How to create a public blockchain ethically

Bill Roscoe

University College Oxford Blockchain Research Centre
OxHainan Research Centre

Introduction

Let’s suppose we have devised a set of protocols plus technology that will guarantee a great blockchain subject to reasonable assumptions about the community it builds and how it behaves.

But blockchains, and particularly public blockchains, are controversial because of

- The use by some of energy-wasting proof of work.
- The lack of regulation and possibility of holding and trading assets anonymously leads to their coins’ or tokens’ use in many criminal activities such as blackmail and money laundering; or simply to hide money from the authorities and taxes.
- There often seems little meaningful foundation for the value placed on the coins or tokens they run on and that have become so fashionable. Thus, they have been compared to pyramid or Ponzi schemes where (respectively openly or covertly) later investors’ money is mainly used to meet the expectations of earlier ones. Thus, enormous swings occur because of sentiment.
- The creators of public blockchains frequently award themselves blocks of tokens that, if they are fortunate, turn out to have gigantic value in the sense of the markets that are established in these tokens. It seems as though these people have printed vast piles of their own currency for themselves and persuaded (perhaps gullible) other people to pay real money for them.
- It can reasonably be argued that, if buying their tokens, people are buying “securities” which are little understood when they carry too large a risk.

On the other hand public blockchains offer a wonderful model of operating a decentralised system in a shared, democratic way with a very attractive way of avoiding having to trust individual parties or some sort of political or commercial oligarchy with common interests. They are potentially important to protect integrity in otherwise corrupt environments. The more centralised power is, the more those under its sway need to make sure that record keeping is invulnerable, just as they need an independent judiciary.

The sequentialisation that blockchains (in general) provide creates solutions to many of the technical issues arising from decentralisation. The “happened first” relationship needs to be unambiguous in many financial, legal and legislative contexts, and blockchains solve this. Of course, the more fairness is built into the way that blockchains do this, and the more independent it is of interested parties, the better.

Thus, perhaps unlike most pieces of non-military technology, public blockchains raise ethical arguments. This paper is about these and related issues of security, trust, and barriers to adoption.

The good properties of public blockchains largely depend on there being a large and diversified population of "good" nodes actively contributing to the creation of new blocks and the maintenance of integrity and security. It follows that banning all but wealthy and sophisticated participants to prevent inappropriate investments might be like throwing the baby out with the bathwater. I will later suggest less drastic measures, but the following question remains:

How much does the good behaviour of public blockchains depend on widespread and carefully managed greed?

A good node is one that follows the rules and does not attempt to attack the security of the blockchain by (for example) trying to create a separate branch, spending assets it does not have, or lying about the status of someone else's block or transaction. We can generalise the above question to

What properties of a mining population do we need to rely on? For example: greed, public spiritedness, and honesty. How important is the belief on the part of participants that they cannot benefit from cheating?

I note that blockchains which, like Bitcoin, are based on massive hash-power have *already* thrown the baby away by concentrating control, by economic necessity, in the hands of a few non-diverse, geographically clustered, mining giants[1]. This raises another issue: innovations can have negative consequences which are both unintended and unanticipated. The more novel, more popular, and less reviewed for risk something is, the more trouble of this sort can be anticipated.

For most of my adult life it never seemed to make sense to buy a computer because in a fast-developing world they were getting both better and cheaper so quickly. Of course, I did buy many but with the certain knowledge that I would be able to get a better one in a couple of years. Blockchains and our understanding of them are also developing rapidly from a technical and security, economic and hopefully ethical standpoint. The problem we have in deciding when to jump onto the public blockchain bandwagon is that, with assets and security tied into the continued community engagement with the one we join, we need either to be satisfied that one we adopt is good enough for all time, or having built in either the decision to move to another one at some juncture, or a secure way of opting for this.

It seems clear that the question of a blockchain being ethical or not comes down to the effects it has, or might have, on the individuals who participate in it, the financial system, society and the environment, plus the thoroughness and clarity with which risk has been explored and explained, and its future laid out. It is also related to how well it guards, and thus contributes to, the ethics of society.

[1] The existence of these mining giants creates a problem. Any well-established commercial community will always do its best to resist a disruptive entrant that threatens their profit model. So traditional banks will resist the entry of blockchain based banking and the existing hash rate mining giants will resist the emergence of public blockchains where their expensive equipment is useless. Turkeys do not vote for Christmas. The dangers of this type of situation are greater in unregulated markets.

By and large, society expects those who judge and govern it to be models of good behaviour. It always desires this.

I believe that the strongest ethical reason for public blockchains is as an integrity server: the strong cryptographic hash of anything placed in it guards against subsequent alteration for as long as the security of the hash lasts. There is consequently an ethical argument for running token economies if these are necessary to get a very wide buy-in to participating in the maintenance of such chains, not least because guarding the propriety of the token economy also guards that of the general integrity function.

Is proof of work unethical?

It is well understood that proof of work is a tool to ration the emergence of blocks: a means of ensuring an orderly consensus and is also an immensely powerful tool to prevent attackers creating fraudulent branches in existing blockchains.

As such it was a piece of genius to introduce the concept into blockchains. However, the very popularity of what it created was its own undoing since the energy it consumes is only necessary if one can think of no alternative. And then it would be a major liability in the balance sheet of using public blockchains at all.

Many people argue that proof of work can be justified by one of the following:

1. Doing useful work rather than pointless hashing.
2. Ensuring that only green energy (typically hydroelectric) is used.

The first of these is only neutral in the amount of work done if the useful work is not valuable to the party doing it, or else it will increase the attractiveness of mining which will further increase the amount of PoW done. The organisation of distributing and doing work for the collective good, *and getting it certified,* is probably beyond feasible, particularly if mining is to be truly democratised.

The second, of course, has opportunity cost in the energy market in general. If green energy is sunk into mining, it cannot be used for other things such as entering the regular energy grid, generating hydrogen from water or other ways of transforming it into storable green energy.

We conclude that PoW is negative because of its consumption of energy and because it concentrates mining power contrary to the basic democratic model that gives its existence so much credibility.

One might as well set up a coalition blockchain amongst those with lots of mining equipment on the basis of how much energy they give to good causes.

For me to consider it ethical, a version of PoW would need to preserve mining democracy rather than creating an oligarchy with big computers and cheap energy, and would have to prove to the world that it was worth it in a holistic sense: the benefits to the world of using whatever version of PoW outweigh its environmental cost.

For the rest of this paper I will assume that, ethically, we can do better than PoW.

Fiat currency versus cryptocurrency

Ultimately, currency is about confidence. To hold currency, you must be confident that you can buy real goods and services with it and that those to whom you have debts will accept it. The great majority of currencies we use are so-called fiat currencies, established and backed by nation states and their central banks. So, their value comes from a combination of legislation and the market's perception of the prospects for the respective state and the prospects for interest rates and monetary control. Countries have credit ratings just as companies and banks do. The economies of countries are hugely influenced by the positions of their currencies.

Fiat currencies have replaced linking nations' currencies to the price of gold: they are no longer linked to any particular commodity, though some countries do pin their currencies to a stronger one or even use another country's[2]. That is generally a sign of weakness.

A national fiat currency is both a symbol of the status of a country and a large lever it can use to control its economy. The foundation of the Euro replacing various national currencies was a very strong political statement and has caused significant economic tensions. The major levers that a central bank has to steady the underlying economy are interest rates and the money supply (the phenomenon of Quantitative Easing seen over the last decade is a big gun in money supply).

While this is little to do directly with ethics, it therefore seems extremely unlikely that nation states and central banks are, for the foreseeable future, going to allow cryptocurrencies *other than ones they control and so are not those of public blockchains* to take over from classical currencies. They are not going to allow the use of novel currencies to eat into their controls and ability to raise tax. Fairness of taxation, at least, is very close to an ethical imperative: look at the current discussions around the big tech companies that massage their profits into low tax jurisdictions, and around Donald Trump's tax affairs. It is widely agreed to be improper for an individual or company to over-aggressively shield themselves from taxation, since doing so is unfair on others and seemingly shirks from one's general responsibility to society.

From a simply pragmatic viewpoint, those setting up public blockchains therefore need to avoid presenting governments and central banks with easy ethical or economic grounds for excluding them. Again, on pragmatic grounds, it will probably be better to make cryptocurrencies both behave like and be branded like conventional securities rather than currencies.

There is of course nothing to stop a blockchain, public or otherwise, from dealing and banking in fiat currencies, where necessary with regulatory approval. I reiterate the basic point that blockchains must contribute to and participate in the existing economy far more than overturn it.

[2] When I went to Ecuador, I was surprised to find that the local currency is the US Dollar. See https://www.usnews.com/opinion/blogs/eric-schnurer/2014/05/02/why-ecuador-and-other-states-dont-use-their-own-money

Why are integrity servers necessary?
Records are now stored in digital form. In most spheres, paper records are now incidental or history. Of course, data and consequently records, increasingly move to the cloud, and inevitably get replicated many times for convenience and security. From the point of view of privacy these facts cause big problems, but that is not directly our concern, particularly since storing a record of data offering bullet-proof integrity need not give it away thanks to cryptographic hashing.

The scenario that blockchain guards against is that of arguments about what the true records are. Once something is firmly established on a blockchain it is immutable provided that the identity of the blockchain itself is agreed and there are rock-solid mechanisms to prevent more than one version of any block (other than very recent ones) existing in a form that anyone who performed reasonably simple checks could agree was genuine.

It follows that both the structure of the chain and the mechanism for selecting new blocks are both very important for this. Understanding the motivations of the population, who are behind block creation and approval is immensely important to achieving the necessary common belief and knowledge, is similarly important. Many people have more trust in masses motivated by the imperative of preserving their assets than they do in governments or industrial associations, hence the appeal of public chains. A diverse population is seen as being less likely to have a common motivation to cheat or to conspire.

Other record systems can prove their own integrity by storing hashes regularly in an integrity server like a public blockchain. Such a system might be a blockchain that is not universally trusted or any system which is thought to be a potential target for attack or bringing down. Or indeed a blockchain that wants to increase confidence in itself can publish the hashes of regular blocks in our public chains. A chain can similarly publish its own hashes in immutable places such as other public chains or even popular newspapers (as does the Estonian record system). The objective is to make changing established values impossible and widely understood to be impossible.

Having an integrity system is invaluable for many parts of life, governance, and business such as

1. Many aspects of public life: essentially ensuring old records are not altered.
2. Auction and tender processes.
3. Financial records.
4. Any system that has to be audited, such as accounts, risk management and procurement.
5. Running elections.
6. The audit process itself.
7. IP and patenting.
8. Clinical trials management.
9. Medical records.

Just because some records are proved to exist by such an integrity system does not prove they can be recovered. Cryptographic hashes cannot be inverted. However, they can decide between two differing versions of a file or allow the rules to place an obligation on a party to exhibit the pre-image of a hash it has signed in an integrity system.

The beneficiaries of an integrity service are those who make their data immutable — because they want to or are required to — those who rely directly on the data knowing it to be what it purports to be, and society as a whole. These parties can either pay for the service in money by participating in the storage and processing that are required.

Is there a role for trusted enclaves?

Two technologies are currently transforming the options for creating trust in decentralised systems: blockchains and trusted enclaves (or TEEs). Blockchains provide order and long-term integrity. TEEs provide privacy preserving computation, and also the assurance that the correct computations are done on the correct data; both even when you do not trust the operator of the computing platform.

I have no doubt that this style of security adds enormously to the capability and usability of blockchain-based systems. For example, identity verification and checking naturally depends on privacy-preserving computation, as does automated regulation which requires access to non-public information such as individual holdings or patterns of trade. Such use is separate from the basic structure and security of the blockchain.

The question we address here is whether such technology can play a part in the security of the blockchain itself. Do we need less participation in mining if blocks are constructed inside TEEs? Or less verification that rules are followed if everything is checked once in a TEE? Can we trust the creation of (pseudo-) random numbers from TEEs? How about the long-term security of blockchains thus created?

I am assuming here that the basic structure of the blockchain, and the strength of cryptographic primitives used, are not affected by this. And that there is a protocol for maintaining the chain in the absence of trusted computing.

A potential issue with TEEs and privacy is that, for either technical or political reasons, some sort of TEE might not be trusted by everyone[3]. To simplify this, suppose that the world is divided into two technical camps, the orange and the green, who are not willing to trust each other's enclaves, respectively mandating the use of their own. (One of the rationales behind such mistrust might be different crypto standards.) The rest of the world just looks at this in frustration.

This creates a big issue for privacy, but less so for rectitude of action, since we would just need TEEs of both colours to agree on verification, on who gets to mine the next block. What this suggests to me is the following

1. The structure of the core blockchain: blocks linked by hashes and possibly hooks should be created independent of the use or otherwise of trusted computing.
2. There is the potential for less replication of housekeeping and calculations if there are a set E of enclave types such that a very high percentage of participants

[3] The same can apply to cryptographic primitives such as those at the core of blockchain. Just as with the discussion here of TEEs, an effort should be made to find a way of satisfying all.

(including consumers of the integrity service) trust at least one of them, and each such calculation is done by a representative of each member of E, *and they all agree. If they fail to agree or no such E exists, the protocol reverts to the standard one.*

3. As with any other blockchain, the evidence must be in place that the whole chain was created with all the necessary agreements and checks, and any of these that would not be up to the standards of today (of cryptography or TEEs, for example), were laid down at times when they could not have been forged and have not been changed since.

Discussion

Imagine we are proposing to build a public blockchain.

Given the obvious questions “Why does the world need another public blockchain?” and “Are you not profiteering by awarding yourself many tokens?”, how are we to answer? Similarly, how can you parry questions about instability, risk, and criminal use.

I do not think that the first of these is really the right question. I think that is “Why do you think you have finally designed one that everyone can have confidence in?” because a convincing answer to that will be justification enough, and the world certainly has enough blockchains that fail this test.

The other issues: greed, risk, and instability are really related to the confidence one too. How can one be confident in a blockchain whose construction was motivated by unconstrained greed, was not believed or preferably proved, to have a stable token economy or did not have risks such as failures of integrity, insecurity, or lack of progress.

The criminality issues require that criminal use is both discouraged and discoverable. It should not be ethical to create a blockchain which does not offer mechanisms — which I expect would involve the combination of KYC, automated regulation and anonymity being conditional on observance of the rules. The test of such a mechanism is that it must be simple enough to understand and incorruptible by either criminals or those with an improper interest in breaking confidentiality.

A cynic would regard many public blockchains as closer to Ponzi schemes than genuine investments. Such a chain could surely not be created ethically. We therefore assert that a public blockchain’s value should be predominantly underpinned, like that of a listed company, by the assets it owns including IP, its trading activities, and the services it sells. It should not be underpinned by the sheer cost in energy or equipment of mining blocks, or by speculation.

We must realise that where tokens are to be underpinned by assets and income generation, tokens that are initially granted (free or at a discount) to founders, investors, guarantors and charitable causes must be justified in the underlying economic model. They must be justified by the economic value of the blockchain for income generation or (in the case of investors and especially guarantors) by assets they place at the blockchain’s disposal. (Here, a guarantor means someone who, like a Lloyd's of London "name", puts their assets at risk to underwrite the commitments and stability of the blockchain.)

In other words, I do not believe that blockchain tokens are a new financial universe where the laws of economics are different. To create one that operates outside the usual laws is unethical without a convincing proof that I am wrong, just as it is unethical to create one that puts up insufficient defences against criminal activities such as money laundering. One argument for my view is that inevitably tokens are traded for conventional money and/or assets with defined value outside the chain, which links the two economies closely.

To quote Gordon Gekko (in the film Wall Street), *Greed is Good,* or at least it is where it is a motivating force to achieve growth, efficiency or in our case a reliable and secure blockchain. But that is the same instinct that inspires people to invest in stocks and shares rather than hiding their money under the mattress or become a name at Lloyds.

If one follows the model I advocate above, it would provide a major defence against accusations of greed, for then founder and investor tokens in a blockchain would be little different from founder and investor shares in a company. They can then justify themselves by *precedent.*

A strong argument for having a large number of tokens in the hands of generally trusted and motivated parties at the bootstrapping of a chain, is the need for a smooth and reliable operation during its initial period. Everyone we want to participate should be confident that these parties (at least most of them) will be active and honest. Often such tokens are "locked" for a long period, meaning they cannot be traded for a long period. That may be part of a solution, as might allowing them to only be used for a specific purpose such as participating in a steady way in the WyS mining model.

My distinguished mentor Sir Tony Hoare, in describing best practice for the creation of safety-critical systems, has often related that engineers who designed ship launches were, after some disasters occasioned by the huge waves these events can generate, forced to take the risk themselves by standing on the opposite bank of the river. The implication of this is that those responsible for systems should share in the (negative) risks as well as the (positive) profits of what they create, because this concentrates their minds on what other stakeholders care about. I have always thought that one of the very worst incentive structures is that adopted by some hedge funds whose managers charge a fairly significant percentage of any profits they make, but make their customers entirely responsible for any losses. An excellent strategy for such a manager would be to set up two funds that make opposite bets, say respectively buying buy and put options in the same securities, equivalent to betting on opposite sides in a football match.

If you award yourself free tokens in a new public blockchain, it is a one-way bet. The worst that can happen is that you make no money.

Concepts such as *conflict of interest* and *insider trading* are well understood in the traditional financial world. Most of these might arise in traditional ways for assets traded on blockchains and it is reasonable to expect that regulation systems for such trading are cognisant of them. Let us now contemplate whether such issues apply to the mechanics of the blockchain itself. Since the mechanisms and trading on a public blockchain are themselves respectively open and public subject to a degree of anonymity, the opportunities for insider trading are limited, but this should still be thought about.

Conflicts of interest arise when someone owes allegiance to multiple parties with distinct interests: perhaps they claim to be backing A when in fact they are backing B. One cannot eliminate this possibility in people who trade or conduct businesses on a blockchain, and therefore cannot eliminate it altogether in parties who conduct the blockchain's own business such as mining.

Similar issues of unfair advantage and divergence of intervention arise between founder token holders and long-term ones, and where there is a distinction, between miners and non-miners. These can manifest themselves in the rules of the blockchain, particularly if these are complex and not all stakeholders understand them. They can be manifest in the choices the participants make when different "good" miners have different priorities. Or they can affect the votes stakeholders make in cases where rules can change. The following can counter these

1. There is much to be said for obliging miners, particularly founders, to make their decisions about the chain and do their block verification, regulation and mining using standard well-understood programs designed to promote the health of the chain and the behaviour advertised as its objectives. Such obligations can be reinforced with penalties on non-compliance.
2. The design and behaviour of the token model and how it relates to the overall behaviour and assets tied to the blockchain must be clear, designed for stability, and thoroughly reviewed.
3. Concentrating mining power in the hands of those with specialist equipment — a consequence of proof of work or other models requiring non-standard computers for economic participation — is very unhealthy and should be avoided at all costs.
4. The stakeholders in a blockchain must be constrained about the rules they can change.

The security of a public blockchain inevitably depends on how sufficient the work being done by "good" participants, namely ones who follow the rules. The assumed threat model must be public and difficult to challenge in the direction of it being insufficiently conservative. The behaviour of the blockchain under this model must be thoroughly analysed by mathematics and simulation. The ability of bad nodes to affect choices made in a blockchain should be considered: this seems inevitable: for example if the good nodes are evenly (in voting power) split on some issue, it will give the bad nodes real power.

If the main purpose of setting up a public blockchain is to establish a beacon of integrity, then it is far preferable for it to be so immediately. I think that a way to achieve this is to set it up with holdings in the hands of a very trustworthy population who are obliged to maintain it using verified software in a public way. Other parties can build up holdings by guaranteeing the value of the cryptocurrency in a carefully defined way, by investing where regulators allow, and by performing maintenance services. This would happen in a controlled way, so the system would evolve from a coalition blockchain designed to be very trustworthy now, to a strong public one.

I imagine that most public blockchains will aim to survive for the indefinite future, but an ethical design should contemplate what happens to it, the assets stored in it and the tokens it generates under all situations that can be anticipated. Such situations might be technical, such cryptography no longer being considered secure or some successful attack being detected. They might be community-based, such as inactivity or insufficient diversity, or they might be regulatory, either due to actions of external powers or because some internal limit has been breached such as capital reserves. Some reactions might be automatic, such as (perhaps partial) suspension or careful move to stronger cryptography. Others might trigger the expectation of a democratic choice according to predefined rules.

Criteria

In summary, I believe that to be considered ethical, a public blockchain should satisfy the following:

1. Its design must be made public and subjected to detailed examination of, for example, its security and economic model, before launch. It must provide a generally accessible integrity service subject to a moderate fee if any.
2. It must not waste resources or equipment beyond what one might reasonably expect from a decentralised system. In particular, the mining model should not be based on buying resources such as energy in the real world.
3. On the assumption that tokens are tradable for “real” money, assets or services, the value of tokens should be anchored against assets or services, or prospects of these, as they would be with a conventional company.
4. It must offer protection against use for money laundering or other criminal use of assets at least comparable with regular banking. This applies to all assets that are or are tradable to real world assets.
5. The initial distribution of tokens should be consistent with the economic model and should help the stability of the chain. All significant beneficial founder holdings should be public. This should possibly apply to *all* significant beneficial holdings: these can probably be identifiable by KYC mechanisms necessary to counter criminality.
6. Charities, with a strong preference for ones independent of other stakeholders, should receive capital and income from the blockchain that have not been earned as independently judged.
7. It should clearly define the options for its own future.
8. The best judges of whether a reward model is fair or not are not those who will benefit from it.

The whole should be thoroughly and independently reviewed.

Conclusions

Blockchains have a lot to offer to society and international trust. If one cannot set up a coalition that will be permanently trusted by all, then a widely used public blockchain where parties hold and generate assets has much going for it. It is necessary to strongly align:

- The interests of the participants who hold assets in keeping the integrity intact.
- The interest of a diverse and, where possible, public spirited community in building a chain with total integrity.

For it to be ethical

- Resource wastage must be eliminated, where *wastage* means that the blockchain must not encourage consumption over and above the basic needs of transacting, holding and security replication of data, and doing an amount of housekeeping of the sort done in typical distributed systems. *This probably means a version of Proof of Stake such as my own Work your Stake model (WyS).*[4]
- It must make wide democratic participation possible and make dominance by a single or related interest groups as hard as possible.
- It must be at least as resistant to criminal usage as the standard mechanisms of society. Consequently, absolute anonymity of those holding assets cannot apply.
- Its model for the value and income generating capability tokens should be understandable and should be designed to be stable with well understood risk.
- Its future should be carefully planned to preserve security and these properties.

Further reading

I have written a number of papers on the construction of blockchains (especially public) over the past three years. These can also be found on this website. I recommend the paper on Green Mining which describes the Work your Stake model backed up by the paper *Taking the work out of blockchain mining:* my anti-forking structure called hooks to help dispense with PoW. Other papers cover secure random number generation, exchange etc.

[4] See *The greening of blockchain mining* also available from www.tbtl.com.