Further thoughts on Work Your Stake (Draft)

Bill Roscoe, Bangdao Chen and Wang Lei

This paper analyses the practicality, economics and options of our WyS mining model. We are not, of course, economists.

We examine the options available to anyone designing a blockchain using WyS and the options there seemingly are to influencing its economy. The model was somewhat based on proof of work, but contains options not available there.

Overview of WyS

In WyS a miner invests coins for the right to mine specific blocks. However the system may choose to spread the efforts of individual miners over a number. We will divide each investment into two parts: a non refundable payment and a refundable deposit.

The role of the payment is to simulate the resources a miner spends in PoW. It also creates the funds for the economic levers we will discuss in the present paper.

The role of the deposit is to secure the correct behaviour of the miner concerned. If she does something demonstrably wrong she can be fined all or part of the deposit, which is otherwise paid back once the obligations derived from this piece of mining are over.

It also makes WyS a much better analogue of PoS. It is unlikely that any but the most gambling-oriented miner would tie up their whole assets in a series of mining tokens and put all of them at risk. However it will be wise for most to risk a small proportion in return for the expected gains. We do not want mining to be dominated by gamblers. By insisting that for every coin risked, 99 are placed as a deposit, we ensure that on-one can have more than 1% of their assets at risk. Of course 99 is an arbitrary choice here. We will see later that a plausible scenario makes the time an average miner has to wait for repayment acceptably small.

Clearly this *payment ratio* is an important parameter for us to choose. Broadly speaking, the closer it gets to zero the more like a pure proof of stake our model becomes, evidently the smaller the rewards for successful mining must be. Also a small payment ratio gives less potential economic levers and charitable donation.

The payment ratio is related to the incentive to miners to behave themselves, as the more deposit there is, the greater this incentive. However we also have the ability to tie multiple mining tokens (paid for together) in together regarding deposit, so the ratio is not the only factor.

WyS uses consensus techniques to control the rate of block formation, and the mining protocol is such that branching is virtually impossible in the sense that well behaved nodes will all choose the same option from a fork.

Rewards are given to miners in the form of minting new coins and gathering mining fees.

Parameters

Any particular episode of mining will tie up the fee and deposit involved for some time interval, with the deposit returned at the end. Clearly this deposit can instead be recycled back into more mining.

The parameters of mining include the fee ratio, this tie-in, which needs to encompass at least twice the maximum interval for a block to become immutable, and the number and nature of the reward coins. All of these can vary with time and the value of coins as valued externally.

A further parameter is what happens to the fees that miners pay. These can be given to good causes immediately, locked or not, or they can be cancelled. Or of course we can choose a mixture. Our assumption is that there will decisions made from time to time about what is done. This distribution might take account of how many new coins are being created as a mining fee since the delta in the total number of coins is this minus the number cancelled, and economic considerations might lead to a target for delta which is positive, negative or zero.

We also have to set technical parameters such as the inter-block time, block size and the way that the hash clock works.

Some of these parameters may vary with time:

1. Adjusting to the condition of the currency and the consequent wish to inflate or deflate it.
2. As the currency changes in value we might want to adjust some fees and rewards.
3. As demand for recording on the blockchain grows we might want to let blocks grow or reduce the inter-block time, subject our ability to cope with these securely.

Such adjustments can be programmed in from the start or they can be controlled (probably within pre-set limits) by some sort of vote. Similarly the code that controls some parameters might be replaced by a voting process. What we need to define at block 0 is the set of rules that determine to what extent and how they vary.

Commitment to mining

For the sake of argument suppose that our blockchain has 1,000,000,000 tokens actively owned (as opposed to unissued or destroyed). We contemplate that there might be two degrees of locking a coin: *locked* meaning that nothing can be done with it, and *semi-locked,* meaning it cannot be transferred but can be used for mining. If the latter we assume that any profits from mining with such coins would also be semi-locked.

The argument for the latter is if we want to make such locked coins "earn interest" in the same way as unlocked one. This an economic decision and one that affects parties' willingness to have coins locked. Of course the effect it has on the likely and worst case distribution of mining power may also affect the decision about having semi-locked coins.

Let us also assume that half of the 1bn coins are being used for mining, and that coins are tied up in mining during the creation of 200 blocks. Thus 2.5M coins are being committed per block in the form of payment or deposits. If a deposit plus fee is 100 coins and a typical commitment is of 100 such units this implies 250 mining commitments made per block. We envisage that for most miners a minimum number of units must be bid in each transaction, and the blockchain assigns these randomly to a range of blocks in the period, say 20-80 of the 200 long commitment interval. This can be implicit and defined by a standard randomisation algorithm that takes as an input the hash puzzle from the block two (say) ahead of the one containing the commitment. The point is that this randomisation cannot be anticipated at the point the commitment is made but at points when the mining is done the result is quickly available to everyone. In particular a commitment of N mining units can easily become N mining tokens, each assigned to a specific block. These tokens are implicit.

Clearly this creates an average of 25,000 mining tokens per block to be mined. These numbers seem to be entirely plausible, arbitrary though some of the choices were. There is much scope for adjusting them.

If we imagine that there might be one block created every 30 seconds, this means that a miner with one token in each mining competition can expect to win roughly once every 8 days.

Of course we still need to decide how much of each 100 coin unit is payment and how much is deposit. Some of this payment should be a transaction fee to the miner who puts the commitment into the blockchain, noting that we may need to overcome the incentivise such adoption to overcome the fact that a miner will not like competition. Such fees will be a zero sum game amongst miners, increasing both costs and prizes for wins. On the grounds that a somewhat random redistribution of assets may not be that popular with all, it may well be wise to keep the value fairly low subject to achieving the objective of having mining commitments included.

The rest of the fee is the amount, the system fee, the blockchain has access to for controlling the currency. Clearly overall, mining needs to be profitable for those who participate, so the sum of the minted coins and expected mining fees will generally speaking exceed the expected total of the system fees.

There needs to be debate and modelling to decide how large to make the minted quantities of new coins and the system fees, as well as how much of the total active currency should be motivated to mine. This can be controlled by the balance between the per-token fee and the expected gain for mining a block. We will discuss concepts of "interest" and "dividends" for blockchains in a later section, which will have a bearing on how mining rewards can be determined.

The main question over the mining parameters above might arise if coins became so valuable that many owners did not own sufficient to submit them in blocks. Such fluctuations have n value may in any case be undesirable.

An alternative is to have nodes place assets into a mining account, which the chain automatically re-invests in mining until, with sufficient notice, the owner removes them. Such an account would need to be updated from time.

Open and closed mining

In the first paper we discussed trust models and the fact that nodes might choose to perform mining and similar tasks in a closed (meaning that they do not reveal their identity or total holding) manner or with various degrees of openness.

The more a miner reveals himself, the more trustworthy we would expect related actions to be, because he has more to lose.

We are imagining a blockchain with KYC in which everyone holding coins has two unique identities:

A. Their true identity IDA as required by traditional banking.
B. A unique anonymous identity IDB which is tied bijectively to IDA. This is a random string of bits.

Neither is generally revealed and agents are free to use a variety of anonymous identities. However the blockchain will reveal which IDB is tied to a trading identity if the latter commits a severe enough crime, and the KYC system will reveal facts relating to IDAs in response to court orders.

We might want to split trading identities into service ones (which can perform mining and other similar tasks) and private ones. Service ones will be more easily linkable to the IDB.

We imagine, however, that agents will be free to do any of the following:

Prove the IDB or IDA and IDB associated with any trading identity they own.
Release all their service identities, in the sense that any service identity can be detected as belonging to the given IDB or not.

The idea here is that the extent of a node's mining power is exposed by the latter, and reputation is more easily damaged by bad behaviour if it is public. If we know a node's full mining capabilities, it is feasible to impose limitations on the influence it can have, for example having no more than 5% of the total or 10% of any one block.

We might want to give extra mining privileges to open miners, namely ones who reveal who they are. For example higher returns, later mining commitment or more focus in aiming mining. We will want to encourage open mining. Indeed we could investigate whether only open or semi-open (I.e. IDB disclosed) holdings might be used for mining, meaning if a holding could not be tied to at least the IDB identity of its owner then it cannot be used in this way. An identity could have both private and semi-open holdings.

Clearly we might have some sort of reputation system based on reviews for open miners or semi-open ones (who reveal IDB), but I would be cautious unless this was based on objective reasons, because of the possibility of gaming.

I can imagine having players with a widely known identity and reputation (e.g. a bank) or a status that makes them supposedly trustworthy (e.g. a notary public) having a special status if their certification was accepted by the community. We might seek, either for mining or some other sort diversity in the parties engaged: geographical, sector etc.

Economy: services, profits, interest and dividends

What is the best analogue for a blockchain in operation? We can think of it as a business which provides a wide variety of services:

1. Custody of assets
2. A platform on which businesses can develop their own applications, providing assurance and security to them.
3. A platform on which parties can develop registry services and databases for wide classes of asset.
4. Provides support for the creation, running and analysis of smart contracts.
5. Plus the existence of its own currency.

In every case there will be multiple users, and it makes sense where having a single party implementing the service is infeasible for efficiency or trust reasons, or the service is naturally linked to another service where blockchain is desirable. Here, not trusting a party might simply amount to not trusting that it will not be hacked.

The blockchain itself is a virtual entity: it exists though the cooperative work of a collection of parties who are all themselves users of it. These parties run hardware and software to maintain it. In essence the owners are those who own coins in it.

It is competing against non blockchain providers of the same or similar services, and must do so competitively in terms of price, efficiency and security. Or course it also competes against other blockchain providers of the same service. For it to make sense in the big wide world it must compete successfully in at least some services.

Coins gain intrinsic value when they are used to pay for goods and services. In our ideal system they pay for:

- Mining: parties pay to mine, and may pay miners to have things included in blocks.
- They are paid so that overall mining pays back more than the cost of doing it. However other users may ultimately pay for the service the blockchain provides to them by transaction fees.

- Services provided within or by the blockchain such as smart contract running, creation and verification, exchange services and running auctions. These would generally carry a payment from the user of the service to the providers.

- Payments for access to data or for operations on the client's data.

- Payment for other assets adopted (perhaps using separate tokens) into this blockchain.

Provided there is a reasonably sized economy comprising mining and services provided at a reasonably-constant real-world cost and, in some cases, real world competition, this should provide a measure of stability to the currency.

Potentially a blockchain can create more demand for its currency by insisting that some or all payments within it are made with that currency. However this might put some users off.

There is of course no reason why cryptocurrency cannot be lent from one party to another. We might want to place restrictions on the use of such loans for mining. Such loans would presumably carry interest, and we could regard mining rewards as another form of interest. A loan agreement will almost certainly be a smart contract.

An interesting question is what sort of banking operations will arise in and operate in blockchains for organising savings and loans. These might operate as individuals like HSBC or Barclays, through which funds flow. They might be banks that organise loans between users in return for a fee and probably offering some form of underwriting. Or they might be some sort of possibly automated collective with no owner as such.

In the conventional world, where a bank issues currency it can make a profit out of this which is called *seigniorage*. This means that the currency issued represents an interest-free loan to the issuer. The position with cryptocurrencies is more extreme in the sense that those issuing coins do not usually have to redeem them for anything. Their value needs to be maintained.

We need to make earning new coins attractive but not too attractive or it might be debased. Mining carries more work and responsibility than lending, and possibly more risk. In mining capital is tied up for a time: we would expect the expected benefits of mining to be a little more than the costs of doing it plus the interest that would have been earned by lending it plus appropriate compensation for the degree of risk.

The continuing value of such a currency is based on confidence, security, and on demand for it to spend on things.

It seems clear that the better the maintenance and security of the blockchain, the better the nature and distribution of miners and service providers is, and the more economical and desirable services are, the higher we can expect the currency value to be. There may thus be a tension between the level of charges levied for services and the desire to keep the currency valuable.

Currency stability

One of the most remarkable characteristics of existing cryptocurrencies is their instability when measured against normal currencies. As far as we can tell this is simply caused by speculation. After all aside from scarcity value and the difficulty value of minting new ones they seem to have little intrinsic value. We believe that the lasting value of new cryptocurrencies will come from the access they bring to excellent blockchain-based services. In essence a public blockchain on our model is a collective amongst the owners of its coins.

The coins can gain or lose value with perception of current and future demand and the consequent safety of investment in them. This value can be protected by IP that prevents others setting up a competitor with the same qualities: this applies both to the blockchain itself and also to apps that may run exclusively on it.

The value can also be affected by the blockchain’s own commitment to managing its own value (this being both a first order effect — the actions the blockchain takes — and a second order one — the knowledge of what will be done). Central banks control the stability of their currencies by controlling various factors such as interest rates and the money supply. A blockchain can potentially do both of these things: it can lend or borrow money at its chosen interest rate, and as we have already noted it can be programmed to release more or less money, or even reduce the money supply, through the WyS mining model.

I believe a cryptocurrency should have a commitment to stability. The provision of services of a rationally assessable value (e.g. determined by outside competition) should help with stabilisation, and a mining model tuned to this also. It might be a mistake to commit to one stabilisation structure for all time when the science and economics of blockchains will probably allow an improved model in future, so it may well be wise to allow for periodic updates by a suitable level of agreement.

Stability here does not necessarily mean being static. Perhaps there are targets, perhaps we try to limit the rate of change.