Cryptoeconomics

a section of

By Eric Voskuil

Libbitcoin.org

TABLE OF CONTENTS

Libbitcoin: Cryptoeconomics

Cryptoeconomics

This section aims to provide clear and useful high level explanations using terminology based on concise formal definitions. Common use terms are retained to the extent that a term is not misleading or cryptic. See the “Cryptoeconomics” side menu for a full list of topics.

Value Proposition

The value of Bitcoin over its alternatives derives directly from removing the state from control over both monetary supply and transaction censorship. Advantages include freedom from seigniorage, foreign exchange controls, and financial surveillance. These allow the money to be transferred to any person, in any place, at any time, without need for third party permission.

These advantages represent cost reduction through the avoidance of tax. Seigniorage is directly a tax, while foreign exchange controls limit its evasion. The state itself often claims political independence as an objective in the interest of limiting this taxing power. Financial surveillance limits tax evasion more generally. **While Bitcoin cannot eliminate tax, or even necessarily reduce total takings, it represents a change in the nature of taxation.** In any case, for those who consider the state a social good, the option to voluntarily fund it remains.

It would be an error to assume these advantages flow from the existence of a more efficient technology than employed by other (state-controlled) monies. The technology is far less efficient, yet it helps people resist state controls. It is this resistance that provides the value.

Axiom of Resistance

In modern logic an axiom is a premise, it cannot be proven. It is a starting assumption against which other things may be proven. For example, in Euclidean geometry one cannot prove that parallel lines never meet. It simply defines the particular geometry.

Proving statements about Bitcoin requires reliance on other systems, specifically mathematics, probability or praxeology, and therefore the assumptions upon which they rely. However Bitcoin also relies on an axiom not found in these systems. Satoshi alludes to this in an early statement:

> You will not find a solution to political problems in cryptography.
>
> Yes, but we can win a major battle in the arms race and gain a new territory of freedom for several years.
>
> Governments are good at cutting off the heads of a centrally controlled networks like Napster, but pure P2P networks like Gnutella and Tor seem to be holding their own.

In other words there is an assumption that it is *possible* for a system to resist state control. This is not accepted as a fact but deemed to be a reasonable assumption, due to the behavior of similar systems, on which to base the system.

One who does not accept the axiom of resistance is contemplating an entirely different system than Bitcoin. If one assumes it is *not possible* for a system to resist state controls, conclusions do not make sense in the context of Bitcoin; just as conclusions in spherical geometry contradict Euclidean. For example, without the axiom, how can Bitcoin be trustless and a sound money? In this case, does sound money not require trust that the state will not subvert its foundations? This leads one to make obvious errors in an attempt to rationalize the conflict.

It is common for people to refer cynically to a Bitcoin-like system that omits the resistance axiom as just another "PayPal", a designation not without merit. Confinity originally attempted to create a system with a similar value proposition to Bitcoin. Having failed to do so it discarded the axiom, building the PayPal we know today.

Custodial Risk Principle

When a contract represents an asset, the contract is a claim against the asset's custodian. This claim is often called a security, with the intended implication that the claim is "secured" against custodial failure to exchange the asset under the terms of the contract. The monetary value of the security is that of the underlying asset minus the exchange and claim enforcement costs.

Custodial risk is a central aspect of any money. The usefulness of a money is limited by the reliability of its custodian. Being human, the reliability of a custodian cannot be assured. In the case of state money, the single custodian is the state. As shown in Reservation Principle state money exists for the purpose of accumulating a reserve. This provides a benefit to the state only because its custodial role can be abrogated both through liquidation of the reserve and issuance of fraudulent securities. In other words, custodial default is the reason for state money.

The monetary value of a unit of Bitcoin is strictly a function of what it can acquire in trade. If no merchant accepts it, a unit is not useful to its owner. Bitcoin is non-custodial, but in the interest of establishing a general principle, one may consider the set of all merchants the collective Bitcoin custodian. As such the custodial risk is spread across the economy.

In the case of Bitcoin, merchants offer their own property in exchange for the money. As such their is no implied securitization of the property. A merchant can cease to accept any money, which reduces the utility of the money. This can be considered a custodial risk, but not a default as the merchant has accepted no obligation to trade for the money. As shown in Fragmentation Principle, changing merchant acceptance is the nature of a split.

As shown in Blockchain Fallacy, "blockchain technology" can offer no defense against custodial default. A "tokenized" asset is a security. The opportunity for fraud or theft by the custodian, either directly or as compelled by the state, is not reduced. **Just as with commodity monies, such as gold, the custodial risk reduction afforded by Bitcoin is not a consequence of technology or contractual obligation, but the size of its economy.** Ironically it is the "security" that is insecure.

Dedicated Cost Principle

The necessary costs associated with mining set a floor for the cost of censorship. A censor must expend greater capital than the necessary expenditures of all other miners combined to maintain the strongest chain, and thereby the censorship regime.

Unnecessary costs incurred by miners contribute nothing to defend against the censor. Such costs constitute true waste, representing nothing more than a given miner's inefficiency. For example, it does not contribute to security if a miner with

misconfigured machines expends a great deal of energy while being unable to win a reward due to the misconfiguration. Any cost that is not strictly required for the optimal generation of hash power is not a necessary cost. A misconfiguration cost can be eliminated by fixing the configuration.

There is a theory that proof-of-work can be made more energy efficient by introducing non-dedicated costs to the mining function. One such example is the discovery of prime numbers. The reason to incorporate such costs is that the resulting discoveries have presumed marketable value. If not there would objectively be no value in the incorporation, invalidating the theory.

Any cost dedicated to the production of independently-marketable value can be offset by selling that byproduct. By analogy, brewers can sell their grain byproducts to farmers. This improves their efficiency by eliminating an unnecessary cost. So to the extent that the resulting byproduct is valuable, it's production does not incur a net cost.

Yet actual (net) costs must rise to the level of rewards, as a consequence of competition. Therefore the same result would be achieved by basic PoW consuming the full reward value and independent energy-consuming operations generating the marketable products. As such the theory is invalid.

Merged mining is typically implemented to resolve the problem of bootstrapping a new coin past the vulnerable stage of low hash rate. This design fails to recognize that hash rate not dedicated to the new coin does not contribute to its security. As the full cost of the hash rate can be recouped by selling it on one chain, there is no cost to censor the other merge-mined chain(s).

Other Means Principle

Bitcoin is an act of resistance, an attempt to "gain a new territory of freedom." Freedom contracts through the constant pressure of compulsory financing of the state. It is typical that freedom is expanded through bloodshed, with the specific objective of reducing state power. Bitcoin cannot eliminate the need for personal risk in achieving this objective. However, through risk sharing it can potentially reduce the inflation tax without spilling blood. This will not eliminate tax generally, however it will reduce state power by making tax significantly more visible.

This conflict between state and individuals for control of money will pass through up to four phases anticipated by the Bitcoin security model. These may overlap and vary regionally, but are each clearly identifiable.

1. Honeymoon
2. Black Market
3. Competition
4. Surrender

The honeymoon phase is characterized by a desire of state agencies to retain regulatory control over the movement of money and securities. To this end pressure is applied at points of aggregation. As pressure on pooled miners and centralized merchants increases, cost rises and utility drops. The money then necessarily becomes more distributed to avoid these expenses.

As it becomes apparent that controls on points of aggregation are insufficient enforcement, and the awareness surfaces that seigniorage is at risk, transaction in and mining of Bitcoin is outlawed. As states collaborate to protect their monies, this may become a global "War on Bitcoin". This may coincide with adoption of an official new money, i.e. Fedcoin. The objective would be to appear to embrace a "safer" money than Bitcoin while retaining the seigniorage and surveillance advantages of electronic fiat.

Assuming sufficient resistance, Bitcoin persists independent of Fedcoin as a black market money. At this point the state concludes that the only effective tactic is to compete as a miner. Given that mining is necessarily anonymous, there is no way for the economy to prevent state participation in mining. Thus Bitcoin enters the competitive phase, with the state attempting a perpetual 51% attack.

Apart from ongoing black market phase enforcement, the competitive phase is characterized by a peaceful hash power battle between the state and individuals. The state operates at a loss due to the rejection of censored transactions. This loss is offset by tax revenue. Fee pressure on censored transactions rises until the state mining tax subsidy is offset by this fee level. **At this point taxes and censored transaction fees both rise until one side of the conflict surrenders.** In this manner Bitcoin can potentially win a war by other means. It cannot be assumed that this surrender will be perpetual. As implied by the Threat Level Paradox, the money is likely to drift into previous phases as the threat diminishes.

Patent Resistance Principle

Unlike copyright, patent is an anti-market force. A true copyright is a contractual agreement between buyer and seller, where patent is exclusively a state grant of monopoly. The patent is not an "attack" by the patent-holder, it is a distortion pooling pressure created by the state.

The process of mining is highly competitive. Monopoly protection in the use of efficient mining algorithms is a strong anti-market pooling pressure. Bitcoin is secured by people resisting anti-market forces. Resistance incurs greater risk when the miner is highly pooled and/or non-anonymous.

If people do not resist such forces there is no security in the money. As the threat level increases the consequence of patent violation becomes no more of a risk than manufacturing miners and mining. As such the impact of patents is inconsequential as it pertains to the security of the money.

Risk Sharing Principle

Bitcoin is not secured by blockchains, hash power, validation, decentralization, cryptography, open source or game theory; it is secured by people.

Technology is never the root of system security. Technology is a tool to help people secure what they value. Security requires people to act. A server cannot be secured by a firewall if there is no lock on the door to the server room, and a lock cannot secure the server room without a guard to monitor the door, and a guard cannot secure the door without risk of personal harm.

Bitcoin is no different, it is secured by people who place themselves at personal risk. Sharing this risk with other people is the purpose of decentralization. A centralized system requires one person to shoulder all of its risk. A decentralized system divides risk among individuals who comprise system security. Those who do not understand the value of decentralization most likely do not understand the necessary role of people in security.

Bitcoin allows people to share the personal risk of accepting and mining coin. It is only the willingness and ability of these people to resist that can prevent coercion of their nodes and co-option of their mines, and this is what actually secures Bitcoin. If people do not accept these risks there is no effective security in the money. If a great many people do so individual risk is minimized. Bitcoin is a tool, not magic.

Reservation Principle

Fiat money exists to collect seigniorage, which is a sovereign tax on the holding of the money. The tax is levied by manufacture of units of the money denominated above its production cost in the money (i.e. counterfeit). The objective of the tax is to accumulate political power, including reserve money.

A reserve is the accumulation of a tax hoard. When the sovereign state collects seigniorage it must do so in another money not subject to the tax. Otherwise the hoard itself would be diminished by the tax. **The existence of a reserve presupposes fiat, as otherwise there is no reason for two monies.** The reserve retains its value despite the purposeful debasement of fiat, and is eventually held primarily by the sovereign. The current (third) United States Federal Reserve was instituted in 1913.

Initially legal tender laws and taxes create demand for fiat, which must be purchased with the reserve money. Such laws were applied broadly in the United States in 1862. As reserves accumulate and fiat becomes widely accepted, the sovereign may seize what remains without immediate risk to future tax on other economic activity. This is done in part by abandoning fiat-reserve convertibility and also by direct confiscation. In the United States domestic confiscation and non-convertibility occurred in 1933 and non-convertibility was applied to nation-state fiat holders in 1971.

It is common for the fiat money of one sovereign state to become the reserve of another. This allows the issuer of fiat held in reserve by others to transitively tax holders of lower tier fiat. This tax is begrudgingly accepted by other sovereigns in exchange for the ability to purchase products available primarily in the global reserve. Control over access to the clearing network) of the global reserve is often used as an economic weapon. The US Dollar and to a lesser extent the Euro make up the vast majority of foreign exchange reserves).

Fiat is typically "paper" and/or metal currency used alongside account-based money. Since issuance is the source of taxation, money accounts, now largely electronic, are brought under control of the sovereign. As such, fiat issuance is primarily electronic yet euphemistically referred to as "printing" money. This occurred in the United States in 1913, coincident with creation of the reserve system. With the Federal Reserve Act the state gained physical control over a national reserve, the power to issue notes against it, and regulatory control over money accounts.

Electronic money is harder for third parties to counterfeit, and gives the issuer visibility into all financial transactions. This visibility is another tax on users of the money, as it compels increased compliance with independent takings. Avoidance of these takings is referred to as tax evasion or money laundering by the state, depending on context. As such we are witnessing the gradual disappearance of tangible paper and metal money.

The only brake on the counterfeit of fiat is the inability or unwillingness of people to pay the tax. This results in abandonment of the money despite legal tender laws. As this takes hold the sovereign typically enacts foreign exchange controls to raise the cost of using alternate monies, though such controls can have high enforcement costs. Eventually fiat goes completely out of use and the reserve money is utilized by both individuals and the sovereign.

With the inability to tax at sufficient levels comes consumption of the reserve, which was ultimately insurance against the failure of the fiat. This is a currency crisis, caused by a sovereign debt crisis, likely resulting in a political crisis. With the establishment of a new money and/or sovereign the cycle continues.

As most global reserve is in fiat, a failure in the reserved fiat will propagate through all state monies. Such a failure can be mitigated neither by transition to another sovereign's fiat nor gold and/or silver. Trade is largely electronic and metals do not travel over computer networks. In other words, Bitcoin is the only money that can resolve a currency crisis in the modern global reserve without a catastrophic global contraction of trade.

As shown in Reserve Currency Fallacy, if Bitcoin was to become primarily a reserve it would cease to be Bitcoin. However it will always be possible to continue it apart from the state, and people may better defend it after learning from past failures.

Scalability Principle

Scalability is the proportional increase in some aspect of performance as more hardware is employed. Bitcoin transaction throughput is perfectly non-scalable as no amount of hardware increases it.

The block size limit consensus rule establishes the arbitrary trade-off between utility and system security. Increased block size marginally increases transaction throughput and therefore the resource cost of transaction validation (i.e. processing, storage, and bandwidth). As the cost of validation increases, economic security is adversely impacted by increased centralization risk. As the trade-off is arbitrary, there is no ideal size.

At any block size the system remains non-scalable due to the necessity of confirmation finality. A finite set of transactions must be selected, which implies that others may be excluded. This exclusion is financially motivated by the opportunity cost of not utilizing deployed mining capital, and is the manifestation of non-scalability. This inherent scarcity necessitates a competitive market for confirmation, and finances it in proportion to demand for the money.

Effective transaction carrying capacity, and therefore utility, can be increased by layering. This represents a *local* and *time-limited* security compromise, in contrast to the *system-wide* and *persistent* security compromise of increasing block size. Either compromise lowers but does not eliminate the utility threshold, which implies the stability property is preserved.

Therefore stability and non-scalability exist at any block size and level of layering.

Consolidation Principle

The need to exchange from one coin in order to trade with merchants of another is a cost. This cost must be non-zero even if automated, as it must consume space and/or time. As such one coin is always "better" (higher utility) than two.

We can reasonably assume that two distinct moneys cannot perpetually have identical utility. Thiers' Law discusses the consequences of better money in the absence of state controls. From this we necessarily conclude that **the better of the two monies will eventually replace the other** in the absence of state controls. As this occurs utility accrues to the surviving coin in the reverse of the manner detailed in Fragmentation Principle.

This does not imply that new coins cannot be created or exist over a significant amount of time. It implies that there is a market pressure toward a single coin. A better money in one situation may not be a better or even useful money in another. For example, gold is not a useful money for electronic transfer and bitcoin is not very useful without a network. One money replaces another in the scenarios for which the former is better.

Fragmentation Principle

The utility of a money derives directly from its ability to facilitate trade, in contrast to barter. If it is not accepted by *any* merchant then objectively it has no monetary usefulness. The more goods and services (including consideration of location) that can be purchased with a money at any given time, the more likely it is that the money has greater utility to any given person.

A split implies that zero or more merchants have stopped accepting the original coin and that zero or more have started accepting the split coin. A "clean" split is a hypothetical scenario in which there is no overlap in merchant acceptance of the two coins, and no change in the set of merchants. A clean split produces two economies from the original set of merchants.

If we assume that the coins are identical apart from the fact of the split, the Consolidation Principle implies that the utility of the combined coins is the same as the utility of the original less the exchange cost. The scenario can be expanded to include merchant overlap. This has no effect on coin utility, as it only shifts the incidence of the exchange cost from buyer to seller.

An increase or decrease in the number of merchants accepting either of the coins is a net gain or loss of combined utility respectively, as it implies the removal or addition of a third-currency exchange cost. In other words the effect is proportional to each of the coins in the split. This factor relates to the particulars of a given split, not to splitting in general.

Therefore a split produces both a shift and reduction of utility, in proportion to the relative sizes of the resulting economies. The Network Effect Fallacy explains why the reduction is not quadratic in nature, as sometimes assumed.

While it may appear that in the shift someone has "taken" value from the original coin, that value has actually "left" to form the split coin. In other words merchants are masters of the value that they provide to a money. Owners have *independent* control over the purchasing power of one unit, based on their level of hoarding. However this affects unit price, not utility.

Upon the split an original unit becomes two units, each with decreased and proportional utility in relation to the original. With bidirectional mandatory replay protection each can be spent at no additional cost. Otherwise the need for self-protection discounts units of the unprotected chain(s).

This analysis is applicable to new coins as well. The difference in the case of a new coin is that original (other) coin units are not spendable on the new chain. As such the new coin is faced with the difficulty of allocating units, which requires work and therefore time. Splits bootstrap this process by subdividing the utility of an existing chain, to the extent its merchants are willing to do so.

Permissionless Principle

Bitcoin is designed to operate without permission from any authority. Its value proposition is entirely based on this property.

A market can be divided into permissioned and permissionless from the perspective of the state. For ease of reference the former is often referred to as “white market” and the latter “black market”. White market trade, by definition, requires permission, and black market does not.

As a simple matter of definition, Bitcoin operations cannot be both white market and permissionless. Any person operating in the white market requires permission to do so. Bitcoin is therefore inherently a black market money. Its security architecture necessarily assumes it is operating without state permission.

The security of Bitcoin does not extend to white market systems. **Any system dependent upon the value proposition of Bitcoin must also be black market.**

Public Data Principle

It follows from the risk sharing principle that system security depends upon covert mining and trade. A coin exists as a mutually-beneficial market between miners and merchants for the confirmation of transactions within blocks in exchange for fees.

The necessarily covert activities are listed by role:

Miner

1. obtain blocks [to build upon]
2. obtain unconfirmed transactions [to earn fees from]
3. create and distribute blocks [to cause others to build upon]
4. receive payment for confirmations [to finance operations]

Merchant

1. obtain blocks [to validate customer payment]
2. obtain unconfirmed transactions (optional) [to anticipate payments and fees]
3. create and distribute transactions [to obtain customer payment]
4. make payment for confirmations [to compensate confirmation]

If blocks cannot be obtained anonymously the system is insecure. The inability to obtain the strongest blocks available to other people is a network partition, which implies localized insecurity. However neither anonymity, nor its opposite identity, can ensure one sees the strongest branch at any given time. In other words, any attempt to mitigate partitioning with the introduction of identity is a false choice that sacrifices system security for the false promise of ensuring localized security.

It is not essential that all miners or merchants see all transactions at any given time. However broad visibility is preferable as it produces the most robust competition for fees and best leading information. In other words, a market where every participant sees all of the transactions all of the time is a perfect market. Asking the network for specific transactions, as opposed to all (or summary information about all), is a source of taint and must be avoided in the interest of security as well.

Creation of blocks and transactions does not inherently expose identity, however public distribution of either is the primary source of taint. To the extent that miners openly self-identify, they are relying on the assumption of a low-threat environment, not contributing to system security. Avoiding taint when disseminating blocks and transactions requires use of an anonymous connection to a community server. This ensures the distribution network never has access to identifying information.

Proof of work exists to preserve anonymity of miners. There is no signature associated with mining and energy is presumed to be ubiquitous. Similarly, the ability to pay anonymously for confirmation is the reason for transaction fee inclusion. It is sufficient to pay a miner directly (off chain) for confirmation, however this exposes the merchant and miner to each other, and makes it more difficult to estimate fees anonymously.

Bitcoin is novel in that all financial transactions can be validated from public data and without identity. Centralized financial systems rely on either trust in (cryptographically-identifiable) connections to other parties or trust in (cryptographically-verifiable) signatures on transmitted data. This is the essence of trust-based systems; certain authorities have secrets that others use to verify that authenticity. **The reason for validation is to eliminate the use of identity and thereby authority.**

Social Network Principle

In the terminology of Paul Baran's 1964 paper on distributed networks the importance of topology in network design is the ability of communications to withstand the loss of a certain number of nodes. A centralized (star) network will fail with the loss of one node. A distributed (mesh) network is more resilient. A hybrid of these systems is considered decentralized.

As a money Bitcoin forms a social graph. Only a person can decide to accept one money or another in trade. A set of people sharing the same definition for a money is referred to as a consensus. Authority in a monetary system is the power to define the money. Bitcoin is a tool that people can use to defend against the tendency toward authority, in order to preserve their agreement and therefore utility in the money.

In distributed systems terminology a Bitcoin "node" is a person and the system is money. It does not matter how many machines the person controls, the loss of that person is a loss of a node in the system (including all of the person's machines). A centralized money cannot withstand the loss of one person. If that one person changes their rules, the original money ceases to exist. As shown in Risk Sharing Principle, Bitcoin relies on decentralization to allow people to resist authority. This decentralization makes the money more able to withstand the loss of more people when faced with state attacks. A loss in this sense is the refusal of the person to trade in the money.

State Banking Principle

Economy is not zero sum. People create and consume property. It is of course possible for the amount of capital to grow and to shrink. Each voluntary trade objectively improves the position of both people in the trade, as otherwise it would not happen. Consequently, the ability to trade freely increases accumulated capital. Assuming no change in time preference, this increase of a person's savings increases the amount invested and hoarded capital. Increased investment results in increased production, in a compounding cycle. It is essential to recognize that money merely facilitates trade, it is not a source of production, and therefore not the source of wealth.

State money is however an effective form of taxation, an anti-market force. State banking is often conflated with free banking, leading to a general distrust of bank lending. In free banking a bank is simply an investment fund. It may offer demand withdrawal of customer accounts, which creates risk for customers of, and investors in, the bank. Free banking manages this risk through the possibility of failure. Yet this scenario is not limited to banking.

Any person (including banks and other businesses) who underestimates or overestimates liquidity requirements is subject to failure. Insufficient liquidity (liquidity crisis) requires the person to borrow again, and excess liquidity produces lower returns (cash drag), an opportunity cost. State banking compels the taxpayer to act as a lender of last resort to member banks that underestimate liquidity requirements. This facility would serve no purpose unless the taxpayer offered such loans at a discount to the bank's market cost of capital, as otherwise the bank would simply rely on capital markets. This is why the interest rate offered to member banks is called the discount rate and why the loans are theoretically considered a "last resort".

The current discount rate in the United States is 3% whereas the historical real market rate of return is between 5.4 and 10% depending on time horizon. A state bank that can borrow at 3% and earn 5.4-10% by lending pockets a 2.4-7% return on taxpayer money. Typically the discounted money is not actually extracted through direct taxation but indirectly through seigniorage. The latter is monetary expansion, which then undergoes necessary market credit expansion. The resulting credit expansion, most visible through bank lending, tends to receive unwarranted blame for what is actually a consequence of seigniorage.

A member bank has a strong financial incentive to underestimate its liquidity requirements as it can profit by borrowing from the taxpayer. This is known as moral hazard and is why state banking regulators set minimum "reserve" (hoarding) requirements for its members. However, the state as well has a strong incentive for member banks to utilize this facility, as it creates revenue for the treasury. As such reserve requirements are typically insufficient and the discount window is used in the everyday course of business, not as a last resort. When a member repays the loan principle with interest to the state (taxpayer), and the state subsequently destroys the principle, the expanded money supply and associated credit expansion both contract.

However, both the state and its member banks have profited, at the rate of 3% and 2.4-7% respectively, at the expense of the taxpayer. The taxpayer was without the principle during the time of the discounted loan, while the member bank was able to invest it for the shared benefit of itself and the state. The cost to the taxpayer is objectively 5.4-10% on every state loan issued to member banks. It matters not whether the taxpayer was taxed directly or indirectly to obtain the lent capital.

No new capital is created through state banking, it is simply a means of taxation. The capital lent by member banks is available only to the extent it is not available to the free market (taxpayer). Furthermore, this capital is directed by the state through its member banks, instead of by the free market. There is a strong political incentive to direct member bank capital in ways that the free market might not, such as toward affordable housing and small business. Similarly, a member bank will accept much higher risk in its investments than its hoard can support. These forces produce malinvestment which manifests as the boom-and-bust phenomenon known as the business cycle. When economy is booming the state and its member banks profit and when busting the taxpayer covers the debt through the discount window. The state then blames this cycle on business (hence the name) in order to obtain more power.

Substitution Principle

A substitute good is one that can be used in place of another. As the price of a product rises, at some level people either move to substitutes or cease use altogether.

While a substitute would be less desirable at the same price as the original product, its lower price offsets this preference. In this manner the presence of substitutes reduces demand for the original good. The substitute competes with the original just as does increased supply of the original.

Given that a coin has fixed supply, it is commonly assumed that no supply side increase can reduce upward price pressure. As shown in Stability Property, Bitcoin integrates transfer fees which necessarily rise with use. This unique characteristic creates downward price pressure by reducing demand. **But this rising cost also makes substitutes viable, creating downward price pressure by effectively increasing supply.**

There is nothing preventing the evolution of multiple similar coins. It is possible for these to exhibit nearly indistinguishable monetary properties, minimizing the substitution tradeoff. As shown in Consolidation Principle, there is always pressure toward a single money, as this eliminates the exchange cost. However this pressure is at odds with rising costs, and at some level of use must give way to substitution (or disuse).

There is a theory that since creation of new coins costs nothing, the substitution principle implies that Bitcoin must become worthless due to unlimited free supply. This ignores the fact that Bitcoin requires people pay to use it. This is as true for a

second coin as it is for the first. Yet increasing supply reduces demand. At some point demand is not sufficient to produce/secure more supply, and as such the theory is invalid. This is the same relationship that holds with commodity monies and indeed all products.

Cryptodynamic Principles

Cryptodynamics is a term coined here for the purpose of easily referring to the fundamental principles of Bitcoin. This is intended to both inform understanding of Bitcoin and differentiate it from other technologies. The principles are the minimal subset of cryptoeconomic principles necessary to achieve this objective.

While the choice of name is not essential, a rationale for it is provided below.

Crypto

“A crypto currency is a [money] that uses strong cryptography to secure financial transactions, control the creation of additional units, and verify the transfer of [units].” - Wikipedia

Dynamics

“Dynamics is the branch of applied mathematics [...] concerned with the study of forces [...] and their effect on motion.” - Wikipedia)

Crypto + Dynamics

Cryptodynamics is the set of forces that secure Bitcoin transactions by controlling (1) definition of units, and (2) transfer of units.

Principles

Security force is entirely human in nature. People must act to secure anything, including Bitcoin. As an economic system, Bitcoin security can only expect people to act in an economically-rational manner (self interest). As such Bitcoin security forces are entirely based on the self-interested actions of individual persons, specifically:

- Risk Sharing
- Energy Sinking
- Power Regulating

These forces depend upon each other in order. Without risk sharing energy cannot be sunk into the system to balance the power of a censor. With these three forces intact Bitcoin can be secure. Without any one of them a technology is not Bitcoin.

It cannot be assumed that, given the incorporation of these forces, a Bitcoin implantation is securable. Furthermore one may be more so than another. **It is only the case that given the incorporation of these forces a technology is a Bitcoin and that without them it is not.**

The possibility of security afforded by these forces can be referred to as “cryptodynamic security”. So, for example, a “permissioned blockchain” violates the risk sharing principle, a strictly proof-of-stake (PoS) technology violates the energy sinking principle, and a money entirely reliant on subsidy for confirmation compensation violates the power balancing principle. None of these are cryptodynamically secure.

Censorship Resistance Property

Resistance to censorship is a consequence of transaction fees. Censorship enforcement is indistinguishable from soft fork enforcement, with majority hash power rejecting non-censoring blocks. Without such enforcement transactions are confirmed on an economically-rational basis despite individual miner subjectivity.

A majority miner is financially profitable. As such there is no cost to acquiring the means of censorship. As mining is necessarily an anonymous role it is always possible for any actor to acquire and deploy majority hash power, and to control it at any given time. As shown in Proof of Work Fallacy, hard forks cannot be used to selectively evict the censor and instead accelerate coin collapse.

In the case of active censorship, fees may rise on transactions that fail to confirm. This fee premium creates a greater potential profit for miners who confirm censored transactions. At a sufficient level this opportunity produces additional competition and

therefore increasing overall hash rate.

If rising non-censoring hash power exceeds that of the censor, its enforcement fails. The censor is thus faced with the choice of subsidizing operations or abandoning the effort. Only the state can perpetually subsidize operations, as it can compel tax and gains from preservation of its own currency regime. **The state must consume taxes to at least the level of the fee premium to maintain censorship enforcement.**

A coin without integrated fees would either fail to a censor or evolve a side fee market. As shown in Side Fee Fallacy it is not necessary that fees be integrated, however fee integration is an important anonymity technique. In either case censorship resistance arises only from the fee premium. The subsidy portion of the block reward does not contribute to censorship resistance because the censor earns the same subsidy as other miners.

It is possible that censorship enforcement could result in a price collapse, causing the censor to incur a loss on operations. However in this case its objective has been achieved, with no opportunity for the economy to counter the censor. This collapse might be achieved at negligible cost by simply demonstrating the intent to censor. It is also possible that a censorship soft fork could lead to a price increase, as white market business embraces the associated state approval. Nevertheless, for the coin to survive, its economy must continue to generate a fee premium sufficient to overpower the censor.

It cannot be shown that the economy will generate sufficient fees to overpower a censor. Similarly, it cannot be shown that a censor will be willing and able to subsidize operations at any given level. It is therefore not possible to prove censorship resistance. This is why resistance to state control is axiomatic.

Consensus Property

People generally think of consensus in the context of a fixed membership, like a jury. In this model consensus implies that all members must agree. But because Bitcoin membership is permissionless and therefore not fixed, there is always complete agreement, as implied by membership. In this model consensus refers to the size of the membership (economy), not a condition of agreement.

A consensus may fragment or consolidate. Generally a larger consensus provides greater utility and greater security by more broadly sharing risk.

Stability Property

Value is subjective and therefore price constancy is an economic fiction. The exchange price of a money is is determined by its supply and demand which is in turn affected by the demand schedules of all people for all products. The stability of a money is not a tendency toward constant prices in all other things, it is a damping relationship between demand for the money and its supply.

We can organize monies into three supply categories:

- Rational supply (commodity)
- Arbitrary supply (fiat)
- Fixed supply (bitcoin)

In any money, destruction of units decreases supply and therefore increases the value of those that remain. Given that there is no financial incentive for loss it does not impact stability.

Commodity money supply increases due to the financial incentive to produce more when price is expected to be at or above production cost (inclusive of capital cost). This relationship between price and supply is predictable despite price (and therefore supply) not being so. Because price is not predictable this monetary inflation cannot be capitalized. Therefore all holders of the money suffer a reduction in value from the increased supply. Competition ensures that this production, financed by existing holders, is limited to capital cost (interest). The feedback of value decrease resulting from supply increase reduces production incentive, creating stability.

Fiat money supply is increased arbitrarily (or taxed as demurrage)) by the sovereign due to the financial reward of seigniorage. When this monetary inflation is predictable it can be capitalized, which discounts the return on seigniorage. As such changes to supply are often not published. Due to state monopoly protection (i.e. production is the crime of counterfeit), competition cannot effectively limit returns. The resulting sovereign profit (tax) is the reward of seigniorage and the reason for fiat. Monopoly protection is the sole economic distinction between commodity and fiat money. The supply increase caused by seigniorage is

mitigated only by political unrest as people resist the consequential value decrease. This unrest initially manifests as capital flight, which is countered by foreign exchange controls.

Bitcoin supply is independent of price. Given that subsidy is predictable it is capitalized and has no affect on price over time. Its purpose is to rationally distribute units and so is eventually phased out. As fees necessarily rise with demand the utility threshold eliminates demand for transaction of value below the threshold. More generally, the fee level rises to the point where monetary substitutes are more cost-effective for a given value transaction. **Stability therefore results from limiting demand directly, in contrast to relying on an increase in supply to do so.** Stability implies that price is bounded, yet it can rise with increased effective transaction carrying capacity of the coin, and with increased utility relative to substitutes.

Utility Threshold Property

Utility is expressed as preference for the coin over substitutes, for transfers of a comparable value. Increasing utility implies a rising fee level. A person with higher marginal utility for a confirmation offers a higher fee than another. Given differences in the market fee level over time, another person may offer a lower fee in expectation of a longer time to confirm. Others will not transact on the chain, relying instead on substitutes.

Increasing utility therefore implies increasing average transfer value, as rising fees will otherwise cause the cost of transfer to exceed the value transferred. Greater depth implies greater confirmation security. Therefore time can be traded for *higher* security against double spend. However time cannot be reduced below one block period to achieve *lower* security. The lowest levels of security are none (unconfirmed) and minimal (one confirmation). There is no trade to be made between these levels.

Higher fees imply higher hash rate, mitigating the need to increase confirmation depth for higher value transfers. **But given there is no way to reduce security for lower value transfers, the useful minimum value transfer rises with utility.** Failure to support transfers in a certain value range implies substitutes are cheaper in that range. This implies the possibility of coexisting moneys to service distinct value ranges. However, similar coins all exhibit this property.

Rule differences in terms of block period or size do not change this relationship. The effect of these coin variations is strictly proportional. Even unlimited size blocks must produce fee levels that price out low value transfers.

Zero Sum Property

Bitcoin mining is a zero sum game. On average the chain grows by one block every 10 minutes, with the full reward controlled by its miner. Miners compete to achieve this reward and will, apart from pooling pressures, each average a number of rewards proportional to hash power. The difference between a miner's cost and this reward over time is the rate of return on capital invested in the mine.

There are two aspects of the zero sum property:

- For the time period between organizations one miner earns a reward and all other miners earn no reward. Neither price, hash rate, difficulty, inflation, fees, nor anything else has any effect on this property.
- The magnitude of rewards, in either coin units or exchange price, has no effect on the rate of return on capital.

Idealized Bitcoin mining is a closed system. Return on capital varies relative to other mines, due to the proximity premium and variance discount protocol flaws, as well as economies of scale and operator efficiency. Yet because these only impact the relative cost of hash power, the proportionality of return rates is affected, not overall returns.

Actual Bitcoin is not a closed system. The market and anti-market pooling pressures of variation and distortion (respectively) are external. Fundamentally Bitcoin exists to defend markets, necessarily pitting distortion against variation (or lack thereof).

When a distortion is applied to a miner in this zero sum system, all other miners are affected. For example, a subsidy (not to be confused with a consensus subsidy) to one miner acts as a tax on all others, and a tax on one miner acts as a subsidy to all others. The subsidized miner operates at a lower cost for the same hash rate, or has a higher effective hash rate for the same cost. The taxed miner operates at a higher cost for the same hash rate, or has a lower effective hash rate for the same cost.

A subsidizer expects no return on capital, otherwise he/she would be considered an investor. Investment is a market force whereby the miner pays a market price for capital. With a higher effective rate of return the subsidized miner attracts more capital than other miners, continuing to expand hash power until there is only one miner. The subsidizer's objective is ultimately control over the subsidized mine.

A tax on mining has the effect of moving all hash power to untaxed mines, beyond the reach of the taxing authority, as capital

seeks market returns. If applied broadly, this can give the authority control through its own mining operation. In other words, the authority can suppress competition. This can also be accomplished through a 100% tax, whereby the authority co-opts mines. The effect is the same, the taxed miner is put out of business, and the proceeds of the tax are applied to control.

The consequences of zero sum mining with inherent pooling pressure are explored in Threat Level Paradox.

Threat Level Paradox

As implied by the zero sum property, presumably the only way to defeat external subsidy is to mine at a capital loss relative to market return on capital. Similarly it seems that only way to defeat tax, up to and including a 100% tax (prohibition), is to mine beyond the reach of the taxing authority, such as in secret. As with all black markets there is an increased cost to subversive mining. Competing against subsidized mining compounds the cost.

If one accepts the axiom of resistance one must assume that both tax and subsidy will be used to reduce the cost of controlling Bitcoin. Using the power to subsidize mining (via tax revenues), states can cause pooling in the region of the subsidy. Once majority hash power is focused the state can use its taxing (regulatory) power in the region to compel censorship.

Therefore in order to enjoy the benefits of a hard currency, it would seem that people will ultimately have to mine at a loss. However, censorship creates the opportunity for others to mine profitably to the extent that people are willing to offset this cost with fees. This black market is Bitcoin's censorship resistance. People pay a higher price for certain transactions, and in order to maintain that higher price the state must also suffer the expense, despite its ineffectiveness.

Paradoxically, this tool works well when money is under attack and poorly otherwise. If there was no internal pooling-pressure these cases would be balanced. But risk distribution is essential to subversive mining, and pooling pressure works *against* distribution. So there is ever-expanding attack surface with no pressure to contract unless effective monetary alternatives are suppressed. The suppression of alternatives raises reward utility to the miner in the region of suppression. The paradox applies as well to centralization pressures.

The expected consequence is that Bitcoin will not be well *prepared* for attacks because it is financially disadvantageous for people in a low threat environment.

Miner Business Model

Miners play a zero sum game within a positive sum economy. They compete with each other, not the economy. Rising utility is the reflection of a positive sum and a natural consequence of trade.

It has been argued that blocks mined in a period of rising price produce outsized returns for miners, at least until the adjustment. This idea is based on the common failure to understand that market prices are not predictable. Wagers on price change are speculative. There is no reason to assume that Bitcoin *speculation* is any more or less effective than any other. However if a rising price was predictable then capital would predict it as well, canceling the possibility of any outsized return.

Bitcoin mining *investment* on the other hand is based on the predictable relationship between profit and competition over time. That relationship predicts that the average of all mining approaches the market rate of interest. As with all markets, shorter time periods are unpredictable in price and longer periods approach market returns. Ultimately time preference controls the market rate of investment return.

So how does a miner achieve outsized returns? It cannot be done with side fee agreements. There is only one way to make a higher-than-market rate of return, which is to have a below average cost of hash power for the coin. This is achieved by either taking advantage of pooling pressures or through superior operational efficiency. Because of the zero sum property, these are offset by lower-than-market rates of return by other miners. The premium therefore declines for an honest miner above 50% hash power, to zero at 100%.

However other miners will eventually exit as their capital seeks market returns. This would leave one miner, bound to market returns. In other words, making outsized returns requires others from whom to capture those returns. The highest return that can be sustained is a function of the greatest opportunity cost others are willing to sustain. This is a function of differential reward utility, as discussed in Threat Level Paradox.

By limiting dividends to market rates of return and reinvesting all remaining reward, a miner can maintain a constant hash power and thereby obtain market returns against a capital base proportional to Bitcoin capitalization. Reinvesting dividends increases hash power and liquidation decreases it. Grinds are liquidated by taking each device offline as it becomes a net negative producer, or discounting those future returns by selling the device.

Mining rate of return on capital is dependent on time preference alone. The relationship between the economy and miners is further explored in Balance of Power Fallacy.

Qualitative Security Model

Decentralization Model

In Social Network Principle it is shown that Bitcoin is a network of human relationships. This can be modeled as a directed graph#Directed_graph) where each vertex represents a merchant and each edge represents a trade for bitcoin. Edges indicate the direction of movement of coin and are quantified in the number of units traded. All owners are presumed to have been merchants at the time of coin receipt, including as miners (selling confirmations) and as recipients of charity (selling goodwill)).

If a person is not personally accepting coin, or does not personally validate coin accepted, the person cannot reject invalid coin. The person is entrusting this task to a central authority. All people utilizing the same delegate are reduced to just one vertex that represents the delegate.

For any period of time, economic security is a function of the number of merchants and the similarity of amounts traded. The strongest economy would be all people in the world trading for the same number of units in the period, an ideal which can be called a "distributed" (or fully-decentralized) economy. The weakest would be one delegate accepting all units traded in the period, which would be a "centralized" economy.

More specifically, the system is most economically decentralized which has the greatest number of vertices (merchants) with the lowest coefficient of variation in the incoming edges (receipts). Defining a *distribution* function as the inverse of coefficient of variation we obtain:

```
economic-decentralization = distribution(receipts) * merchants
```

Similar to economic security, confirmation security can be modeled as an edgeless graph. Each miner is represented by one vertex on the graph. A grinder is not a miner as the grinder has no decision-making ability, only the miner is represented. The total hash power employed by a miner is the weight of the vertex.

For any period of time, confirmation security is a function of the number of miners and the similarity of hash power they directed. The strongest censorship resistance would be all people in the world mining at the same hash power in the period, an ideal which can be called "distributed" (or fully-decentralized) confirmation. The weakest would be one miner with 100% of hash power, which would be "centralized" confirmation.

More specifically, the system is most decentralized in confirmation which has the greatest number of vertices (miners) with the highest distribution in weights (hash power):

```
confirmation-decentralization = distribution(hash-power) * miners
```

Security Model

Decentralization alone is not security. **Security is the product of activity, distribution of that activity, and the fraction of participating humanity.**

```
security = activity * distribution * participation
```

Given that there is no limit to humanity, trade or computation, the level of security in each axis is unbounded. Security is also unbounded with perfect distribution (i.e. infinite decentralization). A minimum level of zero in each is achieved with either no participation or no activity. Economic and confirmation security can thus be defined as:

```
economic-security     = receipts   * distribution(receipts)   * [merchants / humanity]
confirmation-security = hash-power * distribution(hash-power) * [miners    / humanity]
```

Limits of the Model

These relations do not say anything about the absolute effectiveness represented by any value, or the relative effectiveness of any two values except that a greater value represents a greater effectiveness. This is not due to a deficiency in the model. The factors include people, specifically the effectiveness of their individual abilities to resist and their perception of value in the

money. All who validate or mine offer some level of resistance, but there is no implied continuity. We refer to a "level" of security, not an "amount" of security.

As show in in Public Data Principle, anonymity is a tool that aids in defending one's ability to trade and/or mine. As such the level of decentralization can never be measured; the model is a conceptual aid. As shown in Balance of Power Fallacy, the security afforded by each of the two sub-models is complimentary and independent of the other. While people could decide to trade and/or mine independently in the future, the Cockroach Fallacy shows that they are not contributing to security until they do so. The model represents security as it exits in the period.

Proximity Premium Flaw

Latency is the time required for communication. Information moves at a speed not greater than the speed of light and therefore latency cannot be eliminated.

Different distances between miners implies announcements will be known to some before others. While a miner remains unaware of an announcement he wastes capital grinding on a weak candidate. As more time passes it becomes exponentially less likely that the miner will be rewarded for the candidate. Miners therefore compete to see announcements before other miners, as this reduces opportunity cost.

If we were to disperse miners with equal hash rate at equidistant points around the Earth they would experience the same average latency. Yet due to the financial benefit of reduced latency, they would tend to move closer to each other. This force is pooling pressure, and manifests in several ways.

The proximity-based pooling pressure is a consequence of the chain sequence required by consensus rules. **Bitcoin prescribes winner-take-all ordering, which produces disproportionate opportunity cost.** The variance discount is another pooling pressure caused by consensus.

The defense that Bitcoin *intends* to raise is market defense against anti-market forces. To do this it must distribute hash power broadly among people so that it becomes difficult to co-opt. However pooling pressures inherent in the consensus work against this objective. This is why the characteristic is termed a flaw.

Variance Discount Flaw

Variance is the varying frequency of achieving a reward. Variance is inherent to the probabilistic nature of mining and cannot be eliminated.

As a matter of consensus, different hash power among miners implies rewards will be earned by some more frequently than others. With 10% hash power one might expect to be rewarded 10 times more frequently than with 1%, though the multiple is actually greater due to the proximity premium. Actual results are unpredictable and can vary significantly. But it is sufficient here in both cases to assume proportionality. In this example one miner receives a reward every 100 minutes and the other every 1000 minutes. Assuming identical rewards per block, the magnitude of the reward is also proportional to hash power.

Consider then that a tiny miner might have to wait years before any reward. Despite being rewarded proportionally a smaller miner is faced with a perceived deficiency in relation to the larger miner. It can improve cash flow to receive a fraction of the reward more frequently. There is also the possibility that a mine is misconfigured and can never succeed. For these reasons miners discount high variance. Smaller miners will convert their mines to grinds and pay an aggregating miner for reduced variance. This is the rationale behind P2Pool, but because distributed variance reduction is less efficient the pooling pressure remains.

The variance-based pooling pressure is a consequence of singular difficulty as required by consensus rules. **Small miners must compete at high difficulty despite low hash power, which magnifies inherent variance.** The proximity premium is another pooling pressure caused by consensus.

The defense that Bitcoin *intends* to raise is market defense against anti-market forces. To do this it must distribute hash power broadly among people so that it becomes difficult to co-opt. However pooling pressures inherent in the consensus work against this objective. This is why the characteristic is termed a flaw.

Centralization Risk

Bitcoin weakness results from centralization and from pooling. Forces that produce aggregated mining are called pooling

pressures. While pooling weakens confirmation security, centralization weakens the security of consensus rules. Weakness is the result of fewer people with whom to share risk.

Consensus risk is shared among active merchants only, as they are the people who have the ability to refuse trade of property for units that fail to conform to their rules. Financial forces that reduce the number of merchants are called centralization pressures. The problem of delegation is that it is commonly coupled with centralization, as is typical in web wallets. The wallet not only owns the saved units but typically also controls validation of units received in trade. **The latter reduces power over consensus rules to one person for all of the wallets of the service.**

Centralization pressures include:

- Difficulty of use discount.
- On-chain settlement discount.

If exchange is difficult for a customer, the merchant must discount merchandise in order to accept the coin. If exchange is difficult for the merchant, an additional cost is incurred. If referring payments to a trusted third party reduces the size of this discount and/or cost, return on capital is increased.

Transfer incurs fees which also requires a merchant to discount merchandise. If using a trusted intermediary to settle transfers off-chain reduces fees and thereby the discount, return on capital is increased.

Centralization manifests as:

- Payment processors
- Web and other trusting wallets
- Hosted APIs for accessing chain

In a low threat environment the merchant has diminished financial incentive to subsidize Bitcoin security. As the cost of alternatives increases the discount becomes unavoidable. At this point the customer decides the pay a higher price or the merchant closes the business as capital seeks market rates of return.

Pooling Pressure Risk

Pooling pressure is the set of financial incentives for hash rate aggregation, specifically:

- Proximity Premium
- Variance Discount
- Market Variation
- Market Distortion
- Economies of Scale

While latency and variance are unavoidable, consensus rules actually *create* the first two financial incentives. Variation is a consequence of varying market price for mining resources. Distortion is a consequence of varying non-market costs including tax, regulation, subsidy, and patent; the force that Bitcoin is intended to resist. In a high threat environment economies of scale may become negative due to the cost associated with greater visibility but may otherwise be positive.

There are several manifestations of pooling. One is geographic, where independent miners become physically closer together. Another is cooperative, where formerly-independent miners join forces and co-locate grinding. Another is virtual, where miners become grinders and aggregate hash rate to a single remote miner. Another is the existence of relays, which aggregate miner hash power. Another is capital flow, since the higher hash rate associated with greater capital utilization is a form of co-location.

Given a perpetual positive pressure, transaction selection will eventually be reduced to one person's control. It is possible that this is already the case. The risk to Bitcoin is that one person is the sole defense of utility, making successful co-option inevitable. This risk cannot be mitigated by the economy.

Pooling pressure is a Bitcoin analogy to the United States Federal Reserve system. The system was designed to facilitate tax through debasement of a soft money. It offered state support for a monetary proxy in trade for hard money. This combination was designed to create a *pressure* to collect hard money at the central authority. Once this collection was sufficient the state did

away with the pretense and simply seized all remaining hard money. All states have similar systems and cooperate to defend them.

This does not imply that mining is adversarial to Bitcoin. Following the analogy, free banking is not adversarial to gold. Mining is a *necessary part* of Bitcoin. Pooling represents risk, though pooling pressure is not created by miners but by flaws in Bitcoin itself.

ASIC Monopoly Fallacy

There is a theory that Bitcoin ASIC price is controlled by a cartel of miners, creating a disproportionate advantage to mining partners of the cartel.

There is no economic difference between a cartel and a single organization. Changing organizational size is a free market outcome observable as capital seeks optimal economies of scale. If partners receive ASICs at a price that produces a below market return on capital, it amounts to an internal subsidy between partners. The same is true of a price that produces an above market return on capital, with the subsidy in the opposite direction. As such there is no net advantage to such discounting between partners.

Production is generally set at a level intended to produce a maximum rate of return on capital. The only economically rational way for a producer to raise price is to limit production below that optimum. Otherwise higher price implies unsold inventory, resulting in lower net returns. This implies that production must be restricted by the cartel in order to raise unit price for non-partners.

Limiting production leaves an opportunity for other producers to capture customers with a lower marginal utility for the product, as those customers would otherwise be unserved. Thus competition lowers price until the market clears. A free market seeks the clearing price that produces the global return on capital. A current price above this level increases production and below decreases production. It is time preference that determines the rate of return on capital.

Unless production is disproportionately subject to anti-market forces, such as tax or subsidy, everyone enjoys the same opportunity to raise capital and compete in production. If this does not happen it implies that returns on this line of business are consistent with average market returns. Tax and subsidy cause regional distortions but do not eliminate competition. **In other words, monopoly price is only produced by state grant of monopoly power.**

A related theory asserts that purchasing ASICs from this cartel increases its hash power. This is invalid on the basis of the above explanation of monopolistic pricing. The producer's capital will seek the same return in any line of business or investment. There is no reason to believe that the return will be disproportionate in ASICs.

A related theory asserts that the Bitcoin proof of work algorithm produces a pooling pressure, as a consequence of the supposed cartelization. If people truly believe that ASICs are overpriced the rational response is to raise capital and produce ASICs. But in any case market and anti-market forces alone control chip production and as such it does not constitute a protocol-based pooling pressure.

Auditability Fallacy

Solvency of a Bitcoin custodian cannot be audited. A custodian is a person with discretion both in the release of an asset and issuance of securities against it. If both release of the asset and the issuance of securities against it are controlled by consensus rules, then the relationship is not actually custodial. This is the distinction between a reserve and a layer. A layer is protocol-enforced (non-custodial) and therefore has nothing to audit.

A solvency audit requires simultaneous (atomic) proof of both the full amount of the asset held by a custodian and the securities issued against it. In the case of a national Bitcoin reserve this would require complete proof of all fiat (e.g. the security) issued against the reserve, as well as the Bitcoin held in reserve. Even in the case where the security is issued on a distinct public chain the atomicity requirement is not satisfied.

In some cases it may be considered sufficient to waive the atomicity requirement, accepting incorrectness under the assumption that material deviation would eventually be discovered. However in the case of state banking it is insufficient to detect the deviation. Historically it has not been difficult to detect such deviations. The difficulty arises in stopping them.

Balance of Power Fallacy

Power in Bitcoin rests with miners and merchants. Yet these two powers are not "balanced" between each other, as if locked in some sort of checks-and-balances system. Miner power is orthogonal to merchant power. Miners control transaction selection, merchants control validity, and neither can control the other. Not surprisingly, in the original description and implementation these roles were combined.

Power is not the same as influence. Merchants can influence miners by not buying the service. Miners can similarly influence merchants by not producing it. These choices manifest as splits or stalls. However the nature of power is that it can (and often does) ignore influence. The state has power; it can apply coercion and co-option while ignoring influence. Merchants and miners *together* have the power to defend against these aggressions, but neither can do so without the support of the other.

The balance of *power* in Bitcoin is between individuals and the state. Even states create systems that attempt to isolate their moneys from political control. Bitcoin is no different in that sense, incorporating the resistance axiom. Individuals can be miners and can be merchants. With broad distribution of these activities it becomes difficult for state actors to censor this market. **The idea that miners and merchants are in an adversarial position is a failure to understand the Bitcoin security model.**

Merchants purchase a service from miners and as such the two are engaged in trade. Merchants purchase mining services that conform to their rules for a satisfactory fee. They are free to split and miners are free to not mine at all, or to not select particular transactions for whatever reason suits them. Trade is neither adversarial nor asymmetrical, it is voluntary and mutually-beneficial, with all tensions resolved in price.

This failure in understanding leads people to believe that mining can be centrally pooled as long as merchants are not centralized in validation, as the economy can control the behavior of mining, rendering the system secure. This belief is incorrect but unfortunately people are drawing this invalid conclusion from recent events. A closely-related fallacy is the belief that a proof-of-work fork by merchants can control miner behavior.

Blockchain Fallacy

There is a theory that property ownership can be secured by immutable claim-keeping, both against claim loss and Custodial Risk.

Given that a claim is not itself the property, control of the property rests with the custodian against whom the claim is made. A custodian has the ability to surrender or retain the property and is therefore a trusted third party. Abrogation of a claim by its custodian is always mitigated by custodian signature, cryptographic or otherwise, with enforcement of the claim left to its holder.

The theory asserts that immutable claim-keeping provides security against loss of the claim by its owner, as nobody else would have an interest in the loss. However, in order to redeem the claim its owner must produce proof of ownership to the custodian. This requires that the owner not lose the secret that proves this ownership. As such the security of the claim against loss is not mitigated at all, it merely changes form. The theory is therefore invalid on the basis of loss prevention.

Storing a strong reference to the claim can reduce the size, and therefore cost, of its immutable storage. The claim may be in the form of a human or machine contract, and referenced as a one way hash. In either case the validation and execution of the contract is required for property transfer by the custodian. Therefore a referenced contract claim compounds loss risk with additional data, the contract.

As shown in Risk Sharing Principle, people are always the basis of security. People may act collectively to protect the immutability of a money, and therefore any claim data associated with control of the money. However, a custodian is a trusted third party. Immutable claims do not in any way mitigate direct attacks against, or by, a custodian. Where the custodian is the state or is subject to its control, the claim offers no security against the substitution of state authority in place of proven ownership of any claim. The theory is therefore also invalid on the basis of custodial failure.

Bitcoin as a money is non-custodial. Its units do not represent an asset held by a trusted third party. The money is traded directly between customer and merchant. In this sense *all merchants* are custodians of Bitcoin's value. **The blockchain fallacy arises from a misconception of the Bitcoin security model, attributing security to its technology as opposed to its distribution of merchants.** The term "blockchain technology" reinforces this error, implying that it is primarily the structure of Bitcoin's data that secures it.

Causation Fallacy

There is a theory that mining "follows" price, or more specifically, reward value. The implication is that mining is slaved to price, lacking any input into coin utility.

Consider the miner who responds only to historical reward values. This person cannot be the first miner, because the reward has no historical value. No price can be established because no trades have occurred. The miner might have heard that a number of unconfirmed units bought a pizza, but maybe the same units are double-spent. He must anticipate a certain level of future net return on capital that is unknowable until it either materializes, or does not. This is the nature of entrepreneurial risk. The risk must be taken before the product can exist. One might believe that the risk can be shifted to the consumer, with advance ordering. But at that point the consumer has become the entrepreneur, providing the capital for, and assuming the risk of, production.

It is certainly possible for a miner to respond only to historical reward values once history has been established by someone else's risk-taking. But what is the time window and method of averaging that predicts future reward values? The unique ability to predict exchange prices would provide the miner unlimited riches. If it could be done generally, price would never change, as all potential changes would be discounted upon first minting. So either price changes unpredictably, or not at all. In other words, every miner faces the same situation as the first. Historical prices do not exist that can predict future prices.

Assuming a market average return on mining capital generally, both overestimation and underestimation of reward value imply loss in relation to the cost of capital. Given the nature of competition, profits and losses (above and below market returns on capital respectively) experience constant negative existential pressure. In other words, the market attempts to eliminate these errors. But given the unpredictable nature of price, it can never actually do so. Production never seeks demand that exists, which is inherently historical, it always seeks demand that it anticipates. **Production continues to guess at future consumption and in doing so creates the opportunity for consumption** (or creates nothing at all).

One could say that miners converge on a market return on capital and by anticipating highest possible fee values. But merchants similarly converge on a market return on miner capital by seeking the lowest fee value. However, miners must anticipate merchants and risk mining before there can be any utility. So to the extent that there is any asymmetry, mining precedes transacting, just as all production must precede consumption. Assuming otherwise conflates the direction a market seeks with the manner in which it does so.

Cockroach Fallacy

There is a theory that aggregation does not materially reduce the security afforded by risk sharing because miners and the economy will disperse as necessary, similar to the scattering cockroaches disturbed by a light. **The theory irrationally implies that security *actually* exists because it *could* exist.** This is essentially a rejection of the Threat Level Paradox, which implies that security evolves over time under a persistent threat.

The theory relies on grinders switching miner allegiance. This is based on the Balance of Power Fallacy, which incorrectly models miners as the threat. A shift of hash power from one mine to another does not reduce pooling or the risk associated with it. The risk is that states co-opt large amounts of hash power, substantially reducing the cost of attack. It is an error to assume that states do not collaborate in defense of seigniorage.

> The International Monetary Fund (IMF) is an organization of 189 countries, working to foster global monetary cooperation...

As such one cannot assume that any large mine can exist outside of state control. A reduction in pooling requires an increase in the number of miners, specifically those who are willing and able to operate covertly. This requires that grinders suffer the increased cost associated with reduced pooling. Yet people cannot be expected to work against their own financial interest. In order for risk sharing to increase, the financial pressures against it must be reversed. An assumption to the contrary is economically irrational.

The theory also ignores economic centralization and delegation. It is an error to assume the economy can rapidly decentralize, and de-delegation would most likely be infeasible in the case of state attacks as currency controls commonly restrict transfer.

Credit Expansion Fallacy

Credit expansion is the apparent increase in money supply resulting from lending. When a loan is issued the lender and borrower both appear to hold the same money. Due to the inflationary nature of credit expansion, it is commonly treated as an adverse effect on people. Because banks are the most visible lenders this effect is often attributed to banking itself. There is a theory that Bitcoin can eliminate the effects of fractional banking, eliminating credit expansion.

Saving is a general term for hoarding, speculating and investing. Hoarding is owning as required to satisfy expected consumption. Speculation is owning in the expectation of advantageous price change. Investing is lending to another in

exchange for a fraction of the borrower's future earnings. Hoarding and speculating are outwardly indistinguishable, as the sole distinction is in the mind of the owner. As such this topic will, in the interest of simplification, refer to speculation as hoarding. Investment implies a trade between the lender and the borrower. The borrower owns the asset and the lender owns a fraction of the borrower's business. Investment includes both debt and equity contracts as the distinction is strictly financial, having no economic significance.

The distinction between hoarding and investment is essential to the understanding of credit expansion. A hoarded asset is under the control of its owner, as if in a vault, buried in the back yard, or stuffed in a mattress. This is inherent in the meaning of ownership. The lender of an asset is not the owner of the asset, even though a loan is considered savings.

A lender requires liquidity to operate, and as such must hoard a certain fraction of capital. When a loan is created the borrower owns the amount of capital lent. The borrower also requires liquidity, and so hoards a certain fraction of the loan. Any remainder of the loan is necessarily invested. This implies that the borrower has become a lender. The process continues until all capital that exists is hoarded.

The amount hoarded is sometimes referred to as the owner's "reserve", but properly it is the owner's hoard, a fraction of that owner's total savings. This use of the word reserve should not be confused with its use in the state money context of reserve currency. The term "fractional reserve banking" is a reference to the ratio of a bank's hoard to its total savings.

The decision to hoard vs. invest is based strictly on each person's time preference. Time preference is not derivable from any condition. It is, as the name implies, a human preference. Human preferences change and therefore so does time preference. Time preference determines the interest rate which can also be considered the cost of capital. An increase in the cost of capital resulting from increasing time preference causes credit available to contract, and a decrease has the opposite effect. With infinite time preference all capital would be hoarded, ending all production.

It matters not whether a lender is referred to as a "bank", all investment implies the same behavior. If banks operated with a 100% hoard they would not be lenders. This does not imply any reduction in lending, as the rate of lending is determined by time preference alone. **Bitcoin can be lent and as such does nothing to limit credit expansion.** The theory is therefore invalid.

Eliminating credit expansion is equivalent to the condition of infinite time preference, an infinite interest rate, no capital available for production, and no products available for consumption. In states where debt instruments are limited or prohibited by statute (usury laws), investment moves to equity instruments, loan sharking, or an end to production.

Decoupled Mining Fallacy

There is a theory that security is increased by decoupling reward from transaction selection in pooled mining. The theory holds that by sharing only the reward, control over transaction selection shifts to miners with less hash power. This implies a reduction in the variance discount and therefore an increase the competitiveness of smaller mines. Because smaller mines can presumably operate more covertly than larger, this in turn implies that censorship resistance is increased.

The theory fails to recognize that control over transaction selection remains with the pool operator, and is therefore invalid. The sole benefit is variance reduction, but this is only realized by the receipt of payment. As payment is discretionary any condition may be attached. Such conditions may include censorship and identity. Member recourse is to leave the pool for another, just as with a coupled pool. As such decoupled pools and coupled pools are equally subject to co-option.

There is a related theory that transparency of a decoupled pool is greater than that of a coupled pool, facilitating flight of members to non-censoring pools, therefore limiting the dominance of censoring pools. Generously accepting the assumptions of greater transparency and independent miners operating against financial self-interest, we are still left with the fact of co-option. The state can still reserve for itself the ability to operate with the financial advantages of pooling and the theory is therefore invalid.

This fallacy is similar to the Relay Fallacy in that all financial advantage depends on otherwise independent miners granting control over that advantage to a single person.

Dumping Fallacy

There is a theory that selling units from one side of a split coin for units of the other reduces the relative utility of the sold coin. However each sale requires a buyer. As a trade the action is symmetrical and therefore the theory is invalid.

There is a related theory that exchanging units from one side of a split coin constitutes dumping) of that coin, which reduces its

utility. The theory simply misrepresents the concept of dumping. Dumping is state subsidy (not to be confused with Bitcoin subsidy) of a product sold in another state. It is a levy on the taxpayers of the subsidizing state, typically applied to establish market share for the product. In the case where demand is elastic, the subsidy increases sales volume for the product by reducing price relative to the otherwise market price. The lower price increases demand, by capturing buyers with lower marginal utility for the product, until the market clears. In contrast to dumping, trading at market price doesn't reduce price because it is not subsidized.

Finally, there is a related theory that reduction of hoarding) generally reduces exchange prices of the hoarded property. This is true, however a transfer is not a reduction to hoarding levels unless the buyer of the hoarded property subsequently hoards it less than the seller. It is an error to assume this is the case.

Empty Block Fallacy

There is a theory that the mining of empty blocks is an attack. The theory does not require that the blocks are mined on a weak branch in an attempt to enable double-spending, nor does it specify what person is attacked.

Consider the following:

- The term "attack" implies theft. The Bitcoin whitepaper, for example, uses the term only to describe double-spend attempts.
- A reward consists of fees for transactions and a subsidy for the block. The miner who forgoes transaction fees by not including transactions is not rewarded for them.
- The miner's hash power contributes proportionally to the security of the network. The subsidy is compensation for that security during the inflationary phase. The purpose of inflation is to rationally distribute units. The rational distribution is specifically in exchange for hash power, not for transaction inclusion.
- Transaction confirmation is not assured. Fees are the *incentive* for confirmation. Lack of confirmation objectively implies insufficient fee.
- Empty block mining is entirely consistent with consensus rules and cannot be reasonably prevented by a new rule.

For these reasons the theory is invalid. However it is worth exploring the source of the fallacy. Because of the Zero Sum Property, there may be an assumption that mining an empty block "unfairly" takes away the opportunity for transactions to be confirmed.

Other miners retain the ability to confirm transactions in proportion to their hash power. If 10% of the hash power mines empty blocks, then confirmations will take 10% longer on average. However if a miner removes 10% of the total hash power, confirmations will also take 10% longer on average, until the next difficulty adjustment. Yet removing one's hash power is generally not considered an attack.

A miner commits capital to mining, producing hash power. Setting aside the effects of pooling, the miner is subsidized in proportion to hash power produced. Without this hash power other miners would produce the same average number of blocks at proportionally lower difficulty. In other words, *actual* attacks would be proportionally cheaper. So despite not being rewarded for including transactions, the miner is securing previously-confirmed transactions.

Given that the marginal cost of including a transaction is necessarily below average fee levels, the empty block miner is suffering an opportunity cost. This amounts to the miner subsidizing the security of the chain. While this seems economically irrational in the limited context of the coin, it can be rational due to the offsetting opportunity cost for waiting on a new non-empty candidate following an announcement.

While a given miner may consider it advantageous to mine empty blocks, it is within every other person's power to do otherwise. It is the exercise of this opportunity that secures the coin, even against *actual* attacks. **The shaming of those acting on self-interest is ineffective and counterproductive.**

Energy Exhaustion Fallacy

There is a theory that proof-of-work (PoW) may exhaust all energy available to people. PoW converts energy into a monotonically increasing double-spend barrier for any given transaction. This is comparable to the energy expended in securing any money against counterfeit (by its own issuer or otherwise).

The purpose of any security measure is to create a cost necessary to overcome the measure; i.e. a financial barrier. Bitcoin

creates its double-spend barrier by compelling the attacker to replace the branch of the targeted transaction with one of probabilistically greater work. Interestingly, such a replacement raises the barrier to subsequent attackers. **The energy expended is not independently important, the erected barrier is the attacker's necessary *financial* burden.**

The security barrier (S) of a block is the product of unit hash cost (C), hash rate (H), and period (T).

```
S = C * H * T
```

The adjustment varies hash rate to maintain a constant period for a given hash cost and security.

```
T = S / (C * H)
```

A constant period implies that hash rate is inversely proportional to cost for a given security.

```
H ~ S / C
```

As energy supply is reduced its price must increase, which reduces the amount expended for a given level of security. Therefore energy cannot be exhausted by mining and the theory is invalid.

Energy Store Fallacy

There is a theory that the value of energy expended by proof of work is converted to coin value, in effect “storing” the value for later consumption. To the extent that energy has value, the work is sold for bitcoin, and money is a store of value, this is correct.

However the theory errs in the implication that energy value expended in mining is unique in its contribution to value. Miners trade energy for units. However all merchants trade something. Furthermore, mining can be compensated by side fees, ultimately contributing nothing to demand.

All things offered in trade represent demand. **Apart from magnitude, one source of demand cannot be a generally greater determinant of value than another.** As such the theory is invalid.

Energy Waste Fallacy

There is a theory that proof-of-work (PoW) wastes energy. This implies that the level of security provided is greater than necessary or the same level of security can be provided by another externalized proof at a lower energy cost. An internalized proof, specifically proof-of-stake (PoS), is a different security model which is not considered here.

Total hash power is a function of reward, which is a function of fees, which are determined by the confirmation market. If a person considers current hash power insufficient to secure trade at a given value against double spend then the depth requirement increases. Additionally, as shown in Utility Threshold Property, transactions with insufficient value for even single confirmation security are priced out of the chain.

These upper and lower security bounds depend on confirmation cost and are therefore independent of the proof technique. **There is no *necessary* level of security, just a subjective confirmation depth and minimum utility.**

Confirmation security increases with the cost of generating each block. The double spend of a transaction requires that its branch be superseded by another with a probabilistically greater cost. So the only way energy cost can be reduced is by expending the same average per-block cost with a lower energy component.

PoW incurs cost in several forms, including labor, hardware, services, land, etc. Any other externalized proof consumes these same resources, though potentially in different proportion. The question of energy cost reduction is therefore reduced to whether an energy component of the cost of a proof can be replaced by an other resource component with the same cost. However the cost of the substitute resource includes all of is production costs, which must resolve to energy. The theory is therefore invalid.

Additionally, securing any coin has a cost to merchants. As such the fact of its use by them implies that it is preferred over alternatives. This implies the alternatives are ultimately more costly. As all costs are fundamentally resolved in energy consumption, it follows that the money in use is the most energy efficient.

Fee Recovery Fallacy

There is a theory that miners gain financial advantage over other miners by mining their own transactions and "recovering" their own fees.

The theory ignores the opportunity cost of mining block space without collecting payment for it. Payment of a fee to one's self is a financial non-event. Failure to collect a fee is a real cost in the amount forgone, as the cost of mining that portion of the block is uncompensated. **The result is a lower return on capital relative to miners who actually sell their block space.** Given the zero sum nature of mining this disparity allows the more rational miners to increase their hash power by reinvesting the higher return on capital.

A higher volume of transactions competing for confirmation implies an increase in the average fee level. This affects all miners equally, however a fee level increase only produces more competition, not an increased rate of return on capital. If anything, all miners suffer from the reduced utility caused by the increased demand and therefore fee level (which has been financed in full by the uncompensated miner). In other words the consequence is the opposite of that proposed, and the theory is therefore invalid.

There is a related theory that fee estimation tools may be fooled into recommending higher fees than are required. As shown in Side Fee Fallacy this implies a relationship between historical and future fee rates that does not exist, and that all fees are visible on chain, which is not the case.

Genetic Purity Fallacy

There is a theory that a coin is strongest when all validation is performed by a common implementation. According to this theory the complexity of consensus rule implementation implies a likelihood that multiple implementations will diverge, resulting an inadvertent chain split. The split implies financial loss by people on the weaker side. In addition to divergence, a single implementation risks a global stall of the network. The threat of financial loss implies lower utility and therefore system security.

Based on the presumption of high complexity, each update to the "one true client" produces the same likelihood of divergence. Similarly, dependency on external independently updated libraries has the same effect. In other words *it is not possible for there to be just one implementation*. In the case of the initial Bitcoin implementation both upgrade of the client and upgrade of an external dependency have resulted in unintended chain splits and material financial loss. Additionally, zero-day) flaws in this implementation have been published that could have produced a global stall.

A single implementation would produce a weakness directly analogous to that of a living species with genetic uniformity. In the case of a single implementation, both internal and external updates penetrate the economy quickly and deeply. The financial impact of a split is therefore more significant than that caused by a less widely-deployed implementation. In a scenario where ten implementations each supporting an even fraction of the economy there would be risk to at most 10% of the economy for any given update, whereas the update of a single universally-deployed implementation reaches the maximum split risk of 50%. The theory is therefore not only invalid but expresses the opposite of actual behavior.

Halving Fallacy

Bitcoin consensus-rules produce a predictable rate of monetary inflation. This rate is reduced periodically at a point called the halving. There are several step functions in Bitcoin. The halving occurs every 210,000 strong blocks, the difficulty adjustment every 2,016 strong blocks and chain organization approximately every 10 minutes. The numeric values that control these intervals are arbitrary yet the discontinuity is necessary due to the discrete intervals required for proof of work. There is a theory that the halving creates a financial cliff for miners that may lead to a perpetual stall. The theory is based on the confluence of two step functions (halving and difficulty), causing the period of another (organization) to expand dramatically due to coincident reduction in miner profits.

The theory assumes that the difficulty adjustment resets average miner economic profit to zero, allowing only the top half of miners (by profitability) to survive, eventually reducing mining to just a few miners. In other words the difficulty adjustment is considered a positive pooling pressure. However there is no reason to believe that the adjustment reduces *any* miner's profit to zero. The consequence of the assumption is not that there will be *few* miners, but that there will be *none*, due to the difficulty adjustment alone. The adjustment actually does nothing to regulate miner profits, it controls only the organization period. With no adjustment, profit would be unaffected while the organization period and therefore variance would respond to total hash rate. Time preference, which dictates market return on capital, regulates miner profits just as it does in every market.

Consider the case of no price change. In this case there is no reason to expect a change in total hash rate, no adjustments to

difficulty, and we can conclude that the average mine generates the market return on capital. In other words any number of independent miners *can* compete indefinitely (absent *actual* pooling pressures).

Consider also that price changes, difficulty adjustments, and reward fluctuations all effect miner profitability in the same manner. A difficulty adjustment and/or halving is therefore no more important to a miner than a comparable price fluctuation, and exhibits greater predictability. **Miner returns are expected to always average the market rate of return on capital.** Therefore the theory builds on an invalid assumption.

The theory also contemplates that reward may be insufficient to compensate miners for difficulty immediately following a halving. As such they may opt to reduce hash rate, extending confirmation times until fees rise, price rises and/or difficulty adjusts downward. Yet fees and price are determined in a market and can certainly rise to any level that people are willing to pay. There is no way to know what levels the market will support. Yet the two largest halvings have passed with no disruption. Fees and price have both risen, encouraging significant increases in total hash rate. Given that subsequent halvings will produce the equivalent of an exponentially *lesser* price reduction, there is no reason to believe future events will be any more interesting than past.

Finally, by controlling variance, which would otherwise rise indefinitely with price, adjustment is a *negative* pooling pressure. This is the opposite of the effect assumed by the theory.

Hoarding Fallacy

There is a theory that an increased level of hoarding produces an increased level of security in a coin. This is the similar to the Dumping Fallacy but is not necessarily based on a split.

An increase can only be described relative to some base level. If a person can be convinced that there is increased system security in a higher collective hoarding level, the theory holds that the person may decide to hoard more than would otherwise be optimal (i.e. the person's base level). This amounts to an actual individual cost with a presumed socialized benefit. In other words the theory depends on irrational economic behavior, even if the security benefit is actual, and is therefore invalid.

The presumed security benefit to an elevated level of hoarding stems from the theory that an owner has a say in validation and could act to prevent the economy from accepting what the owners collectively consider invalid money. However owners are not acting unless they trade units for something, and in this case it is the merchant who enforces consensus rules. **The possibility that owners could act in unison does not increase this zero level of control.** The theory is therefore invalid.

Finally, the theory implies that less trade in the coin will produce greater security. This is the opposite of the case. As shown in Qualitative Security Model, consensus rule enforcement requires ongoing trade. The price of a unit of the coin is arbitrary, but rises temporarily if individuals are convinced to engage in the fallacy. The benefit of this increase accrues to existing owners. The theory that price can only rise is a related speculative error explored in Lunar Fallacy. Even a provable perpetual price rise would not validate this theory, as it relates only to a temporary relative increase caused by financially sub-optimal individual decisions.

Ideal Money Fallacy

It has been proposed that the existence of an international non-political (i.e. objective) "value index" will result in people compelling states to "value target" their monies against the index, thereby eliminating price inflation. It has also been suggested that Bitcoin is such an index and will precipitate this scenario.

The leverage envisioned is the option to leave certain state monies for others. The movement is from monies of higher inflation to lower, based on comparison with the index. The consequence is that states must increasingly target their individual rates of price inflation to the index. This result is state monies "asymptotically" approaching the condition of Ideal Money represented by the index.

Ideal Money is state money with a zero rate of price inflation:

> ...there is no ideal rate of inflation that should be selected and chosen as the target but rather that the ideal concept would necessarily be that of a zero rate for what is called inflation.

Expression of the theory is both varied and limited (proof is left to the reader). However the above summary expresses all essential elements. Given these limitations it can be helpful to start with generous assumptions. Let us assume that a money can express objective value (see subjective theory of value), that Bitcoin is such a money, and that people generally have the ability to compare the value of Bitcoin to other major state monies. Let us also assume that, despite the apparent contradiction,

people will both generally use Bitcoin in trade (the source of the index) *and* will prefer to use state monies (a necessary premise).

If we also assume that people are free from legal tender laws, and their use of competing currencies does succeed in compelling states to "value target" Bitcoin, seigniorage will be eliminated. However, as shown in Stability Property, the purpose of state money (fiat) is to collect seigniorage, which is a tax. In other words, Ideal Money is a tax collection system that collects no tax. Granting the above assumptions, Ideal Money is the obsolescence of state money. **The proposal fails to consider the reason that fiat exists in the first place.**

Reconsider now the assumptions. Fiat requires the existence of legal tender laws and as such Gresham's Law always governs fiat:

> These examples show that, in the absence of effective legal tender laws, Gresham's Law works in reverse. If given the choice of what money to accept, people will transact with money they believe to be of highest long-term value. However, if not given the choice, and required to accept all money, good and bad, they will tend to keep the money of greater perceived value in their possession, and pass on the bad money to someone else. In short, in the absence of legal tender laws, the seller will not accept anything but money of certain value (good money), while the existence of legal tender laws will cause the buyer to offer only money with the lowest commodity value (bad money) as the creditor must accept such money at face value.

The proposal incorrectly assumes that Thiers' Law) governs. If this was the case people would not use fiat. It also ignores the existence of foreign exchange controls, which exist specifically to prevent capital flight. Such controls strengthen as capital flight accelerates, in order to preserve tax revenue. Finally, such controls materially limit price discovery in the index, making it less useful than the envisioned reference.

The proposal offers no rational explanation for how people will become able to move between state monies in the face of such controls. It assumes that people will better *recognize* the tax, due to the presence of the index and their ability to compare against it, and therefore will more effectively control the state's appetite for the tax. Given the near universal use of gold as a comparably objective index prior to the evolution of global fiat, it is not clear how fiat ever took hold if we can assume people will react to it in this manner.

There is an argument that Bitcoin is an objective index whereas gold is not. This is based on the inflationary supply of gold in contrast to the fixed supply of Bitcoin. This assumes that monetary inflation implies an unstable money whereas fixed supply implies a stable money. As shown in Stability Property, both monies are stable. The argument fails to acknowledge that value, as indicated by the index, is a consequence of both supply and demand. Gold demand is stabilized by inflation and Bitcoin's demand is stabilized by fees.

The theory is therefore invalid. Either fiat will cease to exist or it will collect tax. States only surrender this tax under extreme duress and in such cases only briefly. If anything the "ideal money" will be Bitcoin, and it will not trade freely with state monies (to the extent they remain).

Impotent Mining Fallacy

There is a theory that miners have no power. This is distinct from the closely-related Proof of Work Fallacy. The theory rests on the assumption that miners are subject to economic pressures that preclude sustained effective attacks. This theory leads people to believe that mining can be strongly pooled as long as merchants are not centralized, as the economy can control the behavior of mining, rendering the system secure. The consequence of this invalid theory is complacency regarding the insecurity caused by pooling.

The theory holds that if majority hash power double-spends then merchants will necessarily increase confirmation depth requirements, increasing the cost of subsequent attacks. At some point an equilibrium is reached where greater depths are considered sufficient for exchange. Given that this would preclude double-spending altogether, there would be no advantage to sustaining the attack. The theory accepts that attacks can happen, but not frequently enough to materially reduce utility.

The theory also holds that a miner cannot avoid selecting the highest fee transactions as this reduces relative reward, enriching other miners. This is presumed to lead to a loss of majority power and therefore an inability to continue. This aspect of the theory implies that miners cannot effectively censor.

The theory also considers that selfish mining by majority hash power is feasible, but in the absence of double spending and censorship, there is no adverse consequence to the economy. In this case the majority simply becomes the one miner as all others are unable to retain rewards. Despite lack of competition, hash rate and fee levels are maintained by the ever-looming *possibility* of competition.

Yet miners and merchants are trading partners, engaged voluntarily in mutually-beneficial activity. As explored in the Balance of Power Fallacy, neither can control the other and price is the resolution of all preferences. This would seem to support the theory, however the theory is actually a red herring as it **does not address the threat**. Bitcoin is designed to defend against *non*-market forces, specifically the state. Market forces are never a threat to the market itself.

The pooling of hash power eviscerates security, as states can simply co-opt it. But states can also build their own mines to the same effect. Bitcoin therefore requires both significant hash power *and* distribution of that power among people who are willing and able to risk state controls.

The state is an economically rational actor. Inflation is profitable for the issuer. Bitcoin's widespread use would prevent states from effectively levying the inflation tax. State attacks are therefore expected, and analogous attacks are commonplace. It is practically inevitable that states will subsidize attacks, but even the possibility invalidates the theory.

Inflation Fallacy

Bitcoin consensus rules create a perfectly predictable rate of monetary inflation. There is a theory that this causes loss of purchasing power. However it is essential to consider the effect of predictability on loss, and consequently whether the loss is actual.

Consider the following scenario:

> Alice sells Bob an acre of land for 1 bitcoin, confirmed at height 100. At the time of confirmation the supply of spendable bitcoin is (no more than) 5,000. Less than a week later, at height 1,000, Bob sells the same acre to Charlie for 10 bitcoin. Over this period the 1 bitcoin fraction of supply has been predictably reduced by monetary inflation from 0.02% to 0.002%.

Let us assume that Charlie held his 10 bitcoin for the week prior to his purchase. Charlie could have purchased the acre for 1 bitcoin but a week later it cost him 10. It would appear that over the span of one week Charlie has lost 90% of the purchasing power of his 10 bitcoin. In this example, a 10x increase in monetary supply correlates to a 10x increase in price over the same period (no causation implied). So has Charlie actually lost purchasing power?

Charlie had to first *obtain* the 10 bitcoin. He knew before obtaining his 10 bitcoin that this fraction of all bitcoin would be reduced from 0.2% to 0.02%, or that in a week his 10 bitcoin would be diluted to the current market price of the acre as measured in fraction of all bitcoin. Both Charlie and the seller of his 10 bitcoin discounted the trade for the net present value of the coin.

Charlie suffered no loss of purchasing power in his trade of 10 bitcoin for the acre, because he obtained the 10 bitcoin with full knowledge of its depreciation schedule. Loss would imply that he did not get what he paid for. **It is not possible to misjudge dilution caused by predictable increase in supply.** For example, predictable dilution is *always* used in computing the net present value of stock holdings. A shareholder must assume that option holders will exercise) at a net positive value upon liquidation. Failing to reach this conclusion is simply a computational error, not a loss caused by the exercise.

Commodity currencies such as gold exhibit monetary inflation in relation to exchange price. As price increases competition creates an offsetting increase in supply, keeping return on capital generally consistent with time preference. Bitcoin however exhibits an increase in difficulty in relation to exchange price. Bitcoin price has no effect on supply whatsoever.

Fiat currencies exhibit unpredictable and often secret monetary inflation. The purpose of such moneys is seigniorage, which is a tax on its holders. Unpredictability is important in maximizing the tax as otherwise the new money is immediately discounted to its net present value. Eventually price inflation discounts the money anyway, but not until the sovereign has reaped the benefit. When price inflation becomes fairly predictable, indexation can largely nullify its loss effect. This phenomenon was observed in the hyperinflation of the first Israeli Shekel.

> The linkage system was very successful. In major economies around the world, consumers often feel the pinch of just 2-7% annual inflation. But Israelis, who had to deal with a much higher inflation rate, went about their business practically unaffected. For three and a half decades, their real income was protected by this index-linked mechanism.

However *monetary inflation* can be perfectly predictable. When it is so, and the schedule is public, the effect is entirely offset by capitalization. This computation is orthogonal to demand, which has an independent effect on price. Each exchange includes a discount for net present value (NPV). With positive monetary inflation the NPV of 1 unit is a greater than 1 future unit. As monetary inflation decreases, the NPV of 1 unit increases relative to 1 future unit. With 0% monetary inflation the NPV of 1 unit is the same as 1 future unit.

The benefit of 0% monetary inflation over a perfectly predictable non-zero level is that it simplifies computation. Despite the existence of subsidy, saving does not compensate mining. Subsidy captures discounted purchase price, with the effect of a

transaction fee. Bitcoin inflation exists to rationally allocate units and is otherwise unnecessary complexity. Given that the cap is arbitrary, subsidy is phased out over time.

So unlike victims of seigniorage, Charlie did not suffer an actual loss of purchasing power even if over time it took more units to purchase the same acre of land. The increase in units is a predictable function of the positive (despite decreasing) discount rate. It is neither a tax nor an outsized return on mining, it is simply a mathematical conversion.

Finally, due to perfect predictability, it is not possible for Bitcoin monetary inflation to cause malinvestment or business cycles.

Jurisdictional Arbitrage Fallacy

There is a theory that, since it is unlikely that all states would join a Bitcoin ban, the coin would survive by movement of mining and other activity to such states.

Those who do not comply operate in the black market from the perspective of the banning authority. Another state in violation of a ban is considered a rogue state from this perspective. A ban is a simple political action against which Bitcoin offers no protection.

There is a related fallacy that such an action would be impossibly difficult in the case where Bitcoin is popular. This is the idea that Bitcoin is secured by the vote, which reduces its security model to that of the status quo of state money, eliminating Bitcoin's value proposition.

White market operations are by definition eliminated by a ban. The theory therefore implies that Bitcoin is ultimately secured by the protection of rogue states. This also reduces to security by vote. Furthermore powerful states have many tools to compel others, everything up to and including open warfare. These tools are commonly employed in various wars, such as those on drugs, money laundering and terror. A Bitcoin ban could easily fall under the umbrella justifications for all of these existing international conflicts.

However, Bitcoin is specifically designed to operate without permission from any state. Its continued operation as a black market money may lead one or more states to attempt its suppression through censorship. While this may be attempted by a single state, it is common for states to collaborate in defense of the taxing power of their monies. This is the purpose of the International Monetary Fund.

Such an action can be executed most efficiently from a single geographical location. In this scenario rogue states offer no defense except to the extent that they are not only willing to forego the tax benefit of their own monies, but also to donate tax money to resist censorship. **It cannot be assumed that rogue states can overpower the censoring authority, and any dependence of them reduces Bitcoin to a politically-secured money.** As such the theory is invalid.

Lunar Fallacy

There is a theory that hoarding bitcoin guarantees to perpetual profit. The theory is based on the following economic laws.

- One money is better than two (Metcalfe's Law).
- Better money displaces other monies (Thiers' Law)).
- At fixed supply, price rises with demand (Law of Supply and Demand).
- Potential increase in demand is unbounded (trade is positive sum).

Hoarding is purely speculative, with all returns constituting profit or loss. The money is not lent to another for interest and so is always available for exchange, a benefit that offsets interest forgone.

A corollary to the theory is that no investment in production is required to profit from it. Capital is required for all production. Lenders (investors) earn interest in exchange for time without their capital. Production is the source of trade and therefore all economic activity results from investment. A hoard is defined by its lack of consumption in production. If all people hoarded their capital, there would be nothing to trade and therefore no demand for the money. It seems that the theory is irrational, supporting the idea that Bitcoin is indeed Magic Internet Money.

The theory fails to incorporate the Stability Property of Bitcoin and is therefore invalid.

Network Effect Fallacy

There is a theory that the utility created by an economy varies with the square of the number of its merchants, assuming each merchant offers the same value of goods or services for sale in the one coin. The theory is an application of Metcalfe's Law.

This implies that an even split of the economy reduces combined utility by half. For example, if 20 merchants has utility 400 then 2 networks of 10 of these merchants has utility 200.

However, the ability to exchange any units of one coin for the other collapses the utility of the two economies into a hybrid economy. Due to the conversion cost the hybrid coin has lower utility than would a single, but this cannot be comparable to loss of one of the two entirely unless the conversion cost is unbounded. The theory is therefore invalid.

Private Key Fallacy

Private keys do not secure Bitcoin, they secure units of Bitcoin. **Private key control applies to individual security, not system security.** Whoever controls keys is the owner, and Bitcoin provides security for that owner, even if the keys are stolen. Decentralized validation secures consensus and distributed majority hash power secures confirmation, but private key security is the owner's problem.

Proof of Stake Fallacy

Confirmation security requires a person of authority to order transactions. Bitcoin periodically assigns this authority to the miner who produces the greatest proof of work. All forms of work necessarily reduce to energy consumption. It is essential that such proof be independent of the chain history. We can refer to this as “external” proof.

The only other source of ordering authority is therefore dependent upon chain history, which we can refer to as “internal”. There is a theory that such proof-of-stake (PoS) constitutes a comparable alternative to proof-of-work (PoW) in terms of confirmation security. It is true that both PoS and PoW delegate control over transaction ordering to a person in control of the largest pool of certain capital.

The distinction is in the deployability of the capital. PoW excludes capital that cannot be converted to work, while PoS excludes all capital that cannot acquire units of the coin. This difference has a material consequence for security.

In Other Means Principle it is shown that censorship resistance depends on people paying miners to overpower the censor. **Overcoming censorship is not possible in a PoS system, as the censor has acquired majority stake and cannot be unseated.** As such PoS systems are not censorship-resistant and the theory is therefore invalid.

Proof of Work Fallacy

Merchants purchase mining services that meet their rules for a satisfactory fee. There is a theory that mining services are subservient in this trade. This subservience is sometimes described as "asymmetry" or "users rule". This theory leads people to believe that mining can be strongly pooled as long as merchants are not centralized, as the economy can control the behavior of mining, rendering the system secure. The consequence of this invalid theory is complacency regarding the insecurity caused by pooling.

Miners control transaction selection, while merchants control property offered in exchange. If some part of the economy is unsatisfied with the selections of miners it can offer its property for sale in a split coin with a different work rule that obsoletes all grinding hardware. This is typically described as a proof-of-work hard fork.

According to this theory miners then suffer a catastrophic loss due to the unrecoverable capital investment in highly-specialized hardware. The hard fork may include a difficulty adjustment, allowing confirmation to continue despite a presumed significant drop in hash rate. Due to the lower difficulty and a presumed lack of specialized hardware, more individuals are able to mine. This introduces new miners to the business and reduces pooling.

It has been said that this ability of the economy to foist a capital loss on its trading partners is an asymmetry unique in comparison to other markets. For example, an apple buying community cannot simply "destroy" the orchards of all of its suppliers. The theory **fails to understand that there is no asymmetry in trade**. If all apple buyers decide that they will buy no apples from existing orchards they certainly have that power. Similarly, the orchards have the option to not sell. Price is the continual resolution of this tension. This is exactly the same dynamic that exists in every market.

The theory also **fails to account for lack of identity**. It assumes that the capital loss will cause the exit of existing "bad" miners and the entry of new "good" miners. This is an unsupportable assumption. There is no reason to believe that existing miners will exit nor is there any reason to believe that new miners would not make the same decisions as previous miners given they are in the same business, assuming one could even tell the difference. At least in the apple scenario one knows from whom one is buying apples and can discriminate, this is not possible in Bitcoin.

The theory also **fails to account for the economics of mining**. There is an advantage to proximity that produces greater returns on capital for miners with greater hash power. Larger miners are therefore more profitable than small miners. Larger miners will therefore be better capitalized than their smaller competition. When the rule change occurs the miners that remain will be those who can afford to retool, which will be the largest.

It is irrational to assume that all miners will simply exit. Would we expect all apple growers to be replaced by new apple growers? In mining are not expertise, facilities, energy contracts, process, and non-specialized machinery important advantages over newcomers? Existing miners have an inherent advantage over their supposed replacements. This means they have greater access to capital. So not only do larger miners end up with less competition, all existing miners that remain have an advantage over any new miners.

The theory also **fails to consider that merchants need mining**. Mining is not replaced by splitting, and it retains complete control over transaction selection. So for example if the "bad" miners happen to be states that are attacking the coin, the state itself and co-opted miners will continue with the same disruption, at a lower energy cost. As other miners fail due to what is effectively a 100% tax, the attacker's energy cost continues to decrease. Mining services that are "good" for merchants cannot be produced by splitting.

Finally, the theory **fails to recognize actual consequences**. Based on the previous capital loss experienced by all miners for a given coin, all future miners of its replacement will insure against the likelihood of a similar event. They may self insure, but the increased cost is unavoidable. This will reduce hash rate for the same fee until the possibility of such an event is deemed negligible. So the economy raises its own fees and ends up with the same miners and greater pooling. This is a reduction of security on two levels, with no benefit.

Relay Fallacy

The peer-to-peer network disseminates blocks and unconfirmed transactions. The protocol itself allows nodes to protect against denial of service. Consequently this communication requires no identity. This protection is how the network avoids the need for permission to participate.

However this protection comes at a cost in terms of announcement latency, and because of the advantage to proximity, lower latency translates into higher apparent hash power. Therefore miners compete for reduced latency. One way to reduce latency is pooling, another is to use a more efficient dissemination network. Given that pooling surrenders power to the operator, presumably the latter option is preferable.

One way to improve dissemination is to optimize the peer-to-peer network. The other is to join a distinct network, called a relay, that has lower latency due to elimination of denial-of-service protections, for example:

> [T]he cmpctblock message format was designed to ensure it fits neatly into a UDP-FEC-based relay mechanism. The only difference is that we send it over UDP with FEC... This way, extra hops do not introduce more latency. Sadly, due to the nature of our FEC encoding, we cannot know if individual packets are a part of a legitimate, or any, block, and thus only enable this optimization between nodes run by the same group. - bitcoinfibre.org.

The relay accepts communication from a set of miners, over the peer-to-peer or other protocol. The relay consists of a set of machines under the control of the relayer. It communicates the announcements within its internal network and eventually to the joined miners.

The important security observation is that communication within the relay is under the relayer's control. Due to the removal of denial-of-service protections *central control is necessary* to the scheme. The relayer can delay certain blocks based on miner, region, signal, non-payment, etc. A relayer **sells reduced latency**, and is therefore in the mining business. From a security standpoint it matters not if this service is offered for free. Miners may similarly offer grinders free reduced latency and variance.

Relays are aggregations of miners and miners are aggregations of grinders. The greater the hash power aggregation, the more profitable is the mine, as is the relay. One may consider that grinders are free to leave mines and miners are free to leave relays, and it is of course possible for a grinder to run his own mine and his own relay. But larger aggregations are more profitable, so leaving the largest relay or mine increases relative cost.

A theory holds that relays reduce pooling pressure. This is a fallacy as **any pooling reduction caused by a relay does not disappear but is *transferred* to the relay as a pooling increase.** Relay statistics are never presented alongside mining statistics, masking the power transfer. This leads people to believe that mining is less strongly-pooled than is the case. The consequence is complacency regarding the insecurity caused by the actual level of pooling.

Replay Protection Fallacy

There is a theory that replay protection applied in a split chain increases the relative utility of the original chain. Replay protection is a rule designed relative to another chain and with a directional behavior. The protection makes transactions of the protected chain invalid on the other.

Even without protection it is possible for an owner to spend in a manner that prevents replay in one direction or the other, though there is a fee and/or complexity cost in doing so. A split may reduce, but not eliminate, this cost in one or both directions by activating rules that spends can *selectively* utilize. This is called opt-in, in contrast to mandatory, replay protection. Opt-in replay protection reduces but does not eliminate the cost whereas mandatory protection can eliminate the cost.

The replay of a spend onto another chain is non-dilutive. The common output can be spent on either chain with or without replay. **The only distinction provided by protection is that spends can always be *distinct* on each chain with no extra cost to the spender.** The supply in each chain remains unaffected by protection.

It is a curious misperception that one chain can somehow absorb the transactions of another in a split. All outputs of the common segment remain spendable on both chains. Relay protection only reduces the cost of spending them on the protected chain.

One might assume that the lack of protection makes an owner less likely to spend on the unprotected chain, thereby limiting supply and increasing exchange price. However this assumes demand is unaffected by what amounts to an increase in trading cost. If the owner is not trading because of an increased cost in doing so, then the utility of the coin is not increased but decreased.

The self-protection cost amounts to a one-time demurrage that persists until protection is applied to unprotected units, intentionally or otherwise. This cost is a discount to the utility of an unprotected chain in relation to the hypothetical same chain with protection. This implies *greater* utility of a protected chain relative to the unprotected chain against which it is split than would otherwise be the case. Therefore the theory is invalid.

Reserve Currency Fallacy

There is a theory that Bitcoin will eventually be held by nations as a reserve currency and that individuals will transact using national currencies "backed" by Bitcoin. The theory asserts that transaction volume is insufficient for its use as a consumer currency, but the ability to resist debasement makes Bitcoin an ideal reserve asset. Central banks and their authorized functionaries would issue a dependent currency while holding Bitcoin on reserve. Given that Bitcoin cannot be inflated, the litany of problems produced by state control of money would be resolved, ushering in a new era of prosperity. Transaction fees would be low while transaction volume would be limitless.

Let us consider the scenario as it unfolds. Bitcoin becomes a fairly widely utilized hard currency but struggles with low transaction volume, high fees and long confirmation times. After some unspecified events unfold, national banks end up holding large amounts of bitcoin (BTC) and issue Bitcoin Certificates (BC) to facilitate trade. An auditing process is set up whereby people can verify that the issued BC never exceeds BTC reserves. Legal tender laws are created, requiring people to accept BC as payment for all debts. People purchase BC with BTC so that they can pay taxes and buy stuff from white market retailers. Eventually most BTC is held in reserve.

This scenario should sound familiar, as it is how nation-states ended up with gold and people ended up with paper. The theory is invalid on multiple levels.

The ratio of issued BC to BTC in reserve cannot ever be effectively audited. Even if Bitcoin consensus rules somehow remain, there is no way to know how much BC has been issued, and there is no recourse if debasement is suspected. The member banks have to be trusted to account for their issues, and ultimately this means everyone trusts the central bank to not engage in easing. History demonstrates that this is unlikely, and nevertheless it is no improvement over current state moneys.

So why is it that a person cannot ever effectively audit (validate) BC, as had been possible with the BTC that it replaced? Because that would make BC indistinguishable from the BTC held in reserve. In other words the *reason* there is a difference between legal tender and reserve currency is to enable inflation of the currency in use (taxation) while holding a hard currency

in reserve (savings). States hold hard currency because even other states cannot debase it and issue soft currency so that it can be debased.

Furthermore, for Bitcoin to be a hard money, there must be an actual decentralized Bitcoin economy. Without individuals validating BTC received in exchange, there is nobody to refuse invalid BTC as it comes to be redefined by the state. In this case arbitrary inflation can easily be introduced to BTC as well, invalidating the theory.

It is an error to conflate a reserve currency with a settlement layer. Layering preserves the key security aspect of decentralization, which works against debasement, while "backing" is full abandonment of this protection. It is very possible for Bitcoin to be held by national treasuries, but not for it to be predominantly a reserve currency. People must trade with it for it to be secure.

Scarcity Fallacy

As an *absolute* concept, economic scarcity of a resource implies only that it is not available in limitless supply. Nevertheless, if no person demands it, the resource has no value. A scarce resource under demand becomes property. No degree of difficulty in producing the resource is implied.

Scarcity may also refer to the *relative* availabity of some property. For a given supply, increasing demand implies decreasing availability (increasing scarcity). However, increasing demand tends to increase supply, and thereby availability. Similarly, for a given demand, increasing supply implies increasing availability (decreasing scarcity). However increasing supply tends to decrease demand, and thereby availability. These negative feedbacks stabalize availability and correspondingly price.

A single coin has fixed supply. There is a theory that the fixed supply of Bitcoin is the source of its value. As with Bitcoin, there is a fixed supply of the Mona Lisa, only one is possible. The theory implies that this is the source of value for the famed work of art. However there are countless unique works of art with no demand, and therefore no value. **Bitcoin cannot increase in value only because of absolute scarcity.** To the contrary, it necessarily becomes more scarce as it becomes more highly valued.

An aspect of the theory is that Bitcoin's fixed supply is the source of its utility because it ensures non-increasing availability. However, this requires non-decreasing demand.

Bitcoin is unique in the realm of property in that the cost of transferring it inherently increases with demand to do so. Unlike the Mona Lisa, it is also subject to effective substitution. These forces necessarily create the negative demand feedback seen in property without fixed supply. Given that non-decreasing demand is not assured the theory is invalid. As is common with economic fallacies, the error stems in part from considering just one side of the supply-demand relation.

Another cause of the error is a misinterpretation of the behavior of commodity monies. Because of its lower prevalence on the surface of the Earth, gold has remained more portable in common scenarios than more prevalent materials such as iron and salt. However the portability of electronic money is independent of the number of units in existence. Apart from sufficient divisibilty, the total number of Bitcoin units is entirely arbitrary and therefore unrelated to its utility.

Another cause of the error is a misinterpretation of the behavior of state monies. Through counterfeit laws the state controls the supply of its money by restricting competition. It can therefore collect an inflation tax by unpredictably expanding supply. Without restricted competition supply would expand through market forces, in response to demand, eliminating the tax. In other words the money would behave as a prevalent commodity, with poor portability (at least until renumerated by the state). Poor portability is often an actual consequence of hyperinflation. Like Bitcoin, electronic state money cannot exhibit poor portability. As such its unit quantity is also irrelevant.

Low prevalence (or unit quantity) is not an important monetary property except as it pertains to portability. Scarcity is a function of both supply and demand and therefore cannot be inherent in a money, even with fixed supply. Both commodity money and Bitcoin eliminate the inflation tax, though commodity money is subject to the negative feedback of inflation and Bitcoin is subject to the negative feedback of fee pressure.

Selfish Mining Fallacy

The term "selfish mining" refers to a mining optimization. However, one academic paper frames the optimization as follows:

> Conventional wisdom asserts that the mining protocol is incentive-compatible and secure against colluding minority groups, that is, it incentivizes miners to follow the protocol as prescribed. We show that the Bitcoin mining protocol is not incentive-compatible.

This statement assumes a "prescribed Bitcoin mining protocol" that precludes withholding, which is a straw man. Bitcoin consensus rules are necessarily silent on the timing of announcements.

> We present an attack with which colluding miners obtain a revenue larger than their fair share.

This statement assumes a concept of "fair share" that is foreign to Bitcoin, another straw man. A miner is rewarded based on his/her blocks that reach maturity, not as a proportion of actual hash rate.

These straw men are explicitly attributed to "conventional wisdom". In other words the paper uses them to show that the conventional wisdom is incorrect. However, the paper errs in unconditionally declaring that this *unfair violation of the protocol* constitutes an attack:

> This attack can have significant consequences for Bitcoin: Rational miners will prefer to join the selfish miners, and the colluding group will increase in size until it becomes a majority. At this point, the Bitcoin system ceases to be a decentralized currency.

This is the source of the fallacy. It is not an attack for conventional wisdom to be incorrect, it is an error in the presumed conventional wisdom. Selfish mining implies that Bitcoin exhibits latency-based pooling pressure, though this is a well-established flaw. All pooling pressures tend to reduce the number of miners, exposing Bitcoin to attacks.

Optimizations are not attacks. Pooling increases the *opportunity* for attacks, but opportunity should not be conflated with action. The term "attack" implies theft. The Bitcoin whitepaper, for example, uses the term only to describe double-spend attempts.

Side Fee Fallacy

There is a theory that off-chain fees represent an individual incentive that works counter to system security (incentive incompatible). The theory holds that a merchant paying a miner "on the side" to confirm the merchant's transactions prevents other merchants' transactions from being confirmed, or that it raises the cost of those confirmations, giving advantage to those who accept such fees.

One impact of such arrangements is that an average *historical* fee rate cannot be determined through chain analysis. The apparent rate would be lower than the market rate. This could of course lead spenders to underestimate a sufficient fee. However there is no aspect of Bitcoin that requires future fees to equal some average of past fees. Estimation necessarily compensates, such as by ignoring "free" transactions in full blocks or using standard deviation to identify outliers. But fee estimation is just that, estimation. Actual fee levels are controlled by competition.

Another impact is that disparate relative fee levels can highlight certain transactions as being associated with such arrangements. This can contribute to taint of the merchant's transaction and/or the miner's coinbase. But given the arrangement is a choice made by the creators of these transactions, there is no privacy loss.

There is no impact on market fee rates or the ability of others to obtain confirmations. If the arrangement deviates from market rates then either the miner or the merchant is accepting an unnecessary loss. This is no different than the miner confirming transactions with below-market on-chain fees or the merchant overestimating on-chain fees, respectively. In any case there would be no harm to system security even if all fees were paid off chain.

Bitcoin provides a mechanism for on-chain fees so that a transaction can compensate *any* miner without the use of identity. It is a privacy-preserving convenience. **If miners and merchants prefer to weaken their own privacy by performing additional tasks, there is no basis to consider that undesirable.** This theory is therefore invalid.

Furthermore, unless the miner's hash power is 100%, the merchant must accept a delayed confirmation time inversely proportional to the miner's hash power. The side-fee is offered at the market rate since the miner will incur an opportunity cost otherwise.

There is a related theory that side fee arrangements constitute a pooling pressure. If fees paid are consistent with the market there can be no effect on pooling. Above market fees are a state subsidy, as we must treat the subsidy as not economically rational. Below market fees are a tax, as we must treat the loss as involuntary. These are distortions just like any other state subsidy/tax and are therefore not unique to side fees. As such the existence of side fees does not create a new pooling pressure beyond what exists with on-chain fees, and the theory is therefore invalid.

Stock to Flow Fallacy

Stock-to-Flow historically describes the relationship between capital and income, allowing a future capital level to be estimated from an expected income level. Later this elemental concept was applied to money supply generally.

The ratio of stock to flow is a measure of time. Given a higher ratio, stock will increase more slowly. There is a theory that money with a higher inherent stock-to-flow ratio will suffer less proportional monetary inflation than a money with a lower ratio. The theory holds that the higher ratio implies a “harder” money, defined as inherently more resistant to the effects of monetary inflation.

The theory fails to consider the source of flow rates. It necessarily assumes that the rate of production is simply a property of the substance. But production of anything occurs when the anticipated price makes production profitable. A greater profit potential results in more competition, accelerating supply increase. More people digging for gold increases its flow.

In other words, flow is a function of demand. An anticipated loss results in no production whatsoever. This lack of any flow is not inherent in the substance but a consequence of lack of demand. Given that both supply and demand determine flow, the theory is invalid. This long understood error is not an aspect of the elemental stock-to-flow concept, but a misapplication of it.

Given counterfeit laws, competition to produce state money is restricted, allowing control of supply by the state, independent of market forces. As with other monies, supply and demand are generally unpredictable. A state may “peg” its issue of reserve notes to another money, such as gold. This relation may even hold over many decades. In this case the stock-to-flow ratio would incorrectly indicate a “hardness” comparable to that of gold.

Given that the stock-to-flow of money is the inverted monetary inflation rate, its relationship with monetary inflation is tautological. It does not imply anything about future monetary inflation. It can be used to analyze historical relations, and to calculate future stock based on *assumed* future flow, but it cannot be used to *predict* future monetary inflation. Any statement that one speculation will be more profitable than another based on historical stock-to-flow ratios is an error.

Brand Arrogation

Bitcoin is a brand, not a chain. No person can control the brand. People will use it to describe one or more chains and splits as they evolve. This happens with all monies, including gold and oil which trade at different purities and qualities.

This is consistent with the declaration of Bitcoin, as it binds a set of concepts, not a set of rules, protocols or implementations. People with capital invested in particulars have an inherent desire for brand association, but there is no such thing as a "legitimate" claim to it.

Fedcoin Objectives

As implied by Value Proposition there are two aspects of Bitcoin that make it a target of state controls, both threats to tax revenue.

In combating Bitcoin the state may attempt to introduce a cosmetically similar money, which can be referred to as Fedcoin. This could be introduced as a split or alternative coin. The objective would be to preserve the superficial aspects of Bitcoin while eliminating its value proposition. This would protect tax revenues while propagandizing Fedcoin as a “safer” alternative to Bitcoin. Creation of Fedcoin is not itself relevant to Bitcoin, it is only the act of compelling its use that requires resistance.

The two distinctions from Bitcoin allow the state to arbitrarily create new units (seigniorage) and deny transfer (censorship). The seigniorage objective can be achieved by a hard fork that introduces one new consensus rule. This rule allows the introduction of new units in the case where the state has signed an inflationary transaction. The censorship objective can be achieved by a soft fork that precludes confirmation of transactions that lack state signature.

Preventing the state from compelling the use of these forks is the central purpose of Bitcoin system security. The economy guards against the hard fork and miners guard against the soft fork. The risks taken by these people preserve the value of the money relative to state-controlled alternatives.

Hearn Error

There is a theory that a state cannot ban popular things.

This implies that high transaction throughput allows effective defense against attacks and coercion. This in turn implies that Bitcoin can be secured by accepting the centralizing force of very high transaction throughput.

The theory is invalid, as it rests on a factual error. **It is evident that states actually *prefer* to ban popular things.** The following is a short list of commonly-banned popular things:

- Drugs
- Gambling
- Prostitution
- Religion
- Speech
- Assembly
- Trade
- Migration
- Weapons
- Labor
- Books
- Money

This error may arise from failure to accept the Axiom of Resistance while continuing to work in Bitcoin. This is likely to produce cognitive dissonance. The subsequent search for relief may lead one here. However the error eventually becomes undeniable, which may lead to a rage-quit.

Proof of Memory Facade

It has been proposed that a proof-of-memory (PoM) can replace some fraction of the proof-of-work (PoW) energy cost with hardware, even relying on existing memory devices. As shown in Energy Waste Fallacy, a constant level of security requires a constant ongoing expenditure. Therefore such a system would require a comparable level of hardware consumption to offset any reduction in energy cost. **In other words total energy consumption cannot be reduced, it can only be transferred to hardware manufacture, operation and disposal.**

In December 2017 the estimated annualized cost of energy consumed in Bitcoin mining was $1,628,000,000, based on the approximations of 32.56 terawatt hours consumed at an average $.05 per kilowatt hour energy cost. Contemporaneously this cost level equates to the consumption of 32,560,000 terabyte drives at an average price of $50 per drive. Utilization of existing underutilized memory reduces the unit cost and therefore comparably increases the size requirement.

It is worth considering the economic behavior of a theoretical system in which PoM is determined by an existing (cost free) fixed pool of memory with no expiration or operational costs. As the cost of mining is zero, rewards flow at no expense in proportion to memory (assuming no pooling pressures). Any increase in average fee increases this reward for memory. Capital invested is zero and therefore rate of return is perpetually infinite. Despite unbounded incentive, the assumption of zero expansion precludes competition. But since the proof is externalized, competition cannot actually be restricted. In an actual system hardware manufacture expands perpetually for a given fee level, and this expansion accelerates with fee level increases.

Proof-of-memory is equal to proof-of-work in terms of resource consumption and there is no reason to assume a reduced energy component of that cost. The hardware acts as a proof battery, representing energy provably consumed in its manufacture. This is a facade analogous to the "zero emission" battery-powered car.

Spam Misnomer

The term spam in computing originally referred to excessive Usenet cross-posting and later became a synonym for unwanted broadcast email. While there is no clear distinction between wanted and unwanted email, the messages carry identity, are not fungible, and do not carry payment for processing by the recipient. Bitcoin transactions by comparison are neccessarily anonymous, fungible and carry payment for processing.

While email spam detection is a subjective process, it is necessary due to the lack of payment for processing. This process is

facilitated by identity and lack of fungibility. By contrast, due to anonymity and the fungibility objective, there is no test possible for transaction legitimacy, and due to payment there is no need for it. In other words all valid transactions are equally legitimate and this does not subject nodes to denial of service. A proper name for a transaction with a low fee is “low fee transaction.”

High volume submission of redundant transactions is a typical denial of service problem that is independent of the transaction fee and can be carried out by any person, not limited to the spender. Non-redundant transactions that incorporate mutually-conflicting spends are not a denial of service risk, since they are either rejected as invalid or accepted due to a sufficient fee increment.

Split Credit Expansion Fallacy

There is a theory that the increase of monetary units, as in the case of a split or new coin, creates credit. This is an error that is presumably a consequence of assuming that credit expansion driven by state monetary expansion is a market force. This assumption fails to consider that hard money cannot produce seigniorage.

Seigniorage is a tax. The created monetary units do not represent new capital but instead the dilution of existing units by the state, transferring ownership of the capital that they represent to the sovereign. As this capital is put to use in the subsidy of lending by the state banking cartel, as discounted money and insurance, the cost of capital to the bank's customers is reduced.

This so-called credit expansion is not simply the result of fractional banking as a market force. It is the consequence of the state favoring debtors at the expense of savers. In a free market of banking, banks are simply investment funds. Investors on average obtain a market return on capital and suffer the risk of doing so. In state banking risk, and therefore capital, are rearranged according to political objectives.

Market credit expansion is an increase in the lending of capital, as opposed to its hoarding. Increased rates of lending are a consequence of reduced time preference, and reduce the cost of capital. It is impossible to show that creation of a split or new coin (or anything else) reduces time preference. As such it is an error to assume that these creations either increase the availability of capital or reduce its cost.

Split Speculator Dilemma

In the wake of a split an original coin owner is faced with the choice of retaining or selling units of the original and split chains.

As discussed in Dumping Fallacy there is no way to discourage the existence of one chain or the other by exchanging or hoarding) units of either. Therefore we consider this choice to be strictly a question of how to maximize the value of existing holdings following a split.

Given a position before the split, an owner is impacted by the increased cost of unit conversion, and replay protection as applicable. These are unavoidable future trading costs that reduce the net present value of the units. Therefore these factors are not relevant to the question.

The remaining considerations are based on the **assumption** that the combined coins will increase in price over the contemplated period of time.

The Consolidation Principle implies that the two similar coins must eventually consolidate, reducing to zero the value of one of them over time. If one happens to know which this will be, it is rational to sell it and buy the other. However, given that one may *not* know which coin will survive, there is a chance that the trade would sell the coin that succeeds for the one that fails, sacrificing *all* value in the original units. With no knowledge of the future, selling all or part of one for the other increases the potential reward in proportion to the increased risk. As such it is equally rational to hoard both, which preserves the assumptions that existed prior to the split.

Finally it should be emphasized that both chains could fail, with value consolidating to an independent chain, commodity or state money. This topic intends only to provide a rational decision framework based on assumptions that may not come to pass.

Maximalism Definition

Maximalism is a public relations effort to discourage the formation of substitutes for a given coin. To the extent this is successful it may benefit existing owners by restricting supply and subsequently elevating price. However as people fail to find close substitutes, activity moves to more distant ones. In the case of electronic payments this is generally state money.

Maximalism is distinct from shitcoin awareness in that it is characterized by promotion of one Bitcoin over all others. Proponents

often express the contradictory theory that no other coin could be competitive with their preferred coin. If this was the case there would be no reason to advocate for a single coin.

Shitcoin Definition

A shitcoin is any system that is not cryptodynamically secure yet purports to capture the value proposition of Bitcoin.

Shitcoins are presumed to be scams, though it remains possible for proponents be to well-intentioned yet ignorant of cryptodynamic principles. By way of example, proof-of-stake technologies are shitcoins.

While there may be implementations of Bitcoin that are more secure than others, these are matters of degree. No Bitcoin can be shown to be absolutely secure. As such the term is not reasonably applied to any Bitcoin. By way of example, proof-of-memory technologies may not be shitcoins (despite failure to achieve central objectives).

Glossary

Fundamentals

Person

A decision maker.

Machine

An instruction follower.

Agreement

Bitcoin

The set of principles that secure a Coin from the State.

> The term and principles are defined by Satoshi in “Bitcoin: A Peer-to-Peer Electronic Cash System”.

Consensus

An agreement among People.

> Also the set of people who participate in an agreement.

Coin

A Consensus regarding a mutually-acceptable medium for Trade.

> BTC is a Coin.

Consensus Rules

The set of constraints that define a Coin.

Rule

A subset of Consensus Rules.

Validity

Conformance to Consensus Rules.

Validation

The process of determining Validity.

Enforcement

The act of discarding Invalid data.

Objects

Unit

The atomic transferable value of a Coin.

> The Satoshi is the Bitcoin unit.

Transfer

The change of control over certain Units.

Transaction

A Valid record of Transfer.

Block

A Valid set of Transactions with Timestamp and Proof.

Chain

The Branch with the most cumulative Proof.

Transactions

Script

A set of Operations that authorizes Transfer.

Operation

An atomic declaration of intent.

Contract

A Script that expresses Transfer conditions.

> Public Key Script is an anachronism for this.

Endorsement

A Script that satisfies a Contract.

> Signature Script is an anachronism for this.

Point

A reference to an Output or Input.

Output

An explicit Transfer and a Contract.

Input

An Output Point and an Endorsement.

Previous Output

The Output to which an Input refers.

Locktime

An expression of earliest Transaction Validity.

Blocks

Timestamp

A declaration of the time of Block production.

Median Time Past

An average of preceding Block Timestamps.

Proof

Probabilistic evidence of the cost of Work performed.

Branch

A Valid sequence of Blocks.

Weak

A Branch with less cumulative Proof than another.

> Orphan is a misnomer for this.

Strong

A Branch with more cumulative Proof than another.

Sequence

Confirmation

Inclusion of a Transaction in a Block.

Unconfirmed

A Transaction that does not exist in a Block.

Transaction Pool

The set of Unconfirmed Transactions.

> Memory Pool is a misnomer for this.

Block Pool

The set of Weak Blocks.

> Orphan Pool is a misnomer for this.

Genesis

The first Block of all Branches of a Coin.

Depth

One more than the count of Blocks after a Confirmation.

Height

The count of preceding Blocks in a Branch.

Segment

A contiguous subset of a Branch.

Organization

An Announcement adding a Block to the Chain.

Period

The average time between Organizations.

Layering

Trade using a sequence of Unconfirmed Transactions that can be Settled by either party.

Settlement

Confirmation of layered transactions.

Money

Spend

The initial publication of a Transaction.

Double Spend

The Endorsement of the same Output Contract by distinct Spends.

Subsidy

The issuance of new Units to a Miner.

Inflation

The increase in Supply resulting from Subsidy.

> Also monetary inflation, not to be confused with price inflation.

Fee

An implicit Transfer to a Miner.

Reward

The sum of Subsidy and Fees for a Block.

Coinbase

A Transaction that Transfers a Reward.

Maturity

The Depth at which a Coinbase Output becomes Transferable.

Halving

A reduction in the Subsidy rate (by half).

Difficulty

The level of Proof required for Validity.

Adjustment

A change to Difficulty.

Cap

The limit to Supply over all time.

Price

A moving average of Exchange rates.

Capitalization

The product of Price and Supply.

Economics

Trade

A voluntary swap of property between two People.

Utility

The usefulness of certain property to a Person.

Value

The Utility of a number of Units to a Person.

Supply

The set of all issued Units.

Exchange

The Trade of Units for other property.

Price Inflation

The increase in average Exchange prices over time.

Hoard

To Own for future use.

> This is neither speculation nor investment.

Speculate

To Own in expectation of Price increase.

> Also to borrow in expectation of price decrease.

Lend

To Trade time without Units for property of greater Utility.

> Invest is an alias for this.

Borrow

To Trade time with Units for property of greater Utility to the Lender.

Interest

The rate of increase in Utility from Lending.

Profit

The return on Speculation.

> This excludes interest.

Loss

Failure of Investment to earn Interest.

> This is negative profit.

Volatility

Deviation in Price over time.

Market

The Trade in certain property.

Network

Communication

Conveyance of data between Machines.

Protocol

A set of Communication conventions.

Peer-to-Peer

A symmetrical Protocol.

Client-Server

An asymmetrical Protocol.

Latency

The delay inherent in Communication.

Partition

An inability of certain Nodes to Communicate.

Denial of Service

Using Communication to exploit Protocol or Implementation flaws that degrade performance.

> DoS is an acronym for this.

Components

Mine

A Tool that performs Work.

Grind

A Tool that performs Hashing.

Relay

A Tool that disseminates new Blocks.

Node

A Tool that performs Validation.

Wallet

A Tool that creates Transactions.

Tool

A set of Machine instructions.

Implementation

A specific Tool set.

Actors

Miner

A Person operating a Mine.

Grinder

A Person operating a Grind.

Relayer

A Person operating a Relay.

Merchant

A Person accepting Units in Trade.

> User is a common alias for this.

Owner

A Person controlling certain Units.

> Holder is a common alias for this.

Developer

A Person creating an Implementation.

Claimant

A Person who holds a claim on property under the control of a Custodian.

> Also a lien-holder, shareholder, lender, or depositor.

Custodian

A Person who controls the property of another

Mining

Work

The process of Block production.

Candidate

A potential Block with undetermined Proof.

Hash

An atomic computation to Prove Candidate Validity.

Hash Rate

The rate of Hashing.

Apparent Hash Power

A fraction of Blocks in a Chain Segment.

> Public estimates of miner hash power are based on this.

Majority Hash Power

A subset of Miners with sufficient Hash Power to execute a sustained Attack.

> 51% is a common approximation of sufficient power.

Optimization

A Tool change that reduces the cost of Mining.

Announcement

The first communication of a Block to another Person.

Withholding

The purposeful delay of Announcement.

Honest

A Miner who builds on the Blocks of others.

Selfish

A Miner who is not always Honest.

Variance

The varying frequency of achieving a Reward.

Decouple

A Mine that shares Reward with another to reduce Variance.

Deviations

Fork

A divergence in Consensus Rules.

Hard Fork

A Fork that implies a Split.

> Expansion of the set of potentially-valid blocks.

Soft Fork

A Fork that implies a Split unless Enforced by Majority Hash Power.

> Contraction of the set of potentially-valid blocks.

Split

A Coin bifurcation.

Reorganization

An Announcement promoting a Weak Branch to the Chain.

> Reorg is an abbreviation for this.

Stall

The lack of Height increase over time.

Activation

Starting to Enforce a new Rule.

Signal

A Miner indication via Block data of intent to Enforce a new Rule.

Privacy

Identity

The means to associate Communication with a Person.

Taint

Determination of Ownership.

Security

Power

The relative level of control of a Person over the Chain or Coin.

Economy

The set of all Merchants.

Economic Power

A fraction of all property offered in Exchange.

Hash Power

A fraction of the Hash Rate of all Mines.

Attack

Use of Hash Power to enable Double Spending.

> Confirmation prevention is a case of double-spend enabling.

Co-option

Use of aggression to control Hash Power.

Coercion

Use of aggression to compel Activation.

Distortion

Market aggression that skews the cost of Mining.

Variation

Differences in the resource cost of Mining.

Censorship

Subjective Confirmation.

State

A set of People that uses aggression in place of Trade.

> Typically operates with impunity within geographic limits.

Political

Pertaining to the actions of States.

Weakness

Aggregation

The tendency toward reduced participation in Mining or Validation.

> Implies pooling or centralization.

Pooling

The tendency toward few Miners, including consolidation by Relays.

> Collusion is a common alias for this.

Centralization

The tendency toward few Merchants.

> Merchants directly control validation.

Delegation

The tendency toward few Owners.

> Owners directly control spending.

Partitioning

The tendency toward persistent Partitions.

> Identity implies exclusion.

Correlation

The ability to Taint using statistical Chain analysis.