure known as the Tangle, but validators still have the important roles of deciding acceptance and finalizing transactions. Since the Tangle is a well-structured database (in contrast to a regular mempool, which is a set of transactions with no well-defined order), IOTA 2.0 validators cannot easily censor or order individual transactions.

In this sense, the IOTA protocol is leaderless because token holders can issue blocks containing their own transactions. Thus, users are not required to pay validator fees to gain network access.

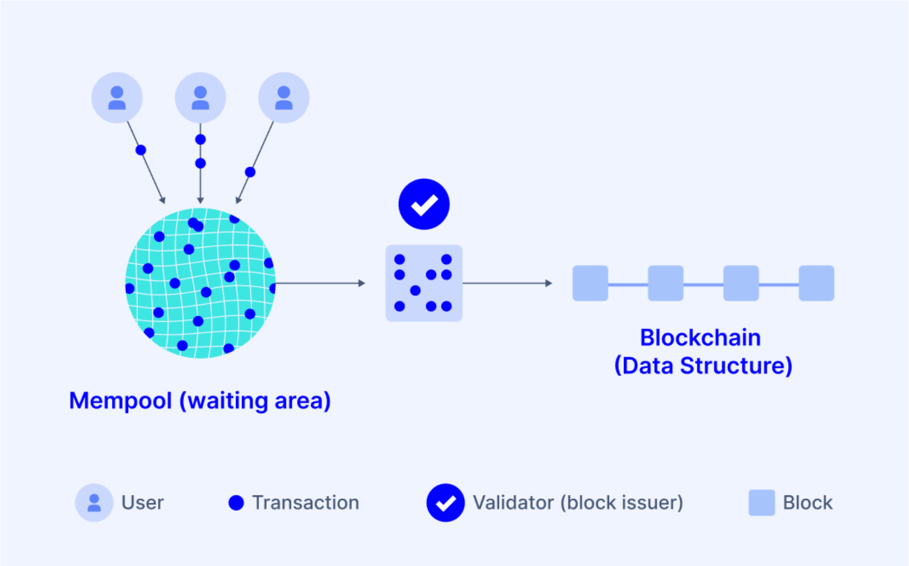

Figure 2.1: Access control in traditional PoS blockchains

However, to protect IOTA from spam and Sybil attacks, we introduce a congestion control mechanism to the block creation process. This congestion control mechanism is based on ownership and consumption of a reward resource called *Mana* which is generated by holding IOTA tokens and contributing to consensus. It is utilized for block issuance, access to network throughput, and protection against Sybil attacks. Thus, a user's right to access the ledger is defined by the amount of Mana they hold, so when we reward users with Mana for certain actions, we effectively reward them with access to the ledger and its functionalities. In this way, IOTA tokenomics is optimized for users who seek to make practical use of the system. Since block issuance is regulated by the consumption of Mana and blocks are the containers of all interactions with the ledger, Mana, in practice, is used to power certain actions in the IOTA protocol, from transferring IOTA tokens to executing smart contracts or minting NFTs.

Initially, Mana generated by token holders and consensus contributors will stockpile to a certain amount. But, as Mana gains value, these stockpiles will be sold and used. Additionally, the rewards and Mana dynamics are designed so that Mana can no longer be stockpiled after the initial bootstrapping phase.

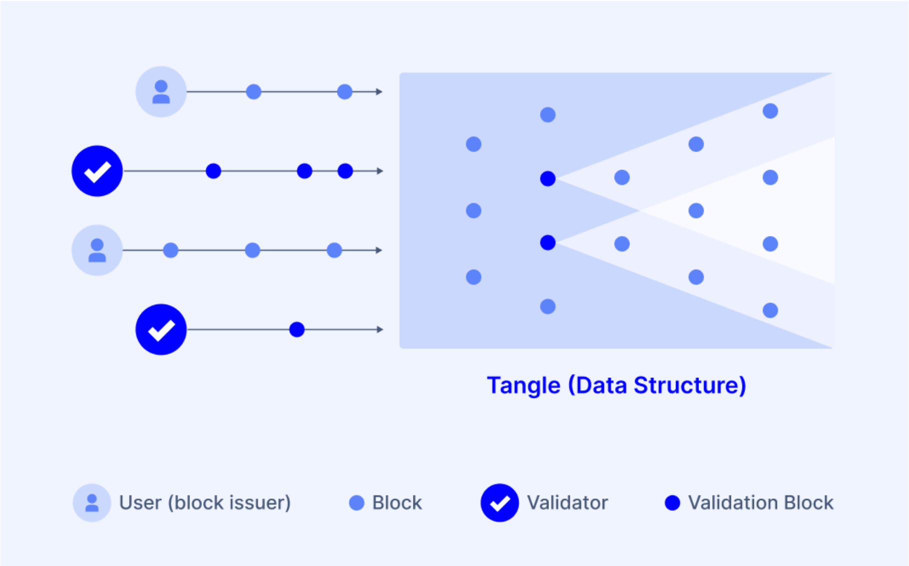

Figure 2.2: Access control in leaderless DAG-based ledgers

After this phase, when the network is considered mature, Mana will be burned quickly after its generation. Moreover, an increasingly higher portion of the token supply will eventually be owned by users wishing to generate Mana that they can utilize. At this stage, users will have a guaranteed share of the network throughput by owning a certain fraction of the token supply.

2.1 A sustainable economic model

A simplified model for the IOTA economy is illustrated in Figure 2.3. The ultimate goal of the IOTA economy structure outlined here is to provide a ledger with high security, utility, and token value through a sustainable incentives scheme.

Utility is multifaceted, and its meaning will evolve as the IOTA ecosystem matures and new applications for the IOTA ledger emerge. Still, block creation will always be the fundamental step required to interact with the ledger and benefit from this utility. Thus, utility drives demand for Mana.

This is where our incentive scheme comes into play: since *Mana can be obtained by holding tokens, delegating, or validating*, any actor wishing to make practical use of the IOTA ledger is incentivized to hold tokens and participate in staking in this way. These activities must be incentivized since they each play a vital role in securing the IOTA network and sustaining the value of the IOTA token. Specifically, holding IOTA tokens, delegating, and validating all require IOTA tokens to be purchased, which sustains the token's value. This directly strengthens the security of our consensus because any actor's voting

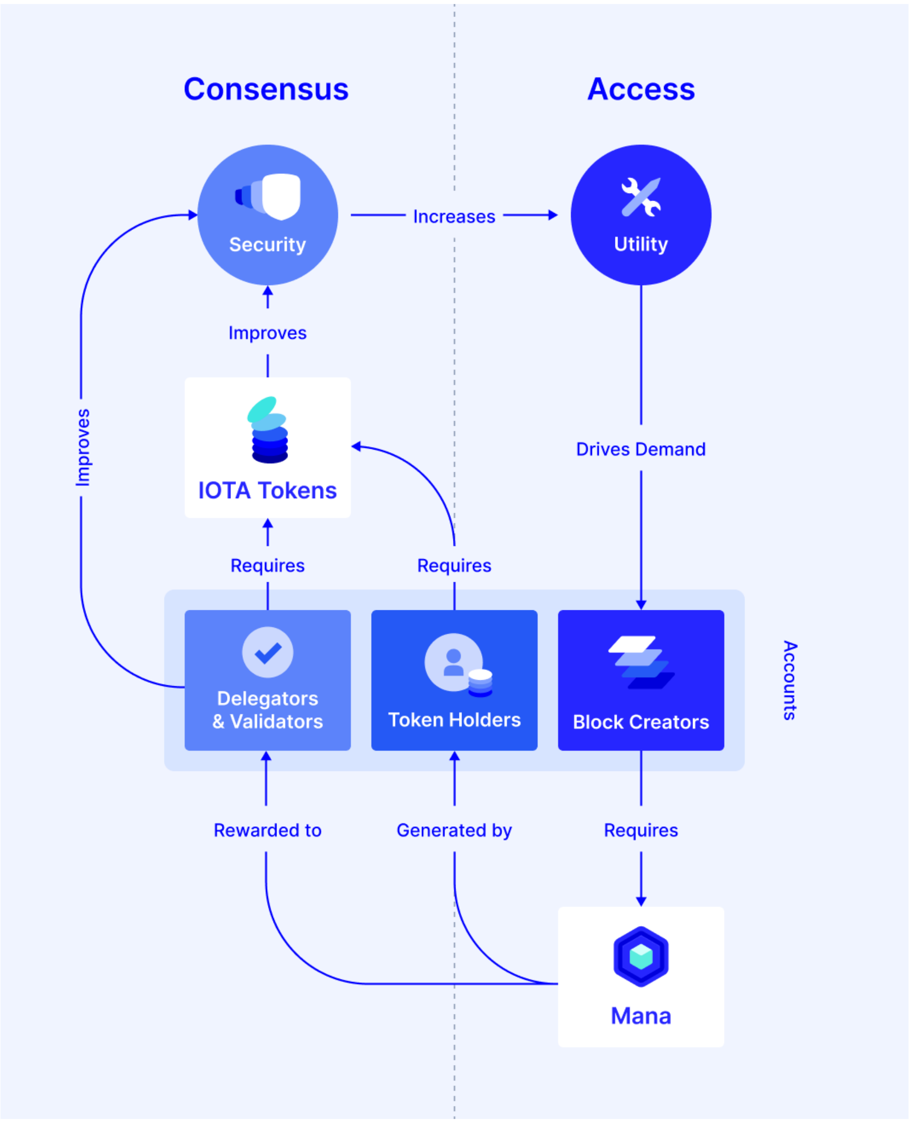

Figure 2.3: A model for the IOTA economy

power depends on their stake, making any malicious activity unsustainably expensive. Additionally, delegating and validating play a direct role in the consensus process; hence, these actors support the network's security. As a final note, increased security improves the utility of the ledger; the most practical utility relies on a secure ledger as its foundation.

Thus, an organic and sustainable dependency is created between these economic actors. By rewarding actors with Mana (instead of the IOTA base token), we attract and reward actors who wish to use our technology rather than draining the value of their tokens through inflationary rewards and fees. Note that a positive feature is that IOTA does not incentivize large sell-offs of Mana due to instabilities in the early stages of the IOTA economy's growth; instead, early adopters of IOTA are rewarded for not leaving the system since Mana will be more valuable when the IOTA economy is well established and when there is a high demand for the multitude of applications built on IOTA.

2.2 Economic actors and their roles

To participate autonomously in the IOTA economy, a user must register an *account* on the ledger, which can then be used for any economic activity within the protocol. The account is a protocol element stored on the decentralized ledger with no association with the IOTA Foundation or other parties. A wide range of objects and metadata can be linked to an account such that it may serve as a form of persistent decentralized identity. However, in this white paper, we are primarily concerned with just two assets held by an account: IOTA tokens and Mana.

These two assets capture all the information required for an account to participate in the IOTA economy. IOTA tokens are the ledger's base asset, a non-inflationary value store. IOTA tokens can be staked and delegated to participate in consensus to earn rewards in the form of Mana. We say that a validator *stakes* their tokens, while delegators *delegate* theirs. The technical differences between these two processes will be introduced in section 4. Mana, on the other hand, is earned by holding IOTA tokens and contributing to consensus. It is spent to create blocks to modify the ledger, among other interactions with the system.

The IOTA economy can be further understood by considering the different classes of economic actors that account holders can embody within the IOTA economy. The roles of these economic actors can be summarized as follows:

- *Block creators* are the consumers of the IOTA economy. The creation of blocks is the fundamental action required to make practical use of the ledger, whether transferring funds between accounts, minting an NFT, or interacting with a smart contract. Block creation requires an expenditure of Mana. Note that users of IOTA are empowered to issue their own blocks using their Mana, in stark contrast to any other blockchain in which users are at the mercy of service providers to have their transactions included in a block.

- *Token holders*, as shown in Figure 2.3, lie at the boundary of access and consensus; they rely on the ledger as a store of value. As a delegated PoS-based system, the security of the IOTA ledger is linked to the distribution and scarcity of the IOTA token. As such, holding tokens constitutes a role in consensus and is rewarded with a passive generation of Mana.

- *Validators* play a unique role in IOTA's consensus mechanism by executing specific tasks that allow the entire network to agree on the ledger state. They contribute directly to the ledger's security and are rewarded for this service with Mana (in addition to the Mana passively generated by holding tokens). Our consensus scheme chooses a subset of validators to provide the validation service within a time period referred to as an *epoch*; we call this subset the *epoch committee*, while we define as a validator anyone who registers to validate, regardless of whether they are selected for the epoch committee in a given epoch (more details on the staking and validation process can be found in Section 4).

- *Delegators* also contribute to consensus, albeit less directly than validators; token holders who are not interested in becoming a validator can still contribute to consensus by delegating their voting power to a validator of their choice (see Section 4 for more details). Delegators are also rewarded with Mana for their contribution to active and well-performing validators in the form of a fraction of the total amount rewarded to said validator (meaning that delegation to an offline or poorly performing validator might not be rewarded at all).

All token holders, regardless of whether they participate in staking or delegation, are rewarded with Mana. An extra reward is given to delegators and validators for their contribution to consensus. The reward given to validators is more significant than that given to delegators. This reward difference is fair since an active node is required to perform validation services, while delegators do not need to maintain this node. Since delegation is an inexpensive activity, all tokens would be staked or delegated in a perfect scenario. However, delegation is not enforced to respect the token holders' particularities, and a token holder will always have the option to abstain from delegation.

2.3 Wealth redistribution

Wealth flow among system participants is a crucial factor contributing to the sustainability of our tokenomics scheme. IOTA's unique approach differentiates it from other cryptocurrencies in this respect, thanks to a combination of no fees being paid to validators from the base IOTA token and rewards being distributed as an access-related secondary asset.

To illustrate the contrast with most traditional PoS-based DLTs, let's examine a typical wealth flow. In this context, we focus on internal wealth dynamics rather than external flows. In Figure 2.4, we present a simplified model of this internal wealth flow. This depiction highlights two primary drivers: user fees paid to validators and rewards distributed in base tokens. Over time, these factors naturally increase the share of base tokens held by validators, shifting the balance away from users.

We want to emphasize that validators certainly deserve rewards for their valuable contributions to the system and the real-world costs they incur. However, this wealth flow tends to reduce the number of tokens held by users without a clear lower limit.

The absence of a lower limit can be a subject of debate. Some systems have reward mechanisms that decrease over time until they reach zero (as seen in Bitcoin, also known as the *halving* of block rewards). However, the safety and practicality of such designs need to be thoroughly verified, (as discussed in [4]). When the transaction fee income surpasses rewards from newly minted tokens, block issuers might be tempted to engage in questionable practices, as forking the chain when a block collecting a large amount of fees is mined.

Thus, a provenly secure incentive scheme should ensure a continuous flow of wealth to block issuers, theoretically allowing their share of tokens to eventually dominate the total supply.

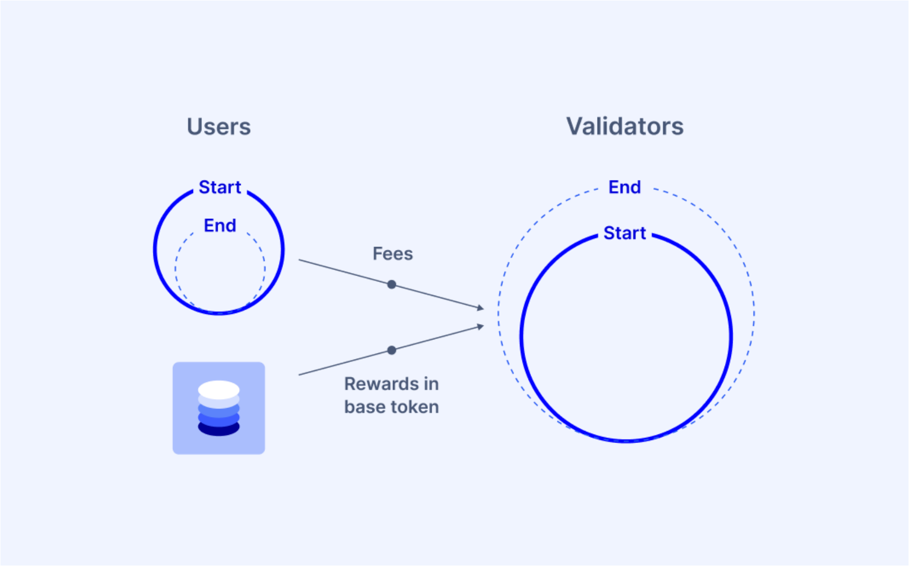

Figure 2.4: Wealth flow in traditional PoS-based systems

On the contrary, IOTA's wealth flow operates on a fundamentally distinct principle. As depicted in Figure 2.5, we can see a simplified internal wealth flow model unique to IOTA.

It's important to note that these flows don't involve a transfer of assets from one type of actor to another. Instead, two distinct categories of actors receive and burn Mana, and their holdings of IOTA tokens remain unaffected by how they use the system. However, it's crucial to recognize that this alone doesn't automatically imply that wealth redistribution is inherently beneficial.

Validators do indeed require some form of gain to incentivize their participation. However, this profit doesn't necessarily have to come at the expense of token holders. We envision two potential scenarios for someone to become a validator:

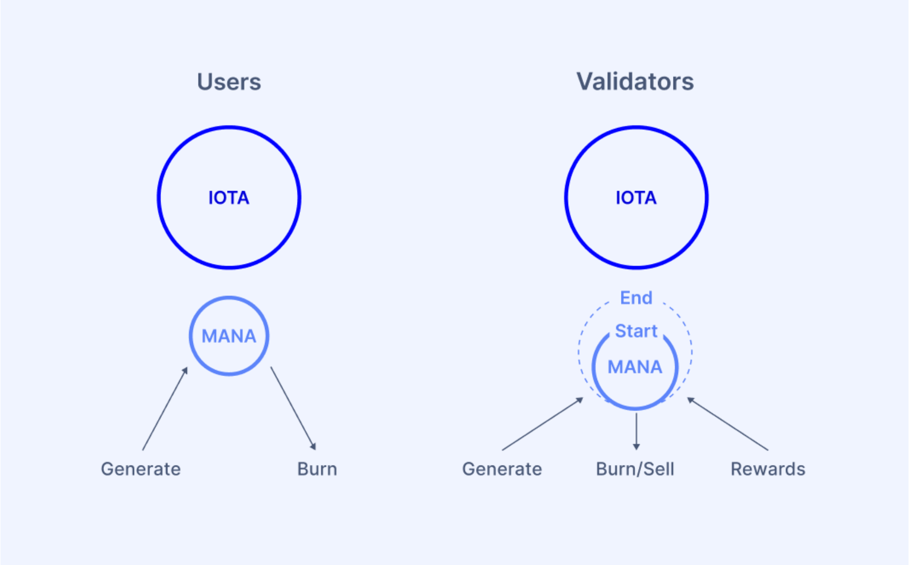

Figure 2.5: Wealth flow in the IOTA economy

- *IOTA Application Host*: In this scenario, the validator is a provider looking to integrate the network into their applications seamlessly without encountering any friction from transaction fees. It's clear that, in this case, no wealth is extracted from token holders and transferred to the validators.

- *Mana Seller*: Alternatively, a validator can be an entity interested in selling Mana for a profit. But if token holders already possess a guaranteed amount of Mana for their intended use, who will buy Mana from the validator? This is where our approach becomes interesting: validators can profit by selling Mana to individuals that in some sense are external to the system. This can include users who periodically require additional network access beyond what their token holdings allow or even users who don't hold tokens at all. This way, token holders' wealth isn't diminished by design, and validators can still reap their rewards without negatively impacting the token holders.

3

Mana

As illustrated in Figure 2.3, Mana is the resource required to access the IOTA ledger and update its state by creating blocks. As a spendable resource tracked in the ledger state, Mana can also be used to power smart contracts, DeFi applications, and other services[1].

The mechanics of the Mana system can be better understood by thinking of a user's IOTA token holdings as a pear orchard and Mana as pears (see Figure 3.1). Just as pears are a tangible result of cultivating an orchard, Mana is generated proportionally to the base token holdings of each user in the IOTA network. Let's explore this analogy in detail:

- *IOTA Tokens as an orchard:* In this analogy, IOTA tokens are represented by the orchard in which pears are grown. IOTA tokens are like the fertile ground on which Mana can be cultivated. An orchard is a valuable asset to own because of its potential to grow pears and the fact that suitable land of this kind is finite. Similarly, the value of IOTA tokens is derived from the fact that its supply is finite and it has the potential to produce Mana.

- *Mana as pears*: pears represent Mana in this analogy. Just as pears are the tangible rewards of tending to your pear orchard, Mana is the tangible reward of holding IOTA tokens and actively participating in the IOTA network. Pears are the edible outputs of an orchard. Similarly, Mana is the fruit of IOTA tokens and it can be consumed to make practical use of the IOTA ledger.

- *Proportional growth*: the amount of pears you can harvest is directly linked to the size of your orchard. Analogously, the amount of Mana you can generate is directly linked to the number of IOTA tokens you possess. The

[1] For instance, our L2 smart contracts use as gas either a native asset from L1 or an L2 token. For each IOTA Smart Contract (ISC) chain, deciding on the asset to be used is a configurable choice, thus making it possible to have ISC chains in which Mana is used to pay for execution.

more land you cultivate, the more pears you can harvest, and the more IOTA tokens you hold, the more Mana you can generate.

- *Active participation*: just as an orchard owner can improve their pear yield by caring for their land, IOTA users can generate more Mana through active engagement with the network. This engagement might involve validating transactions or contributing to the network's security by delegating voting power, both activities that benefit the IOTA ecosystem.
- *Generation of value*: in the analogy, pears have value since the farmer can either eat those pears or sell them. Similarly, Mana will have value in the IOTA network and serves as an incentive to users' contribution and active participation.

In summary, the analogy of Mana generation and pear harvesting through land ownership illustrates the concept of proportionate resource generation in the IOTA protocol. Just as owning more land allows you to cultivate more pears, holding more IOTA tokens and actively participating in the IOTA network will enable you to generate more Mana.

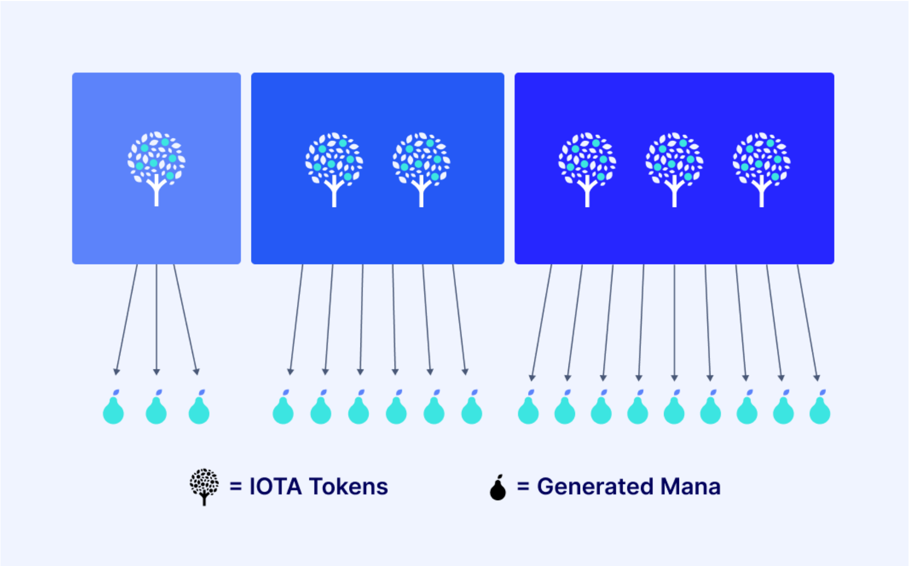

Figure 3.1: Mana as a pear orchard

Demand for Mana is driven by its utility within the IOTA protocol and the promise of future utility that will be created as the IOTA ecosystem develops and new applications emerge. As Mana can only be generated through holding tokens, delegating, or validating, demand for Mana drives demand to participate in these productive activities, further strengthening the IOTA protocol. The role of Mana in the IOTA economy can be further understood by considering how it is obtained and used, as we shall outline in the remainder of this section.

3.1 Obtaining Mana

Mana can be obtained in different ways. It can be generated by holding IOTA tokens, earned for participation in consensus (as either a delegator or a validator), or received from other Mana holders.[2] In any of these cases, Mana is generated at the protocol level; it is not distributed or allocated by the IOTA Foundation or any other entity.

Mana generation by holding IOTA tokens

Mana is constantly generated over time by holding IOTA tokens. The Mana generated by holding tokens is not allocated automatically and periodically by the protocol to avoid excessive and purposeless updates to the ledger state. This Mana will only be accounted for in the ledger when the output holding those IOTA tokens is consumed. You can think of unspent outputs as having *Potential Mana* attached to them because they will generate actual Mana whenever they are spent. Potential Mana is not explicitly stored anywhere but can be easily calculated from the ledger (specifically, from the set of unspent outputs).

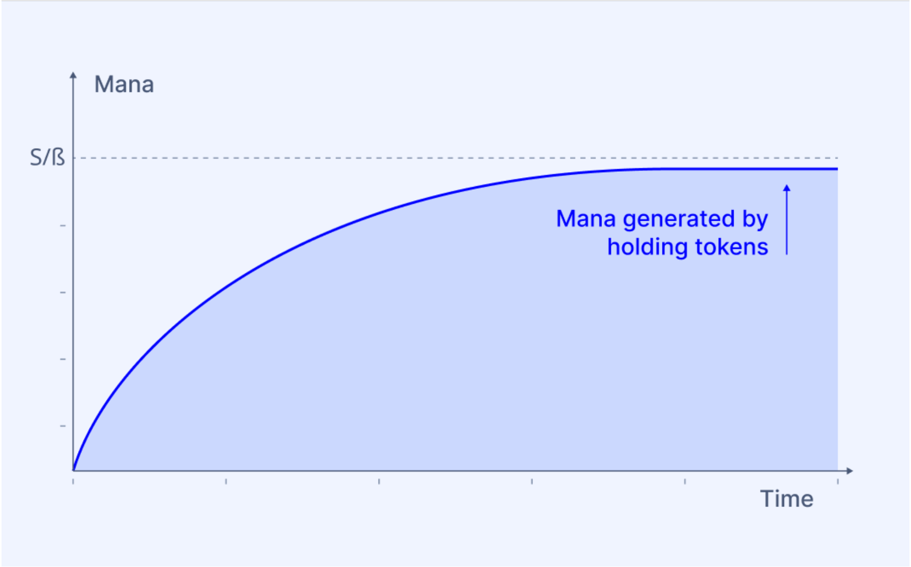

Figure 3.2: Mana generation by holding IOTA tokens

The amount of Mana generated is given as follows:

$$\text{Mana Generated} = \text{IOTA Value of Inputs} \times f(\text{Time Held}), \tag{3.1}$$

[2] A Mana marketplace will not be developed by the IOTA Foundation in the foreseeable future. However, Mana can be sold and purchased in third-party applications or exchanges like any other asset in our ledger.

where Time Held represents the difference in timestamps between the transaction that generated the inputs and the transaction that consumes it. The function f is designed to prevent an excessive accumulation over time, making Mana stored for long periods less valuable[3] (see Figure 3.2). We discuss more details of how Mana is handled and the exact formulas used in Appendix A.

The concept of Mana being generated by holding IOTA tokens can also be extended to other on-chain assets (such as NFTs) if these assets also include deposited IOTA tokens. Holding these assets will generate Mana for the user at a rate proportional to the user's IOTA tokens deposited according to (3.1).

To understand how Mana generation happens in more detail, consider Figure 3.3, which illustrates a generic transaction payload. Note that the figure is merely an illustration and does not reflect the exact content of a transaction as implemented in IOTA. Each consumed input yields Mana based on its IOTA value and the time it has been held. In the example above, the total amount of IOTA tokens in the inputs is 6, the same amount of IOTA tokens in the outputs. Analogously, the amount of Mana stored in the inputs is 3, and the potential Mana attached to the inputs (not shown in the image since it's not explicitly part of the transaction but of the inputs themselves) is 2. This results in a total of 5 Mana to be sent to other accounts or to be stored in new outputs.

Mana Generation by Delegating and Validating

Mana is also rewarded for participating in consensus via delegation and validation. While these rewards are not directly derived from IOTA tokens, the reward amounts are part of the information locally tracked by nodes so that the corresponding Mana can be claimed individually by all reward recipients. All matters related to staking and rewards will be discussed further in Section 4.

Purchasing Mana

As explained above, Mana is passively generated by holding IOTA tokens and contributing to consensus. Nevertheless, any user who requires additional Mana can purchase it, and users who do not wish to hold IOTA tokens for any reason can purchase Mana directly from someone else instead. To do this, the Mana seller must communicate with a Mana buyer to construct a transaction in which the buyer transfers tokens to the seller in exchange for stored Mana transferred by the seller. Alternatively, buyers and sellers can use third-party applications or exchanges to transact. We emphasize that the IOTA Foundation will not provide the structure for a Mana market in the foreseeable future; however, Mana is a resource like any other that lives in the ledger, thus not presenting any conceptual impossibility of being traded.

[3] In practice, the function f corresponds to periodical decay factors applied to a constant generation rate over time

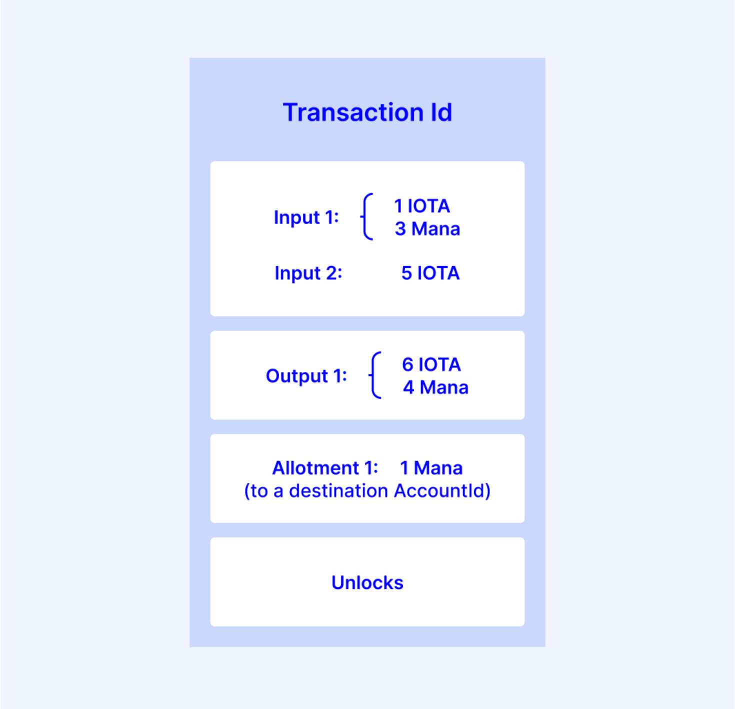

Figure 3.3: A generic transaction payload

3.2 Spending Mana

Contrary to how ledger access is managed in blockchains, block creation in IOTA is not reserved for a select group of validators or miners. Instead, any IOTA account holder can use their Mana to create blocks autonomously. Furthermore, users do not pay this Mana to other powerful actors to create blocks on their behalf as in fee-based systems; instead, they use their Mana in the protocol to take advantage of useful functionalities without intermediaries.

Specifically, Mana is burned each time a block is created, meaning it is subtracted from the Mana balance of the block creator. The network's congestion levels determine the amount required to be burned, and nodes dynamically regulate this amount as the blocks are received. A detailed description of this congestion control is out of the scope of this document, but the interested reader can find more information about this topic in [5]. The main difference between a Mana-based system and a traditional fee-based one is that users can interact with the ledger and issue blocks without spending IOTA tokens. Furthermore, their Mana generation is proportional to their IOTA token holdings, meaning that (on average) a user can reserve a share of the network throughput according to their fraction of the IOTA token supply [4]. Thus, a single IOTA token

[4] During periods of low network activity, a user might obtain a larger share of throughput

investment guarantees a specific throughput for as long as the user wants.

Figure 3.4 illustrates a simplified outline of the block structure, showing generation of Mana from a generic payload (for example, a simple IOTA token transfer) and the burning of Mana due to the creation of the block. This dramatically simplifies block creation from the user's perspective, providing infrequent ledger users with a simple means of transferring crypto assets without expressly having to acquire Mana and without paying fees to other network users to provide block issuance on their behalf.

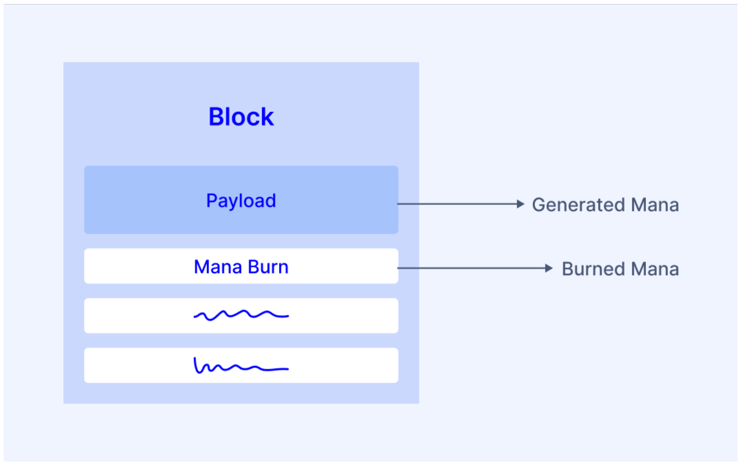

Figure 3.4: A simplified illustration of a block containing a generic payload to add to the ledger. Mana generated by the payload can be credited to specified accounts, and the Mana burned is debited from the block creator's account.

3.3 Particularities of Mana dynamics

Three primary factors impact the price and supply dynamics of Mana:

- The rate at which Mana is generated from holding tokens and participating in staking.
- The rate at which the global Mana supply is decayed.
- The rate at which Mana is burned to create blocks to make practical use of the distributed ledger.

Firstly, consider the generation rate of Mana from held tokens and from staking and delegating. Protocol parameters determine each of these generation rates, and increasing those rates will increase the Mana supply, which can lead to a decrease in the Mana price. However, with a significant abundance of Mana

than the guaranteed level.

in the system, it becomes cheaper for actors to congest the network, which would increase the Mana burned per block. Thus, we consider the Mana generation rate a scaling parameter that dictates the Mana supply but does not affect the system's usability.

The second parameter that controls the dynamics of Mana is the global decay rate, which is also set by the protocol. Decaying Mana encourages spending and discourages hoarding of Mana, which is similar to the effect of increasing Mana generation and reward rates over time: both approaches decrease the value of hoarded Mana over time. The larger the decay rate, the smaller the total supply of Mana will be. However, unlike in the case of generation rate, decay impacts the qualitative behavior of the system. A decay rate set too low might not incentivize Mana spending at all, whereas a decay rate set too high would make the users' Mana too dependent on their amount of IOTA tokens and almost independent of the time for which those tokens are held. Thus, we consider the decay rate as an important parameter, which profoundly impacts user behavior. Interested readers can find technical details of how Mana decay is implemented in a non-gameable way in Appendix A.

The final and least controllable aspect of Mana dynamics is the burn rate of Mana due to the creation of blocks. The *Mana cost* required to issue a block is an adaptive value that increases during high congestion periods and decreases during low congestion. This reference Mana cost has some impact on the total burn rate of Mana. Still, the burn rate primarily depends on the demand for block creation and ledger resources, as well as on the myriad factors that affect the perceived value of Mana at any given moment. As such, the burn rate is difficult to predict or control. However, as demand increases, the reference Mana cost increases, and so does the burn rate until we reach an equilibrium point at which the burn rate equals the generation rate.

It is out of the scope of this document to provide a deep analysis of the monetary equilibrium points of the system, but simulations point to an equilibrium in the Mana generation and burn rates in all studied scenarios. Naturally, this equilibrium point changes quantitatively depending on how the protocol parameters are set, which means that the Mana spent per block tends to be dictated by the protocol parameters. However, the Mana spent per block is not a direct measure of the price of a block, since the Mana supply can also substantially change with protocol parameters; what would actually be a good measure of it is how much (in fiat value) one is willing to pay for the Mana cost of a block.

4

Staking

Our incentive model includes a staking mechanism (which requires the validator's tokens to be locked) and a delegation mechanism (which leaves the delegator's tokens unlocked and free to be spent), native to the protocol. When there is no need for differentiation between them, we call both mechanisms *staking* in a more general manner. Some points make this an exciting and desirable mechanism to be included in an incentive scheme of a DLT, which are summarized below:

1. **Permissionless participation.** Our mechanism is designed in such a way that any party is free to participate in staking (with no minimum stake needed) and is rewarded accordingly (see next point).

2. **Fairness in rewards.** Our reward function is designed so that all staking participants are compensated fairly according to their contribution to the successful functioning of the system (see Appendix B).

3. **Commitment to the security of the network.** Our validator setup requires staked tokens to be locked up (and incentivizes this through increased rewards compared to delegating), which makes the protocol less susceptible to exploits since it tends to stabilize the voting power distribution and obliges validators to bear a financial risk when engaging in misbehavior[1].

4. **Liquid delegation.** Our delegation system does not require the locking of tokens. This allows token holders who aren't willing to perform the computational job of a validator to also commit to the functionality and security of the network (albeit to a lesser extent) while still having their tokens unlocked (see Section 4.3).

[1] One of our design principles is that only objectively measurable misbehavior (i.e., behavior that we can be sure is deliberately malicious) can be punished with slashing. In our case, we have introduced the possibility of slashing Mana to the protocol.

5. **Shared security between layers.** Our staking system allows for funds in Layer 2 to participate in the security of Layer 1, by delegating the voting power that would otherwise be locked in IOTA Smart Contract chains. By enabling delegation of funds participating in Layer 2, the success of our Layer 2 chains (and, consequently, an increase in their security) does not undermine the security of Layer 1.

6. **Decentralization of stake.** Our rewards formula disincentivizes the concentration of delegated stake on a single validator by reducing rewards for delegators in such cases. This encourages delegators to choose validators with less delegated stake, leading to a stable ratio of delegated and validator stake in the system.

7. **Low barriers to adoption.** Staking is a widely known and relatively simple mechanism. For any type of user, adopting a network with a known functionality is easier than learning about a new subject before deciding to participate.

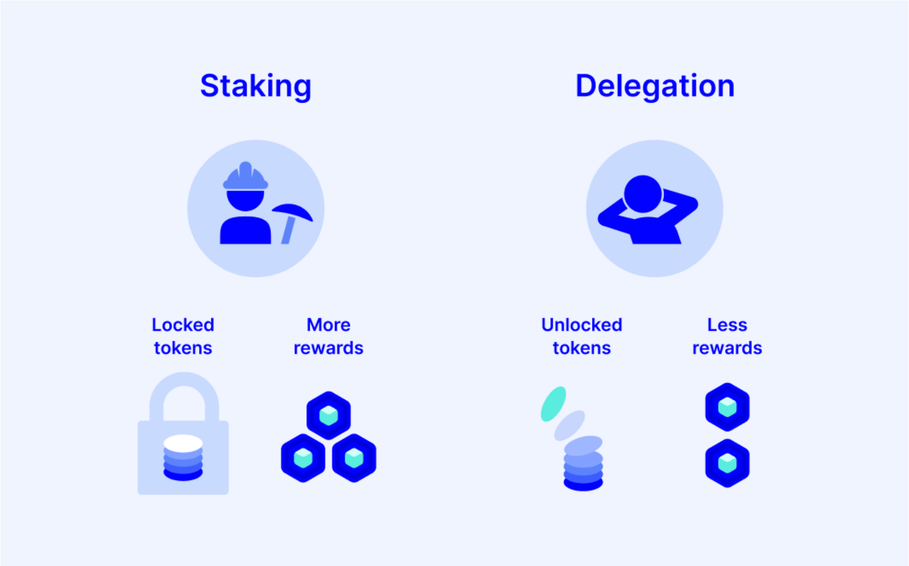

Figure 4.1: Staking vs delegation

4.1 Validation blocks

The IOTA consensus protocol includes a committee selection algorithm executed at each epoch to select a subset from the registered validators (the registration procedure is detailed in section 4.2). Each registered validator has an associated weight, which is a function of their staked tokens and the number of tokens delegated to them. This weight is the key parameter used for selection.

Members of the selected epoch committee are expected to issue special blocks (called validation blocks) in a timely manner to keep the consensus protocol running without issues. Validation blocks differ from other blocks, as they should be as small as possible, containing only strictly necessary data (as parents or signatures). Validation blocks issued by nodes not selected for the committee are discarded and not gossiped.

Additionally, the protocol's congestion control mechanism, which regulates the allocation of block creation among accounts, ensures that members of the epoch committee receive a guaranteed throughput allowance without burning any Mana. This ensures that all validation blocks are disseminated through the network even during heavy congestion and minimizes the cost to validators for providing this essential service. The extra throughput provided for validation blocks issued by the committee is designed to be small compared to the throughput allowed for basic blocks.

Issuing validation blocks (and, consequently, the validation of the DAG) is considered an essential part of the validator's role, as issuing these blocks corresponds to casting their votes on the ledger's state. Thus, the reward scheme is tied to the issuance of these blocks, i.e., a validator selected for the committee must correctly issue validation blocks at epoch n to receive the rewards relative to the said epoch. However, the specification of the number of validation blocks that must be issued per epoch is out of the scope of this document.

4.2 Staking for validation

Actors who wish to directly participate in the validation mechanism must register themselves as validators by adding a *staking feature*[2] to their account. The registration is only considered successful once the blocks that add that feature are confirmed. This way, the actor signals their interest in participating in future validator committee selections. The staking feature locks a user-chosen amount of IOTA tokens until a user-chosen end epoch. Alternatively, the validator can choose not to specify an end epoch, and they will be considered locked until the validator signals that wishes to unstake. After this signaling, the locked funds must be unbonded before they are unlocked by the protocol (i.e., validators must wait until the end of the epoch after the unstaking request to have their funds actually unlocked). Once an account has a staking feature whose end epoch is the current one (or has already passed), it is no longer considered for validator selection. Keep in mind that during the unbonding period, the account is still considered staked. After the unbonding period, the funds can be accessed again.

Validators selected for the committee will receive rewards after the end of their mandate depending on how well they performed their duties: the value

[2] A *feature* is a concept of our output designed introduced in TIP-18 ([6]). One can think of a feature simply as some additional information attached to outputs or accounts.

rewarded shall depend on whether the validator issued the expected rate of confirmed validation blocks, whether their votes helped to confirm other blocks, on the size of their stake, and on how much voting power was delegated to them by third parties. Section 4.4 defines the specific rewards distribution to committee members. When the staking feature that was added during registration is removed from the account, the validator can claim their rewards in this same transaction. The rewards will be larger than zero only if the validator was selected to the committee and performed their duties.

Note that registration is necessary but not sufficient to be eligible for the committee. To be considered a candidate for the committee selection, the registered validator must, among other requirements, be active (i.e., issue at least a block) within a certain period. However, the exact eligibility conditions are outside the scope of this document and can be found in [7].

4.3 Delegation

Delegation is implemented with a special output type called *delegation output*. In practice, the delegation process can be seen as a regular token transfer to one of the delegator's own addresses, specifying the account ID of the chosen validator for each output.

Under this delegation mechanism, token holders do not need to lock their tokens to delegate their voting power. Instead, they can use or re-delegate their tokens to other validators of their own volition. If any actor wishes to delegate to a different validator, they can simply re-delegate to the preferred entity (there is no need to un-delegate to delegate again). Note that the effects of re-delegation (e.g., changes in a validator's total stake) will only take place in the following epoch. For that reason, delegators do not lose their rewards for the epoch when moving funds[3].

Any rewards for delegation are claimable by the owner of the delegated IOTA tokens after the end of the epoch, regardless of whether their validator claimed their rewards or not. Section 4.4 defines the specific rewards distribution to delegators. To prevent too many token holders from delegating to the most powerful validators, the reward function is designed to incentivize the distribution of delegated stake among validators. For details of how we incentivize this behavior, see Appendix C.3.

Token holders can select any validator of their choice to delegate their stake. However, we expect them to choose validators based on their own expected rewards and their perception of the validator's reputation. Their perception of the validators' reputation is subjective and, thus, unpredictable. Nevertheless, by the construction of our reward formula, delegation rewards are maximized by choosing validators who consistently perform their validation duties and do not concentrate a proportionally significant delegation stake.

[3] Technically speaking, this requires a transition of the delegation output to a *delayed claiming* state; however, implementation details of outputs is out of the scope of this document. For further information about this subject, see [7].

Finally, our protocol enables what we call *liquid delegation*, since delegated tokens can be moved at any time. Furthermore, it enables the delegation of tokens participating in L2, since any ISC chain can delegate its voting power to a validator and gain rewards accordingly. For more on this subject, see [7].

4.4 Properties of our reward scheme

A well-designed reward distribution function that provides the correct incentive mechanism is essential to achieve the behaviors we desire in the system. Our reward distribution is designed to present the following main features:

1. Non-gameability of the locking rules: there are incentives for the validators to stake their funds instead of delegating them to themselves.
2. Incentives for high-quality validation services: there are incentives for the validators to perform their expected consensus-related activities correctly.
3. No incentives for the concentration of validators' stake.
4. Incentives to spread delegation among all validators.
5. Larger incentives to early contributions to increase the network's security.
6. Minimum guaranteed profitability (or, the 1:2:3 rule).

In the following paragraphs, we elaborate on these features and describe the design of our reward distribution mechanism. The exact reward calculations are in Appendix B.

Properties 1 and 2: Non-gameability of the locking rules and validation services. The influence of validators on consensus depends on their staked value and the outputs delegated to them. We define a validator's *supporters' pool* (or simply *pool*) as the validator plus whoever delegates[4] to them. For example, suppose that for epoch n, the consensus module selects the committee based on the state of the end of epoch $n-1$. Then, we take this information collected at the end of epoch $n-1$ and define the stake of validator i's pool as the sum of the validator's stake plus the sum of all value delegated to them at that point in time:

$$\text{Stake}_i(n) = \text{validator's stake} + \text{delegated value in epoch } (n-1).$$

If rewards were simply linear to the stake, validators could game the locking condition by creating a separate account and delegating the account's tokens to themselves, which would be more attractive since staking is more restrictive

[4] In this document, if a token holder delegates to two different validators, without loss of generality, we consider them as two different delegators, so each delegator will only be part of a single pool.

than just delegating. To prevent this, the reward distribution privileges pools with a larger locked stake by assigning a lower weight to delegated funds[5].

Finally, the rewards are also proportional to a *performance factor* that measures the quality of the service of each validator during each epoch. Thus, validators who do not issue the required validation blocks for epoch n will not receive the rewards corresponding to epoch n, nor will their delegators.

Properties 3 and 4: No incentives for centralizing the funds of validators and delegators. According to the construction of our reward formulas, there is no incentive to centralize validator funds. Proof of this property is provided in Appendix C.2. Technically speaking, this means that two different validators do not get more rewards by combining their stake.

Analogously, the concentration of delegator funds is disincentivized by the construction of our reward formula. As more delegated stake is concentrated on the same validator, the rewards for delegators become smaller and the delegators are incentivized to re-delegate to validators with less delegated stake. Assuming that actors are rational, the system would stabilize around a constant ratio of delegated and validator stakes among pools in the long run. Proof of this property is provided in Appendix C.3.

Property 5: Incentives to early contributions to the network's security. To distribute rewards at the end of each epoch n, we first define the target reward $R(n)$[6]. This target reward takes the form of

$$R(n) = \begin{cases} Ae^{-Bn}, \text{ if } n \leq N \\ C, \text{ if } n > N \end{cases} \tag{4.1}$$

where A, B, and N are positive parameters, and $C = Ae^{-BN}$. The qualitative behavior of this form of reward function is depicted in Figure 4.2

The target rewards function has two qualitatively different regimes: a decaying phase (or *bootstrapping phase*) and a constant regime. The presence of two regimes is justified by the existence of two behaviors we want to incentivize at different times.

The network's utility in its early stages is expected to be lower than in its later stages. Thus, the utility of the network by itself might not be sufficient to incentivize enough users to participate in the validation process. Therefore, in this bootstrapping phase, validators and delegators are more incentivized to contribute by receiving a contribution premium. Since the network utility is lower in this phase, the Mana will not be as valuable as in an established,

[5] To summarize, we introduce a parameter that determines how influential staked tokens are compared to delegated tokens. An analysis of how this parameter affects the rewards can be found in Appendix B.1.

[6] The target reward is the expected reward to be distributed at the end of the epoch. Note that, in the case of a randomized committee selection, the actual reward per epoch can be larger or smaller than $R(n)$, but in expected terms, one can think of $R(n)$ as the total reward distributed in epoch n.

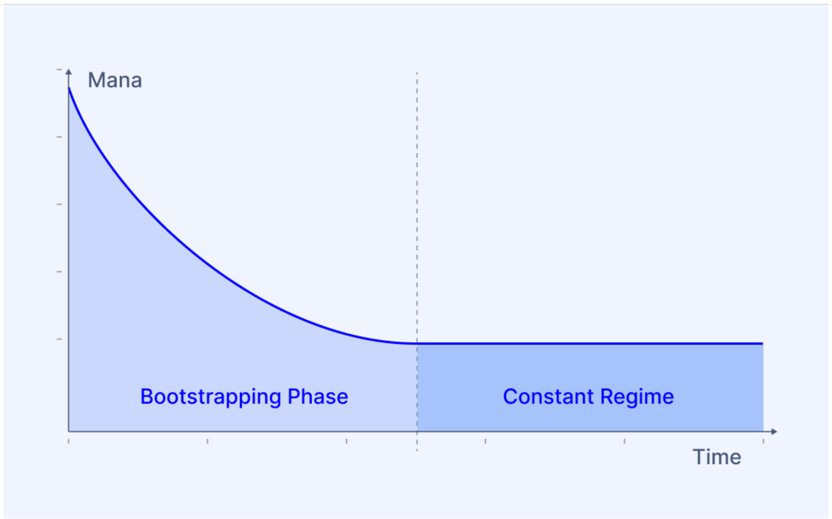

Figure 4.2: Qualitative behavior of the target rewards over time

mature network so the validators need to be able to hoard Mana to sell in the future. The decaying regime of the rewards is designed in such a way that Mana can be accumulated over time at the early stages, even in the presence of the Mana decay factor, as it can be verified in Appendix D. This is due to the fact that the contribution premium in the bootstrapping phase exceeds its decay; thus, one can hoard Mana. The possibility of hoarding Mana during the bootstrapping phase incentivizes validators to contribute to consensus (and, thus, to the network's utility) as early as possible.

On the other hand, the constant regime is designed for a mature and high-utility network when the protocol optimizes the incentives for our users. The combination of a constant rewards regime and Mana decay incentivizes the spending of Mana and makes any hoarding of it over time unfruitful. In this phase, the share of Mana given to consensus contributors stabilizes around a particular value controlled by the protocol. Appendix D provides a mathematical analysis of the different reward regimes.

Property 6: Minimum guaranteed profitability. Our reward scheme follows the *1:2:3 rule*. This means that if someone holding a certain amount of tokens gets x Mana per epoch and delegates this same stake or runs a validator node, the person would get at least $2x$ and $3x$ Mana, respectively.

Provided a validator performs their duties correctly, the minimum of $3x$ has no conditions attached; no matter their stake or the stake of the other validators, their profitability of at least $3x$ is guaranteed. Since rewards are calculated based on the share of the total stake instead of the absolute value of the stake, if the total staked value in the system is low, the validators can expect to gain substantially more than $3x$. This property is highly desirable since it incentivizes validators the most when the system needs more stake, thus attracting new validators to the economy and increasing the system's security.

In the case of delegators, their profitability will depend on how much stake is already delegated to their pool. Pools with an excessive delegated stake (compared to their locked stake) will reward delegators less. As described in Properties 3 and 4, this will incentivize delegators to delegate to less populated pools. Thus, rewarding $2x$ is not guaranteed to delegators *independently of their behavior*. Instead, by design, it is guaranteed that at least one of these less populated pools will reward their delegators with at least $2x$. Analogously to the validator's case, delegators can expect to gain substantially more than $2x$ if the total staked value in the system is low.

For a detailed analysis of this mechanism, see section C.4.

Rewards distribution

After calculating the target reward $R(n)$ per epoch n, given by 4.1, it is possible to define the actual reward $R_i(n)$ that should be distributed to each pool i. We reiterate that the pool reward distribution is not exactly linear to its stake. Thus, given that two pools have the same total stake, our incentive scheme will offer more rewards to the pool with the higher validator stake. The exact formulas are defined in Appendix B. The pool rewards are also proportional to a performance factor that measures the quality of validator i's services at each epoch.

After calculating the pool combined rewards $R_i(n)$, a fixed cost c_i (declared in the validator registration) is discounted from the rewards and assigned first to the validator. Since the validator's fixed cost is public information, delegators will know whether a validator is declaring a reasonable value for it. Thus, it would not be rational to delegate to a validator with an unreasonably large fixed cost. Furthermore, if the fixed cost is larger than $R_i(n)$, the validator is punished by not receiving any reward, so validators are incentivized to keep their fixed costs at fair valuations to attract delegators and not be punished.

From what is left after discounting the fixed cost, we discount the profit margin (which is set by the protocol, as a percentage of the pool reward), which is also assigned to the validator. After this second discount, the delegators and validators share their rewards proportionally to their stake.

4.5 Retrieving rewards

The reward to be distributed to each participant is an objective measure. With the exception of the performance factor, it can be calculated in advance, since it refers to public information from a point in time in the past, as depicted in Figure 4.3. Furthermore, nodes commit to this value. The performance measures are calculated *a posteriori*, but use public, objective information. Thus, every node can calculate and verify that the owed rewards are correct.

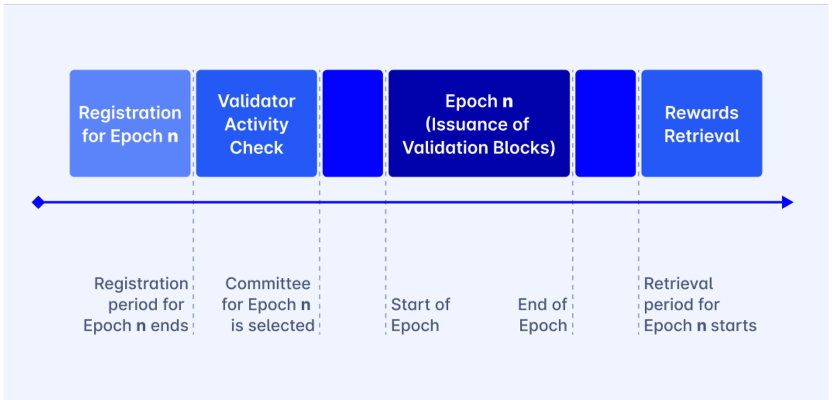
Registration for Epoch n
Validator Activity Check
Epoch n (Issuance of Validation Blocks)
Rewards Retrieval
Registration period for Epoch n ends
Committee for Epoch n is selected
Start of Epoch
End of Epoch
Retrieval period for Epoch n starts

Figure 4.3: Summary of the timing of the locking and rewarding mechanism

To retrieve their rewards from epoch n, the validators must consume their staking feature at a point in time after the end of epoch n. Both the committee selection and the rewards for epoch n are calculated relatively to the same point in time (in epoch $n-1$). Additionally, the information about rewards and performance factors must be tracked at the node level. To prevent the nodes from storing unreasonably large amounts of information about these rewards, the claims must be made before a specific expiration period (1 year).

5

IOTA token supply and demand

The health of any token economy is reflected by its base asset (which is the IOTA token in the case of the IOTA economy) and by its price, which, as in traditional economics, is dictated by the principles of supply and demand. Therefore, the fundamental goal of our tokenomics must be to maintain and strengthen the IOTA token through the sustainable design of these two aspects.

The critical factor to consider regarding supply is that the IOTA token has no inflation, i.e., the total supply is fixed as no rewards are given in the form of IOTA tokens. This is in stark contrast to most other DLTs where inflation is used to incentivize profit-driven block producers to secure the ledger. The fact that the IOTA token supply is fixed greatly simplifies our tokenomics from a supply perspective. It makes IOTA very well suited as a store of value since it is not a naturally (i.e., by design) depreciating asset.

However, a fixed supply alone does not make the IOTA token valuable; it also needs demand. Sustainable demand for any DLT's native token must stem from the core utility provided by the ledger, and IOTA is no exception. The tokenomics scheme presented here exemplifies this principle. In the following, we explain in detail what we expect to be the drivers of demand for the IOTA token, each corresponding to valuable features of the DLT.

Mana generation

Access to the ledger without paying fees in the base token has been at the heart of IOTA's value proposition since its inception, and this crucial feature will remain in IOTA 2.0. It is important to note that *no fees* does not mean that access to modify the ledger is free. Instead, this access is allocated to those who contribute to the protocol in the form of Mana. Requiring no IOTA token fees will generate demand from a broad spectrum of prospective IOTA token holders, from infrequent and casual users (whose access needs will be satisfied by the Mana passively generated from their tokens) to large organizations whose business models are built on top of the IOTA protocol and who operate validators to support their block creation requirements.

Long-term Investments

In the context of DPoS blockchains, in which delegators and validators stand to gain significant returns over a short period in the form of token rewards, return on investment typically refers to yield gains associated with a relatively short-term investment. Such systems based on short-term profits can be seen as unsustainable in the long run (see [8]) because staking rewards are ultimately paid for by token holders, either through explicit fees or by the devaluation of their token holdings due to inflation, or even both. In other words, exploitative users purely driven by profit stand to gain the most from these schemes; as soon as they can cash out their rewards, they leave the network's other users with a deficit. A common way for a DLT to incentivize (or, rather, force) long-term investments is to require the locking of assets to gain rewards, as we see in blockchains such as Ethereum[1]. However, such long-term locking requirements create significant opportunity costs for stakers, which must be compensated with distributed rewards.

IOTA 2.0 takes a different approach to encourage long-term strategies. *Since owning a fraction of the IOTA token supply guarantees a share of throughput in the network, owning tokens brings highly appealing rewards to actors with long-term strategies who seek to build lasting infrastructure on IOTA.*

Governance

A fundamental characteristic of a decentralized DLT is that a single entity is not making all network decisions. Those who hold IOTA tokens (and therefore have an economic incentive to make good decisions) can participate in protocol governance. Actors that use their IOTA tokens to participate in governance have an increased perception of their tokens' value. This effect may be hard to measure, but tokens with governance rights have advantages over tokens that do not offer these functionalities.

An essential part of governance incentivization and implementation is the technical possibility of delegating governance rights to someone else. Many token holders use this functionality in other networks because they may not follow a project closely and do not feel knowledgeable enough to make educated decisions about the project's future. However, this delegation scheme must differ from the delegation of stake to receive rewards. Suppose we would only have this single delegation functionality (for both rewards and governance); in that case, we would have a single class of users (i.e., validators) dictating all decisions in the network. As they have different incentives than token holders, this may lead to unfavorable governance decisions for most holders. Thus, token holders can independently decide who gets their consensus (validator delegation) and governance weight (governance delegation). This could be, of course, the same

[1] `https://ethereum.org/en/staking/solo/`

entity, but it also could be distinct entities with entirely different natures and motivations.

6

Conclusion

The core concept behind the IOTA 2.0 incentives and tokenomics is to reward contributions to the network with access. In this paper, we have presented the key components of the protocol that realize this concept. We described the different roles embodied by users of the IOTA ecosystem, including block creators, token holders, delegators, and validators, and how each of these roles fits into IOTA's economic model. We introduced Mana, the reward resource that can be used to access the system and make use of its capabilities via our Mana-based congestion control. We presented our access-based reward scheme for stakers and delegators, the properties it achieves, and how it achieves them. Finally, we focused on the IOTA token itself, and how our economic model will impact it in terms of supply and demand.

The fact that holding IOTA tokens generates Mana eliminates the need for fees for token holders in IOTA 2.0. This is complemented by the fact that rewards for all contributors are provided in the form of Mana, thereby avoiding inflation or deflation of the base token. The fixed IOTA token supply, combined with the absence of fees and inflation, removes the inequitable wealth distribution found in many cryptocurrencies where end users lose value to block producers and validators.

Furthermore, because of the virtuous circle of token holders generating their own Mana that they can use to create their own blocks, a whole raft of value-extracting middlemen is removed from IOTA 2.0's tokenomics. Our DAG-based consensus protocol enables leaderless access, which further protects end users from censorship and empowers them to take full advantage of all that distributed ledger technology has to offer.

Because we reward all protocol contributors with Mana, we tie rewards to the real tangible utility of the IOTA ecosystem, which is embodied by the demand for access. This discourages contributors from cashing out their rewards early to make a quick profit, encouraging long-term commitment to the protocol instead. This is because Mana will be most valuable when there is a high demand for access to the IOTA network and its resources. In this way, IOTA 2.0's incentives are aligned with providing a powerful distributed ledger that is useful and accessible for everyone.

Rewarding participants with Mana and not a cryptocurrency also provides flexibility to anyone unable or unwilling to participate based on regulatory reasons. This opens up IOTA 2.0 to an even broader class of users than other projects.

IOTA 2.0's tokenomics challenges the norms that have evolved in the crypto space in recent years, norms that have led to the exclusion and exploitation of so many. We believe that a cryptocurrency should not be an exploitative environment and that digital autonomy should not just be for those with large reserves of money and hardware. It should be for everyone.

Appendix A

Technical details of Mana

As discussed in the opening sections of this document, Mana is a resource that lives in the IOTA ledger. To provide a more precise understanding of where Mana resides within the ledger, we must delve into some underlying motivations.

First and foremost, let's examine the process of burning Mana. The reason we would opt for Mana to be an account-based asset is its effectiveness in serving as a deterrent against spam attacks. This approach simplifies the separation of congestion and spam control from transaction execution and validation. Without an account-based system, both transactions would have to be executed for a "double-burn" of Mana to render them invalid. This would need the propagation of blocks containing these transactions throughout the network. Thus, Mana would primarily serve as a mechanism to prevent ledger pollution but would not directly contribute to spam control during transaction execution and dissemination.

On the other hand, we aim to leverage the IOTA UTXO ledger's inherent characteristics to enable parallel processing of transactions involving simple Mana transfers.

Consequently, we've devised a solution where Mana can exist in either of two forms. Initially, when claimed, Mana is stored in outputs, akin to IOTA tokens. This form of Mana stored in outputs is freely transferable and tradable. However, it lacks the capability to be used for block creation. To be used for block issuance, this Mana must first be allocated to an account, wholly or in part. Given that users must have an account to issue blocks, this allocation process poses no significant hurdle. Subsequently, the Mana associated with the account can be employed for block issuance.

It's essential to note that once Mana is linked to an account, it can no longer be stored back in an output and, as a result, cannot be traded. For more details about Mana and its forms, see [9].

A.1 Mana decay

As mentioned in Section 3.3, we have introduced a global decay rate on all forms of Mana. The rationale behind implementing Mana Decay is to encourage spending while discouraging hoarding for extended periods. This approach ensures that all Mana in the system undergoes continuous decay consistently, making it a "non-gameable" system where users cannot exploit the system to gain more Mana. For example, if the Mana linked to an account's decay was slower than the Mana stored in an output, a user could keep all their Mana linked to their account.

We introduce an exponential decay of parameter β, set according to the desired Mana behavior. Suppose some Mana represented in the system with amount $M(t)$ at time t is updated at time $t + \Delta$, where Δ is a positive value representing the time since this Mana was last decayed. The Mana amount is updated as follows:

$$M(t + \Delta) = M(t) \exp(-\beta\Delta). \tag{A.1}$$

Decay in accounts balances and Mana in outputs: Now, let's break down how this decay affects account balances and Mana in outputs. Suppose you have an output created at time t, holding M Mana and no IOTA tokens, so there will not be any generation of Mana associated with this output. If this Mana is moved at time $t + \delta$, the new output will hold the decayed value of M. The Mana allowed to be stored on the new output will be then $M \exp(-\beta\delta)$, by A.1.

The same principle applies to Mana associated with accounts. If an account held M Mana at time t and no burning or allotments of Mana occurred, at time $t + \delta$, this account will hold $M \exp(-\beta\delta)$. Notice that frequent updates of account balances might be impractical and even useless unless the account receives new Mana or burns Mana for block issuance. As an example, take this same account from the example above. If someone had updated the account balance to account for its decay twice (once at $t + \delta/2$ and once at $t + \delta$), the resulting balance would be the following:

$$\begin{aligned} M(t+\delta) &= M(t+\delta/2)\exp(-\beta\delta/2) \\ &= M\exp(-\beta\delta/2)\exp(-\beta\delta/2) \\ &= M\exp(-\beta\delta) \end{aligned}$$

which is the same result as updating the account's balance just once, at time $t + \delta$.

Decay and mana generation by outputs: Now, let's consider the generation of Mana by outputs. Suppose you have an output created at time t, holding no Mana but S IOTA tokens. This output in principle would be only exposed to Mana generation. However, this Mana generation has to also be exposed to the same decay rate as the Mana in outputs; otherwise, one could profit by delaying the claiming of the generated Mana. Thus, we also apply the decay *while it is being generated.*

This is generated as follows:

- Each IOTA token will generate γ Mana per slot[1].
- Then, a decay factor of $\exp(-\beta\Delta)$ (where Δ is the epoch duration) is applied to the already generated Mana at the end of each epoch.

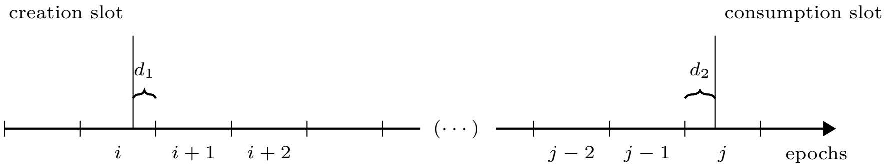

Figure A.1: Auxiliary time division scheme for Mana Generation

To model exactly what happens to the generation of Mana from an output holding S IOTA tokens, created at *creation slot* and consumed at *consumption slot*, see Figure A.1. Assume d_E is the number of slots per epoch, and that $j > i+1$.

The Mana generated at epoch i, (i.e. $S\gamma d_1$), should be decayed $j-i$ times, since it crosses $j-i$ epoch boundaries. Generally speaking, the Mana generated in any epoch k between $i+1$ and $j-1$ (i.e. $S\gamma d_E$) crosses $j-k$ decay boundaries, so it must be decayed $j-k$ times. Finally, the Mana generated in epoch j (i.e. $S\gamma d_2$) is not decayed at all. Adding these values, we find the following value for the Mana generation of such an output:

$$
\begin{aligned}
&S\gamma d_1 \exp(-\beta\Delta(j-i)) + \sum_{k=1}^{j-i-1} S\gamma d_E \exp(-\beta\Delta k) + S\gamma d_2 \\
=&S\gamma d_1 \exp(-\beta\Delta(j-i)) + S\gamma d_E \exp(-\beta\Delta)\frac{1-\exp(-\beta\Delta(j-i-1))}{1-\exp(-\beta\Delta)} + S\gamma d_2
\end{aligned}
$$

Analogously, if $j = i+1$, the Mana generated will be $S\gamma d_1 \exp(-\beta\Delta n) + S\gamma d_2$; if $j = i$, it will be $S\gamma\delta$, where δ is the difference between the creation and consumption slots.

Decay and rewards: Finally, recall that Mana is also provided to validators and delegators as rewards for their contributions. Those actors can choose when to claim their rewards. Thus, decay must be applied identically to the distributed rewards as to the Mana stored in outputs, otherwise validators could delay the claiming as much as possible in order to maximize their Mana gains. Then, suppose a validator has the right to rewards of value R, from a certain

[1] A slot is a subdivision of an epoch; each epoch is defined as 2^{13} slots

epoch n, which ends at time t. If said validator claims the rewards at time $t+\delta$, then, the Mana to be distributed at the claiming time will be:

$$R\exp(-\beta\delta). \tag{A.2}$$

A.2 Validity conditions for transactions

Because Mana is associated with unspent outputs, the fact that Mana decays and is generated while stored makes a transaction's validation logic different regarding IOTA and Mana values. Let $\mathcal{I}$ be the set of inputs of a transaction and let $\mathcal{O}$ be the set of outputs. If IOTA_i is the IOTA value of output i, then the validation regarding IOTA values is done as usual:

$$\sum_{i\in\mathcal{I}}\text{IOTA}_i = \sum_{i\in\mathcal{O}}\text{IOTA}_i$$

On the other hand, the validity conditions for Mana must take into account all the decay and generation factors defined above. Furthermore, Mana can be *forfeited*, meaning that if users wish not to claim the Mana associated with the generation due to their IOTA holdings, they can just not store it in outputs.

Finally, we want to emphasize that the validity conditions cannot be inconsistent between different nodes due to different architectures.

Since floating point operations might lead to inconsistencies among nodes due to the different possible rounding behaviors in different architectures, fixed point arithmetic (which does not expose the nodes to these rounding divergencies) must be used in all the validity condition-related algorithms. In particular, all the Mana and rewards calculations have to be done with fixed-point arithmetic. For more information about how the fixed point arithmetic should be implemented, see [9].

Appendix B

Rewards calculation

In the following, we specify the key details of our proposed approach to staking. This requires some notation, which is summarized in Table B.1.

Table B.1: Notation for staking

V	validator set
m	profit margin of validators
c_i	fixed costs of validator i
p_i	token value staked by validator i
P	total value staked by validators in V
D_i	token value delegated to validator i
D	total value delegated to validators in V
S_i	token value staked (by locking and delegation) of validator i's pool
d_{ji}	token value delegated by account j to validator i
r_i	probability of validator i being in the validator set
$R(n)$	total rewards targeted for the epoch n
$R_i(n)$	rewards to validator i's pool (i.e., to validator i and their delegators)

To properly calculate the rewards, we first define the target reward $R(n)$ per slot n, given by:

$$R(n) = \begin{cases} R\exp(-\beta n\Delta), \text{ if } n \leq \frac{T}{\Delta}, \\ c, \text{ if } n > \frac{T}{\Delta,} \end{cases} \tag{B.1}$$

where c is a constant set in Section D.1, R is a protocol parameter set in Appendix E, Δ is the length of a epoch, and T is the duration of the network's early stage. The motivation behind this reward function, together with its properties, is explored in the Appendices D and E. To sum up, the decaying regime incentivizes the validators to contribute as early as possible, while the constant regime stabilizes the Mana given to validators around a share θ of the total distributed Mana.

Given the target reward $R(n)$ in a given epoch n, we proceed to define how rewards are distributed among participants. Most of the variables used (as the stake of each participant) assume different values for each epoch, but because we are fixing the epoch to n, to make the notations clearer, we drop the dependency on n for these variables. We assume it is implicit that we refer to all these variables at the fixed epoch.

Let r_i be the probability of actor i being chosen for the committee at epoch n, $R(n)$ the total rewards targeted for said single epoch, and α a parameter ranging from 0 to infinity. Additionally, φ_i (within the range of 0 to 1) represents the *performance factor* that assesses the quality of validator i's services at epoch n. To start, let's determine the distribution of the total target rewards, $R(n)$, among the various pools. When validator i is selected to participate in the committee for epoch n, the rewards for pool i are calculated as follows:

$$R_i(n) = \frac{R(n)}{1+\alpha}\frac{1}{r_i}\left(\frac{p_i + D_i}{P+D} + \alpha\frac{p_i}{P}\right)\varphi_i. \tag{B.2}$$

An in-depth analysis of how the variable α influences the rewards is presented in Section B.1 of this Appendix (in simple words, the larger the α, the larger the incentive to lock tokens to stake).

Moreover, this is a universal formula that applies to any committee selection method. For committees chosen through random processes, r_i must be computed according to the specific procedure, taking into account factors such as selection weights, caps, and the possibility of a validator holding multiple committee seats.

In the case of fixed committees and committees based on the top stakers, a slightly different formula is employed:

$$R_i(n) = \frac{R(n)}{1+\alpha}\left(\frac{p_i + D_i}{P_C+D_C} + \alpha\frac{p_i}{P_C}\right)\varphi_i. \tag{B.3}$$

In this scenario, P_C and D_C represent the locked and delegated stake within the committee, respectively. It's important to note that if all committee members have a performance factor of $\varphi_i = 1$, the total reward distributed per epoch will be as follows:

$$\sum_C R_i(n) = \sum_C \frac{R(n)}{1+\alpha}\left(\frac{p_i + D_i}{P_C+D_C} + \alpha\frac{p_i}{P_C}\right) = R(n). \tag{B.4}$$

In this case, the sum encompasses all committee members, and the reward distributed per epoch consistently matches the target reward. For the sake of generality (as it is always possible to set $r_i = 1$, $P = P_C$, and $D = D_C$), we will assume that the committee selection is randomized going forward.

Since a validator is selected to be part of the committee at a certain epoch with probability r_i, and, in that case, the reward given to this validator would be $R_i(n)$, the expected rewards $\mathbb{E}(R_i(n))$ at epoch n will be given by:

$$\mathbb{E}(R_i(n)) = r_i R_i(n) = \frac{R(n)}{1+\alpha}\left(\frac{p_i + D_i}{P+D} + \alpha\frac{p_i}{P}\right)\varphi_i. \tag{B.5}$$

The equation above provides rewards to pools in proportion to their share of the total stake within the system. Consequently, in situations where the total stake within the system is relatively low, the potential rewards for those who choose to delegate or stake their tokens are significantly magnified. This characteristic is highly desirable, as it serves to attract stakeholders when the overall stake in the system is limited. At such times, the system is most in need of enticing participants to enhance its security.

After calculating the combined reward given by equation (B.2) the fixed cost c_i declared in the validator registration is discounted from $R_i(n)$, and given first to the validator. If the fixed cost is larger than $R_i(n)$, no rewards are given to the validator as a punishment. Moreover, to attract delegators, validators are incentivized to keep their fixed costs at fair valuations. From what is left after discounting the fixed cost, we discount the profit margin (which is set as a percentage), which is also given to the validator. After this second discount, the delegators and validators share their rewards proportionally to their stake. Mathematically, this means that the reward to each actor will be given by:

$$\begin{cases} R_i^i(n) = \min\left(R_i(n), c_i + mR_i(n)\right) + \max(0, (1-m)R_i(n) - c_i)\dfrac{p_i}{p_i + D_i} \\ R_i^j(n) = \max(0, (1-m)R_i(n) - c_i)\dfrac{d_{ji}}{p_i + D_i}, \text{ if } j \neq i \end{cases} \tag{B.6}$$

B.1 Effects of α on the incentive mechanics

In this appendix, we analyze how the choice of the parameter α can affect the system in a qualitative manner. We begin by recalling the combined expected reward function to a set of delegators and a validator with D_i and p_i delegated and staked, respectively:

$$\mathbb{E}(R_i(n)) = \frac{R(n)}{1+\alpha}\left(\frac{p_i + D_i}{P + D} + \alpha\frac{p_i}{P}\right)\varphi_i. \tag{B.7}$$

We assume that the validation tasks were done correctly (i.e., $\varphi_i = 1$), and define $S := \sum\limits_{i \in V} S_i$, $K := P/D$ and $k_i := p_i/D_i$ (which implies $S = P + D = (K+1)D$ and $S_i = p_i + D_i = (k_i + 1)D_i$). Then, the combined rewards can be expressed as:

$$\mathbb{E}(R_i(n)) = \frac{R(n)}{1+\alpha}\frac{S_i}{S}\left(1 + \alpha\frac{k_i(K+1)}{K(k_i+1)}\right). \tag{B.8}$$

The derivative of this function on α is positive if and only if $k_i > K$. Note that K can be seen as a weighted average of the pools' k_is (more specifically, we have that $K = \sum\limits_{i \in V} \frac{D_i}{D} k_i$), so if a pool has $k_i > K$, another pool must have $k_i < K$.

Furthermore, the limit of this function when $\alpha \to \infty$ equals $R(n)\frac{S_i}{S}\frac{k_i(K+1)}{K(k_i+1)}$. This means that an increase in α benefits validators with a larger share of locked staked value compared to total stake (i.e., with larger k_i's) and decreases the rewards of validators with smaller k_i's, which implies that larger α's mostly tend to incentivize (locked) staking instead of delegation.

To better exemplify this effect, assume all pools have the same stake S_i. For the sake of simplicity, we fix $K = 1$. Defining $\beta = S_i/S$ we have:

$$\frac{\mathbb{E}(R_i(n))}{\beta R(n)} = \frac{1}{1+\alpha}\left(1 + \frac{2\alpha k_i}{k_i+1}\right). \tag{B.9}$$

Figure B.1 represents the above equation as a function of α, calculated for selected values of k_i's. As expected, for $k_i = 1$ (meaning that $k_i = K$), the curve is a straight horizontal line, meaning that it does not depend on α, as stated in the last section. Also, since in this case, all validators have the same S_i, K must correspond to the average k_i among all validators; thus either all $k_i = 1$, or a combination of validators with different k_i (some larger than 1, some smaller than 1) must coexist. We finally conclude that an increase of α corresponds to an increase in the incentive to stake, since in this case, the rewards will benefit even further the validators with larger k_i's when compared to the ones with smaller k_i's. However, the larger the α, the less sensitive it is to increases (recall that on infinity, the scaled rewards defined above stabilize on $\frac{k_i(K+1)}{K(k_i+1)}$, as stated in the previous section).

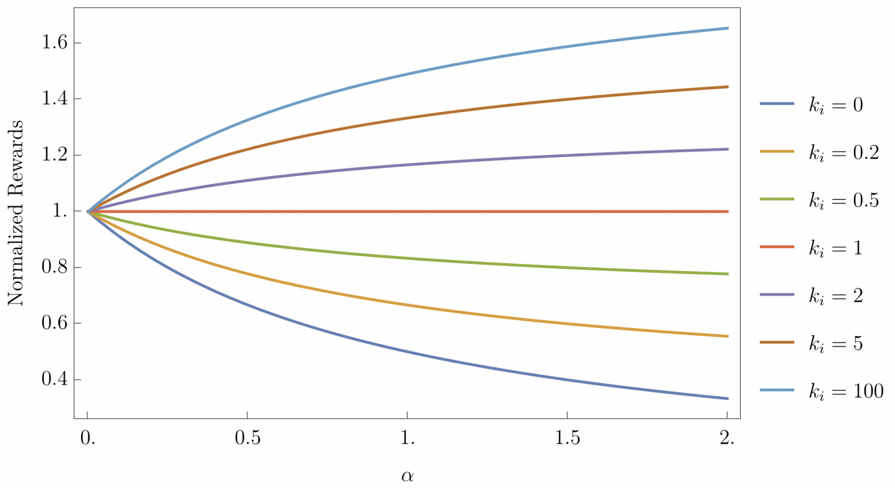

Figure B.1: $\dfrac{\mathbb{E}(R_i(n))}{\beta R(n)}$ as a function of α, for selected values of k_i's

The case of homogeneous k_i's

Pools with the same share of delegated value compared to stake (i.e., with the same k_i) will receive combined rewards proportional to their stake S_i (without dependence on α). Note that, if all validators have the same $k_i = k$, then $K = k$ too, resulting in:

$$\mathbb{E}(R_i(n)) = \frac{R(n)}{1+\alpha}\frac{S_i}{S}\left(1+\alpha\frac{k(k+1)}{k(k+1)}\right) = R(n)\frac{S_i}{S}. \tag{B.10}$$

Also note that, as indicated in Appendix C.3, rational actors will make the system tend to an equilibrium of k_is, so it is expected that the formula above will hold in practice. This does not mean that validators are not incentivized to stake, because, in the homogeneous k_i case, a smaller stake implies a smaller delegation and, consequently, a smaller S_i.

Appendix C

Incentives compatibility of the reward function

A well-designed reward function should naturally align with the desired incentivization goals of the system. To better understand this alignment, we will refer to the requirements established in Chapter 4.4:

1. Non-gameability of the locking rules: there are incentives for the validators to stake their funds instead of delegating them to themselves.
2. Incentives for high-quality validation services: there are incentives for the validators to perform their expected consensus-related activities correctly.
3. No incentives for the concentration of validators' stake.
4. Incentives to spread delegation among all validators.
5. Larger incentives to early contributions to increase the network's security.
6. Minimum guaranteed profitability.

In the following sections, we will provide evidence that the reward function defined by equations (B.2) and (B.6) does indeed meet these requirements. We will skip the proof for properties 2 and 3, as they straightforwardly follow from the reward formula. For the sake of simplicity, we will assume that $\varphi_i = 1$ for all validators, and assume that $(1-m)R_i(n) - c_i > 0$ (i.e., validators are rational actors who provide reasonable values for c_i). Then, we can calculate the expected rewards for each actor as follows:

$$\begin{cases} \mathbb{E}(R_i^i(n)) & = c_i + m\mathbb{E}(R_i(n) - c_i) + (1-m)\mathbb{E}(R_i(n) - c_i)\dfrac{p_i}{p_i + D_i} \\ & = (1-m)c_i\dfrac{D_i}{p_i + D_i} + \mathbb{E}(R_i(n))\dfrac{p_i + mD_i}{p_i + D_i}, \text{ for validators} \\ \mathbb{E}(R_i^j(n)) & = (1-m)\mathbb{E}(R_i(n) - c_i)\dfrac{d_{ji}}{p_i + D_i}, \text{ if } j \neq i, \text{ i.e., for delegators} \end{cases} \tag{C.1}$$

C.1 Non-gameability of the locking rules

We aim to analyze the rational choice for validators seeking to maximize their rewards while adhering or not to the rules governing token locking. To do so, we will evaluate the rewards distributed to a specific validator during a fixed epoch under two different scenarios. Suppose the validator holds a total of $p+x$ tokens and can opt for one of the following strategies:

- Add the additional tokens x to the already locked portion p (**Strategy 1**).
- Delegate the additional tokens x to themselves, along with d tokens that were delegated to them by others (**Strategy 2**).

Let's break down each strategy:

Strategy 1: In this scenario, the reward allocated to the validator's pool would be calculated as:

$$R_{P1} = \frac{R(n)}{1+\alpha}\left(\frac{p+d+x}{P+D+x} + \alpha\frac{p+x}{P+x}\right) \tag{C.2}$$

From this pool reward, the validator will receive:

$$c_i + m(R_{P1} - c_i) + (1-m)(R_{P1} - c_i)\frac{p+x}{p+d+x}, \tag{C.3}$$

which accounts for their fixed cost, profit margin, and their share of the remaining rewards.

Strategy 2: In this case, the reward distributed to the validator's pool is determined as:

$$R_{P2} = \frac{R(n)}{1+\alpha}\left(\frac{p+d+x}{P+D+x} + \alpha\frac{p}{P}\right) \tag{C.4}$$

From this pool reward, the validator will receive:

$$\begin{cases} c_i + m(R_{P2} - c_i) + (1-m)(R_{P2} - c_i)\dfrac{p}{p+d+x} & \text{from validating} \\ (1-m)(R_{P2} - c_i)\dfrac{x}{p+d+x} & \text{from delegating} \end{cases} \tag{C.5}$$

Comparison: Let $R_1(n)$ represent the expected validator reward for Strategy 1, and $R_2(n)$ for Strategy 2. We can calculate the difference between these rewards as follows:

$$R_1(n) - R_2(n) = c_i + m(R_{P1} - c_i) + (1-m)(R_{P1} - c_i)\frac{p+x}{p+d+x} \tag{C.6}$$
$$- c_i - m(R_{P2} - c_i) - (1-m)(R_{P2} - c_i)\frac{p+x}{p+d+x} \tag{C.7}$$
$$= m(R_{P1} - R_{P2}) + (1-m)(R_{P1} - R_{P2})\frac{p+x}{p+d+x} \tag{C.8}$$
$$= (R_{P1} - R_{P2})\frac{md+p+x}{p+d+x} \tag{C.9}$$

Using the previously calculated values, we find:

$$R_1(n) - R_2(n) = \frac{R(n)\alpha}{1+\alpha}\frac{md+p+x}{p+d+x}\left[\frac{p+x}{P+x} - \frac{p}{P}\right] \tag{C.10}$$
$$= \frac{R(n)\alpha}{1+\alpha}\frac{md+p+x}{p+d+x}\frac{x(P-p)}{P(P+x)} \geq 0 \tag{C.11}$$

Hence, we can conclude that staking the tokens instead of delegating to oneself is financially advantageous for the validator. It's worth noting that in the above relationship, equality (rather than inequality) would only hold if the validator is the sole participant in the network. This situation is not expected in practice, making the choice to lock tokens the default rational one for honest validators.

C.2 Nonexistence of incentives to centralization of validators funds

We demonstrate that when two validators, each with locked values p_1 and p_2, and delegated values D_1 and D_2 pool together, their combined expected reward remains unchanged, equal to the sum of their initial rewards. For simplicity, we assume that the fixed cost is negligible compared to the reward, setting $c_i = 0$

Let $R_b(n)$ represent the expected collective reward before pooling, and $R_a(n)$ after pooling:

$$R_b(n) = \frac{R(n)}{1+\alpha}\left(\frac{p_1+D_1}{P+D} + \alpha\frac{p_1}{P}\right) + \frac{R(n)}{1+\alpha}\left(\frac{p_2+D_2}{P+D} + \alpha\frac{p_2}{P}\right) \tag{C.12}$$

$$R_a(n) = \frac{R(n)}{1+\alpha}\left(\frac{p_1+p_2+D_1+D_2}{P+D} + \alpha\frac{p_1+p_2}{P}\right) \tag{C.13}$$

Hence, $R_a(n) - R_b(n) = 0$, demonstrating that the collective rewards remain the same. In this zero-sum game, the validator (and their delegators) with a smaller ratio $\dfrac{p_i}{D_i+p_i}$ will benefit from pooling, while the other will not.

Now, let's define $R_b^V(n)$ as the expected validators' reward before pooling, and $R_a^V(n)$ after pooling:

$$R_b^V(n) = \frac{R(n)}{1+\alpha}\left(\frac{p_1+D_1}{P+D} + \alpha\frac{p_1}{P}\right)\frac{mD_1+p_1}{p_1+D_1} \tag{C.14}$$

$$+ \frac{R(n)}{1+\alpha}\left(\frac{p_2+D_2}{P+D} + \alpha\frac{p_2}{P}\right)\frac{mD_2+p_2}{p_2+D_2} \tag{C.15}$$

$$R_a^V(n) = \frac{R(n)}{1+\alpha}\left(\frac{p_1+p_2+D_1+D_2}{P+D} + \alpha\frac{p_1+p_2}{P}\right)\frac{m(D_1+D_2)+(p_1+p_2)}{p_1+p_2+D_1+D_2} \tag{C.16}$$

Now, we can analyze the difference between these rewards:

$$R_a^V(n) - R_b^V(n) = -\frac{R(n)}{1+\alpha}\frac{\alpha(1-m)(D_2p_1-D_1p_2)^2}{P(D_1+p_1)(D_2+p_2)(D_1+D_2+p_1+p_2)} \leq 0 \tag{C.17}$$

This indicates that under the analyzed conditions, validators would not profit from pooling.

C.3 Equilibrium of delegated stake among all validators:

In this section, we demonstrate that when delegators aim to maximize their profits by switching validators, they tend to stabilize the ratio between locked and delegated tokens, denoted as k_i, for different validator pools. This equilibrium results in a uniform value, denoted as $\bar{[k]}$, which is equal to $P/(P+D)$, where P represents the total staked tokens, and D represents the total delegated tokens. We provide formal proof that rational redelegation of tokens decreases the disparity in the k_i ratios among affected pools.

The general problem: Let's consider the following problem: Suppose a delegator is currently staking with a validator who has a staked value of p_1, and D_1 tokens have been delegated to this validator by others. The delegator is presented with the option to redelegate their funds, with a value of d, to another validator with a staked value of p_2 and D_2 delegated tokens. For simplicity, we assume that the c is zero. The expected rewards in the current state are denoted as $R_B(n)$, and the rewards after redelegation as $R_A(n)$. Then, the difference in rewards is calculated as follows:

$$R_B(n) - R_A(n) = \frac{R(n)\alpha}{1+\alpha}\frac{(1-m)d}{P}\left(\frac{p_1}{p_1+D_1+d} - \frac{p_2}{p_2+D_2+d}\right) \tag{C.18}$$

$$= \frac{R(n)\alpha}{1+\alpha}\frac{(1-m)d}{P}\frac{(D_2+d)p_1-(D_1+d)p_2}{(p_1+D_1+d)(p_2+D_2+d)} \tag{C.19}$$

This difference represents the delegator's incentive to switch validators, which is valid if and only if the following condition holds:

$$\frac{p_1}{D_1 + d} \leq \frac{p_2}{D_2 + d} \tag{C.20}$$

Let's analyze two generalized scenarios:

Scenario 1: In this scenario, the ratio k_1 (locked tokens to delegated tokens) in the current pool is greater than or equal to the ratio k_2 in the second pool before redelegation:

$$\frac{p_1}{D_1 + d} \geq \frac{p_2}{D_2} > \frac{p_2}{D_2 + d} \tag{C.21}$$

In this case, C.20 does not hold, and the delegator will never profit from that redelegation.

Scenario 2: Here, the ratio k_1 is smaller than k_2 before redelegation In this scenario, it is not guaranteed that redelegating *all* the d tokens will lead to profit. However, the delegator can always redelegate *part* of their tokens. In particular, assume the delegator redelegates xd, while keeping $d(1-x)$ in the current delegator. In this case, the difference in rewards can be expressed as:

$$R_B(n) - R_A(n) = \frac{R(n)\alpha}{1+\alpha}\frac{(1-m)d}{P} \tag{C.22}$$

$$\times \left(\frac{p_1}{p_1 + D_1 + d} - \frac{p_2 x}{p_2 + D_2 + xd} - \frac{p_1(1-x)}{p_1 + D_1 + (1-x)d}\right) \tag{C.23}$$

The derivative of this difference with respect to x at $x = 0$ is negative, indicating that there is always a redelegation value xd that would make the delegator profit from redelegation. This redelegation reduces the difference between the k_i ratios of the pools. Importantly, this argument is always valid as long as $k1 < k2$ and ceases to be valid if $k1 >= k2$. Furthermore, the value of xd that results in a profitable redelegation is never large enough to make k_1 larger than k_2 after the redelegation. To prove that, note that in the case k_1 becomes larger than k_2 after the redelegation:

$$\frac{p_2}{p_2 + D_2 + xd} < \frac{p_1}{p_1 + D_1 + (1-x)d} \tag{C.24}$$

and then

$$R_B(n) - R_A(n) > \frac{R(n)\alpha}{1+\alpha}\frac{(1-m)d}{P} \tag{C.25}$$

$$\times \left(\frac{p_1}{p_1 + D_1 + d} - \frac{p_2}{p_2 + D_2 + xd}\right) \tag{C.26}$$

On the other hand, if the delegator had chosen xd such that $k_1 = k_2$ after the redelegation, $R_B(n) - R_A(n)$ would be exactly the value in the right hand of the inequality above. Thus, it is always more profitable for the delegator to redelegate in such a way that, after redelegation $k_1 = k_2$ instead of $k_1 > k_2$. To sum up, rational delegators, in the long run, tend to stabilize the k_i parameters among different pools, leading to an equilibrium.

C.4 Profitability of validators and delegators

In this section, we establish the presence of minimum profitability for validators and provide assurance of guaranteed profitability for delegators[1].

Consider a system with a fixed validator stake denoted as P and a constant delegator stake denoted as D. Let's assume that p represents the stake of a specific validator. We will examine how the expected profit of a validator depends on the value d delegated to them. Our expected validator reward at epoch n, denoted as $E(R_V(n))$ (with $c = 0$ for simplicity), is expressed as:

$$E(R_V(n)) = \frac{R(n)}{1+\alpha}\left(\frac{p+d}{P+D} + \alpha\frac{p}{P}\right)\left(m + (1-m)\frac{p}{p+d}\right) \tag{C.27}$$

The above expression attains its minimum when:

$$\frac{d}{p} = \frac{\sqrt{\alpha m(1-m)P(D+P)} - mP}{mP} \tag{C.28}$$

We define the profit margin m in such a way that the minimum validator profitability is reached when the pools exhibit homogeneous delegated stake, meaning that $\frac{d}{p} = \frac{D}{P}$. This leads to a profit margin m given by:

$$m = \frac{\alpha P}{D + (1+\alpha)P} \tag{C.29}$$

With the profit margin defined as above, we can determine the minimum expected reward, denoted as $E(\bar{R_V}(n))$, for a validator with stake p:

$$E(\bar{R_V}(n)) = \frac{(\alpha+1)}{(\alpha+1)P + D}R(n)p \tag{C.30}$$

For delegators, the expected reward can be calculated as:

$$E(R_D(n)) = \frac{R(n)}{1+\alpha}\left(\frac{p+d}{P+D} + \alpha\frac{p}{P}\right)(1-m)\frac{d_i}{p+d} \tag{C.31}$$

We introduce a parameter $k = \frac{Pd}{Dp}$. Then, the equation becomes:

$$E(R_D(n)) = \frac{R(n)}{1+\alpha}\frac{(1-m)(\alpha(D+P) + Dk + P)}{(D+P)(Dk+P)}d_i \tag{C.32}$$

[1] The *minimum* profitability of the validators is assured no matter how much stake is delegated to them; the profitability of a certain delegator is guaranteed as long as said delegator makes a good choice of validator to delegate to.

This expression decreases as k increases. However, there always exists a pool with $k \leq 1$. As a result, a delegator can consistently choose a pool that ensures rewards of at least $E(R_D(n))^*$, where:

$$E(R_D(n))^* = \frac{(1-m)}{(D+P)} R(n) d_i = \frac{1}{D+(1+\alpha)P} R(n) d_i \tag{C.33}$$

In section E, we provide further details to ensure that the expected rewards for these actors follow a specific quantitative behavior, known as the 1:2:3 rule.

Appendix D

Behavior of the mixed rewards

The presence of a decay factor in the Mana rewarded to validators and delegators might lead to a lack of incentives for these actors to collaborate in the early stages of the network. To balance the effects of decay and to effectively incentivize said actors, we define a reward function that (in the beginning) decreases with time. Thus, we have designed a reward mechanism that, even in the presence of exponential decay, rewards early adopters proportionally more. We call this mechanism *mixed rewards* since it presents two different regimes, a decaying regime and a constant regime, which will be analyzed in the rest of this section.

Decaying Regime In this regime, we use the following decaying reward $R(n)$ formula:

$$R(n) = R\exp(-\kappa n\Delta) \tag{D.1}$$

where Δ stands for the duration of an epoch and κ is a positive incentivization factor. To exemplify how this mechanism can effectively incentivize early adopters, assume a system with a single validator from $t = 0$ to $t = T = N\Delta$, which never spent their Mana. We analyze the amount of Mana $M_n(T)$ in their account at time T, that was distributed as rewards at each epoch $n < N$:

$$\begin{aligned} M_n(T) &= R(n)\exp(-\beta\Delta(N-n)) \\ &= R\exp(-\kappa n\Delta)\exp(-\beta\Delta(N-n)) \end{aligned}$$

The total Mana $M(T)$ in their account is given by:

$$M(T) = \sum_{n=1}^{N} M_n(T) = R\exp(\beta\Delta)\frac{\exp(-\beta\Delta N) - \exp(-\kappa\Delta N)}{\exp(\kappa\Delta) - \exp(\beta\Delta)} \tag{D.2}$$

Then, the share of distributed Mana $\frac{M_n(T)}{M(T)}$ related to the rewards of each epoch equals:

$$\frac{M_n(T)}{M(T)} = k_N \exp(-(\kappa - \beta)\Delta n) \tag{D.3}$$

where $k_N = \exp(-\beta\Delta(N+1))\frac{\exp(\kappa\Delta)-\exp(\beta\Delta)}{\exp(-\beta\Delta N)-\exp(-\kappa\Delta N)}$. In this case, for $\kappa \geq \beta$, the validator is effectively incentivized to join the network early, since the share of Mana acquired in the first epochs dominates its total Mana. Figure D.1 represents the accumulated value of Mana $M(T)$, for $\beta\Delta = 0.1$, $\kappa\Delta = 0.12$ and $R = 1$ and different epochs k when the validator joined the network:

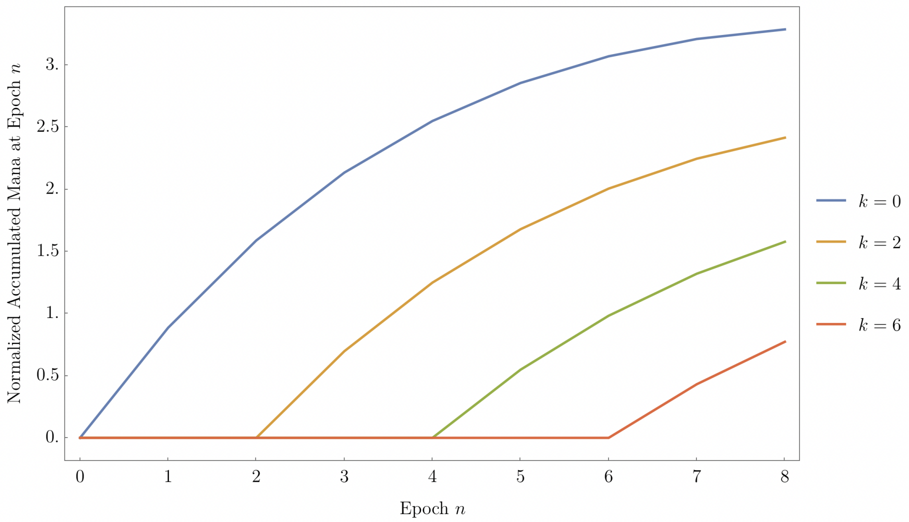

Figure D.1: Accumulated Mana for different epochs k when the validator joined the network, as a function of the present epoch (Decaying Rewards)

Now suppose a validator can only sell their Mana at epoch 8. In that case, it is in the interest of said validator to join at the very start, since joining at time 0 will reward them proportionally more Mana than joining at time $t = 6$. Notice that this conclusion is based on the fact that an early validator might want to sell their Mana at some point. Since we also assume Mana will hold no monetary value at the early stages of the network, the validator will save part of their Mana to be sold at some point in the future. However, similar arguments apply for validators that want to use the Mana generated in the early stages of the network, since larger shares of Mana imply a larger share of access rights.

Constant Regime In this regime, the total reward target $R(n)$ at epoch n is a constant R. To exemplify how this mechanism affects early adopters, assume a system with a single validator from $t = 0$ to $t = T = N\Delta$, which had never spent their Mana. We analyze the amount of Mana $M_n(T)$ in their account at time T, that was distributed as rewards at each epoch $n < N$:

$$M_n(T) = R\exp(-\beta\Delta(N-n))$$

The total Mana $M(T)$ in their account is given by:

$$M(T) = \sum_{n=1}^{N} M_n(T) = R\frac{e^{-\beta\Delta(N-1)}\left(e^{\beta\Delta N}-1\right)}{e^{\beta\Delta}-1} \tag{D.4}$$

and the share of distributed Mana $\frac{M_n(T)}{M(T)}$ related to the rewards of each epoch equals:

$$\frac{M_n(T)}{M(T)} = k_N \exp(\beta\Delta n) \tag{D.5}$$

where $k_N = e^{-\beta\Delta}\dfrac{e^{\beta\Delta}-1}{e^{\beta\Delta N}-1}$. This means that the Mana distributed in the first epochs of this regime does not contribute significantly to the total Mana of this validator, so this regime does not incentivize early validators as much as the decaying regime does. Figure D.2 represents a normalized value of Mana $(M(T)(1-\exp(-\beta\Delta))/R)$, for $\beta\Delta = 1$ and different epochs k when the validator joined the network:

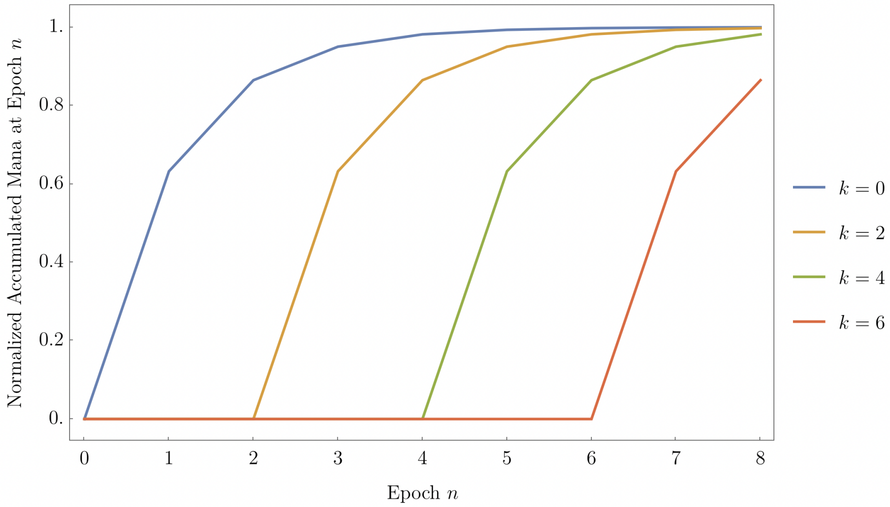

Figure D.2: $(1-\exp(-\beta\Delta))\dfrac{M(T)}{R}$ for different epochs k when the validator joined the network, as a function of the present epoch (Constant Rewards)

We conclude that, in the constant regime, there is no incentive to store Mana to be used or sold in the future. We expect that, at this stage, the network is already mature and rewards are rather quickly used to acquire access rights.

D.1 Mixed rewards

To capture the properties of both regimes described above at different stages of the network, we propose the following mixed reward function:

$$R(n) = \begin{cases} R\exp(-\beta n\Delta), \text{ if } n \leq \frac{T}{\Delta} \\ c, \text{ if } n > \frac{T}{\Delta} \end{cases} \tag{D.6}$$

where $c = \frac{RT}{\Delta}\left(\exp(\beta\Delta) - 1\right)\exp(-\beta(\Delta + T))$ and T is a time long enough so that the early adopters are expected to have already been given the chance of either spending or selling their early acquired Mana and the validators at this time no longer need to be proportionally more incentivized (meaning we assume that the utility of the network by itself provides enough incentives to the validators). Note that the maximum total Mana in the system given to validators at time T is:

$$M(T) = \frac{RT}{\Delta}\exp(-\beta T) \tag{D.7}$$

For $n > \frac{T}{\Delta}$, the total Mana distributed for validators is ruled by $M(\Delta n) = M(\Delta(n-1))\exp(-\beta\Delta) + R(n)$, i.e.:

$$\begin{aligned} M(T + \Delta) &= \frac{RT}{\Delta}\exp(-\beta T)\exp(-\beta\Delta) && \text{(D.8)} \\ &+ \frac{RT}{\Delta}\left(\exp(\beta\Delta) - 1\right)\exp(-\beta(\Delta + T)) && \text{(D.9)} \\ &= \frac{RT}{\Delta}\exp(-\beta T) = M(T) && \text{(D.10)} \end{aligned}$$

thus, by induction, $M(T + k\Delta) = M(T)$ for all $k \geq 0$, and the total Mana distributed to validators is constant after time T.

Figure D.3 represents the normalized accumulated value of Mana $M(n\Delta)/M(T)$, for $\beta\Delta = 0.1$, $\Delta T = 8$ and $R = 1$ and different epochs k when the validator joined the network:

In this case, at the early stages of the network (i.e., for $n\Delta \leq T$), it is in the interest of a validator to join at the very start, since joining at epoch 0 will reward them proportionally more Mana than joining at epoch $n = 6$, for instance. For $n\Delta > T$, on the other hand, the new validators can "catch up" with the early joiners, and the total Mana distributed to validators will tend to stabilize around a given value. We assume T is large enough so the network is already stable and that its utility is already a good enough incentive for validators to join.

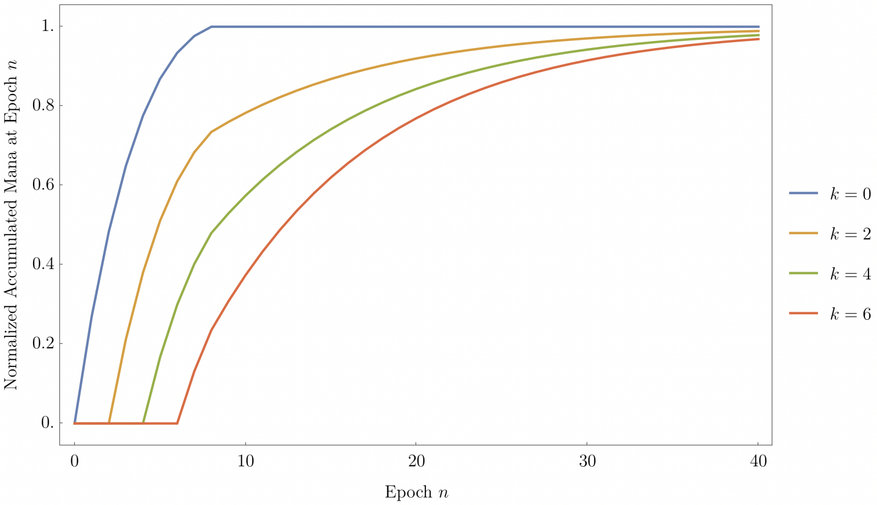

Figure D.3: Accumulated Mana for different epochs k when the validator joined the network, as a function of the present epoch (Mixed Reward)

Appendix E

Parameter setting

To implement the proposed mechanism, the following parameters are set:

1. T: duration of the network *early stage*
2. α: parameter in the reward function that dictates the incentive to lock tokens instead of delegating (and indirectly, the minimum profit margin m)
3. θ: fraction of the total Mana that is distributed to the consensus contributors (the validator's pools), when compared to the Mana distributed to holders and polls combined
4. β: the global decay parameter of Mana
5. R: scaling parameter of the distributed Mana to validators

We assume that Δ (epoch length) is given, being defined by other protocol modules. We claim that only the four first ones (T, α, θ, and β) can be directly set, whereas R is indirectly calculated using the other ones. In this section, we define how the last parameter is calculated, besides providing a rationale of how T, α, α, and θ are set.

Setting β and T : When determining the values for T and β, it's important to strike a balance. We aim for T to be sufficiently long to allow validators a reasonable opportunity to sell their rewards. Estimating the point at which network access becomes valuable isn't a straightforward task. Nevertheless, for the sake of practicality, we assume a conservative estimate of *3 years* for this to occur.

The decay factor, β, is another crucial parameter that requires consideration. Our goal is to maintain a reasonably stable level of Mana for holders by the time T is reached. Once T is reached, we can assume that the network has entered a more stable phase. At time T, the holder's Mana will already be at approximately $(1 - \exp(-\beta T))$ of the maximum possible Mana in their hands.

Therefore, setting $T = 1/\beta$ ensures that by that point, they will have around 63% of the total potential Mana available in the system. Since Mana is also being spent, we consider this fraction to be large enough for good usability.

Calculating R : In the long term, our objective is to allocate a fraction θ of the total Mana to validators. Using the parameters T and β as previously defined, we can deduce that validators' Mana will eventually stabilize around $\frac{RT}{\Delta}\exp(-\beta T) = \frac{RT}{\Delta}\exp(-1)$, while token holders' Mana will stabilize at $\frac{S\gamma d_E}{1-\exp(-\beta\Delta)}$, where S represents the total supply of IOTA tokens and the remaining parameters are defined in Section A.1. Consequently, we can express this relationship as follows:

$$\theta = \frac{\frac{RT}{\Delta}\exp(-1)}{\frac{RT}{\Delta}\exp(-1) + \frac{S\gamma d_E}{1-\exp(-\beta\Delta)}} \Longrightarrow R = \exp(1)\frac{\theta}{1-\theta}\frac{S\gamma d_E \Delta}{T(1-\exp(-\beta\Delta))} \quad \text{(E.1)}$$

Figure E.1 illustrates the accumulated normalized Mana $(1-\theta)M(t)/(ST)$ over time for both token holders and validators, showcasing various selected values of θ.

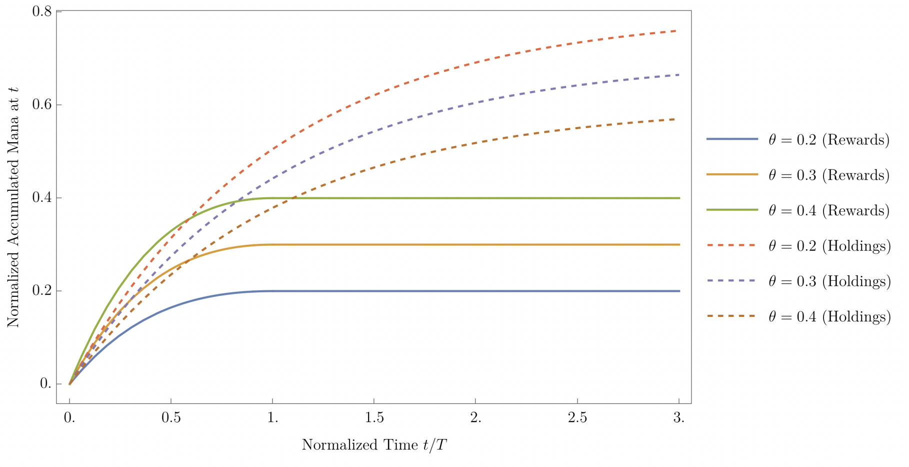

Figure E.1: Normalized Mana $(1-\theta)M(t)/(ST)$ as a function of t/T and different values of θ, for validators (full lines) and token holders (dashed lines)

Setting α and θ :

We aim to ensure a minimum level of profitability for both delegators and validators. Referring to Appendix C.4, we can recall that the expected rewards for these stakeholders in an epoch n, under the assumption of delegated stake equilibrium, can be described as follows:

$$\begin{cases} \frac{(\alpha+1)}{(\alpha+1)P+D} R(n)p, \text{ for validators} \\ \frac{1}{D+(1+\alpha)P} R(n)d, \text{ for delegators} \end{cases}$$

Here, P represents the total validator stake in the system, D denotes the total value delegated in the system, p represents a specific validator stake, and d represents a particular delegator stake. It's essential to note that these values will reach their minimum levels during the later stages of the network when $R(n) = c$. On the other hand, if these actors simply held their tokens without participating, their earnings per epoch would be as follows:

$$\begin{cases} \gamma d_E p, \text{ for validators} \\ \gamma d_E d, \text{ for delegators} \end{cases}$$

We can establish a "1:2:3" rule, indicating that a delegator should receive at least twice the amount of Mana they would receive if they didn't delegate, while a validator should receive at least three times the Mana they would receive if they didn't validate. This "1:2:3" rule can be expressed as:

$$\begin{cases} 2\gamma dp \leq \frac{(\alpha+1)}{(\alpha+1)P+D} cp, \text{ for validators} \\ \gamma dd \leq \frac{1}{D+(1+\alpha)P} cd, \text{ for delegators} \end{cases}$$

A potential solution to these inequalities is to set $\alpha = 1$ and θ in such a way that:

$$\gamma d \leq \frac{c}{D+2P}$$

In the worst-case scenario, when $P = S$, setting $\gamma d_E = \frac{c}{2S}$ should ensure that the above inequality always holds. Therefore, we have:

$$\gamma d_E = \frac{c}{2S} = \frac{\gamma d_E}{2} \frac{\theta}{1-\theta} \Longrightarrow \theta = \frac{2}{3}$$

Bibliography

[1] B. S. Frey and F. Oberholzer-Gee, “The cost of price incentives: An empirical analysis of motivation crowding-out,” *The American economic review*, vol. 87, no. 4, pp. 746–755, 1997.

[2] P. Cagan, “The monetary dynamics of hyperinflation,” *Studies in the Quantity Theory of Money*, 1956.

[3] P. Daian, S. Goldfeder, T. Kell, Y. Li, X. Zhao, I. Bentov, L. Breidenbach, and A. Juels, “Flash boys 2.0: Frontrunning, transaction reordering, and consensus instability in decentralized exchanges,” 2019.

[4] M. Carlsten, H. Kalodner, S. Weinberg, and A. Narayanan, “On the instability of bitcoin without the block reward,” 10 2016, pp. 154–167.

[5] A. Cullen, P. Ferraro, W. Sanders, L. Vigneri, and R. Shorten, “Access control for distributed ledgers in the internet of things: A networking approach,” *IEEE Internet of Things Journal*, vol. 9, no. 3, pp. 2277–2292, 2021.

[6] “Tip-0018,” https://github.com/iotaledger/tips/blob/main/tips/TIP-0018/tip-0018.md, accessed: 2023-10-09.

[7] “Tip-0040,” https://github.com/iotaledger/tips/blob/main/tips/TIP-0040/tip-0040.md, accessed: 2023-10-09.

[8] N. Dimitri, “Monetary dynamics with proof of stake,” *Frontiers in Blockchain*, vol. 4, 2021. [Online]. Available: https://www.frontiersin.org/articles/10.3389/fbloc.2021.443966

[9] “Tip-0039,” https://github.com/iotaledger/tips/blob/main/tips/TIP-0039/tip-0039.md, accessed: 2023-10-09.