- DRAFT -

A Study of Decentralized Markets on the Zcash Blockchain

Kevin McCabe and Aleksander Psurek, George Mason University

This past summer (2021), the Computational Experimental Economics Laboratory (CEELab) at the Interdisciplinary Center for the Study of Economics (ICES) at George Mason University (GMU) hosted the fifth annual research experience for undergraduates. The research projects were chosen to model blockchain institutions using the Microeconomic Systems (MES) approach used in economic science. The sponsors for this workshop were the International Foundation for Research in Experimental Economics, the Agoric Foundation, and the Electric Coin Company (ECC). The ECC team consisted of three undergraduate students from Spain, China, and the United States together with team mentors Aleksander Psurek (graduate student) and Stephan Kunath (graduate student), and Professor Kevin McCabe director of the CEELab.

In their research, the ECC team used mTree, a Python library developed at the CEELab, to build a model of bitcoin-like cyber currencies. The mTree library allows economists to design a MES and study the system using agent-based simulations and human subject experiments. The MESframework was pioneered by Leonid Hurwicz, Stanley Reiter, and Roy Radner to develop a theory of decentralized incentive-compatible mechanism design. Vernon Smith then used the framework as a model for the design and implementation of economic experiments. In his 1982 and 1985 articles, Smith specified necessary conditions for controlling preferences and replicating economics experiments. The mTree library, implemented by Stephan Kunath as part of his Ph.D. dissertation, uses an Actor system to facilitate the implementation of a concurrent computational model. MES simulations allow us to compute the consequences of the model for given agent strategies, while MES experiments allow us to test the theory using cash motivated subjects.

For the summer project, the ECC team designed and implemented an MES model of a blockchain cyber currency based on a proof of work consensus. The MES consists of an environment, which defines the state of the system and the technologies for changing state, and one or more institutions, which define the rules for how agents can access the technologies to change state. The ECC team implemented a blockchain technology with the economic agents, users, and miners, who use the blockchain. The system's state is the current set of resources, including the blockchain currency and the initial ownership of these resources by agents. The agents are assumed to have preferences over the different states of the system and interact through the blockchain institution (protocol) and complementary institutions to exchange goods and services for cryptocurrency.

Theoretic Design

The definition of agent preferences is critical to the subsequent analysis of the performance of a MES. We decided to break the value of owning a cryptocurrency into two types of preferences. The first preference for valuation is in terms of the long-term public asset value of the cryptocurrency. This valuation assumes that sometime in the longer run, tokens are defined as a lottery with different probabilities of different potential future values of the tokens. Let $L_A = (y, \alpha)$ define a lottery on tokens to be paid sometime in the future. L_A is understood to mean a token holder will receive y for each token with probability α and 0 for each token with probability $1-\alpha$. The expected value and fundamental value of this lottery is $e_A = \alpha \cdot y$. If our token holder is risk-averse, then the expected utility of holding the token until the future date is given by $eu = \alpha \cdot u(y) + (1- \alpha)u(0)$, where u(x) is the utility of the monetary outcome x.

People may trade tokens speculating on capital gains. If a trader buys a token at price p_t and sells the token at price p_{t+1} then the capital gains on that token is simply $(p_{t+1} – p_t)$. Such speculative behavior is known to produce price bubbles in laboratory asset markets and may explain some of the price volatility in cyber currencies.

The second valuation mechanism is the indirect private value model, which assumes that the value of a medium of exchange is the value of its purchasing power. Specific cryptocurrencies would then have different potential indirect private values. A simple version of this model is Alfred Marshall's model of supply and demand. In this model, we can think of a token as offering a medium of exchange for a good that a buyer could pay a seller in exchange. We can use the fundamental value of the token as their store of value, allowing buyers and sellers to trade tokens, knowing they will be valuable in the future. Let V be the maximum a buyer is willing to pay (denominated in tokens) for good, and let C be the minimum amount (denominated in tokens) that the seller is willing to accept. Assuming no other transaction costs, the buyer and seller could agree on a price p (denominated in tokens) between V and C, with the buyer earning surplus $V – p$ and the seller earning surplus $p – C$ with the gains from exchange being $V – C$. Given this model, we can write a buyer's surplus, $v = V – p$, as the value of a transaction sending p tokens to the seller and the seller then completing the transaction. If we let β be the subjective belief that this will happen, then the buyers expected value $EV = \beta V – pf$, where f is the fee for making the transaction. Buyers will choose transactions that maximize the sum of the expected value of transactions.

Given these preferences, we can define the actual performance for the MES as the cumulative gains from exchange. We can then define market efficiency as actual performance divided by maximum gains from exchange that could have been achieved.

Institutions

The primary institution of interest is the blockchain protocol, which includes the consensus mechanism, the block rewards mechanism, and the transacting mechanism, which includes a history of all transactions on the chain. All these features are common to the basic blockchain infrastructure, but additionally, we model privacy, the type allowed by Zcash, and its shielded transactions. The economic value of such transactions is the protection of property rights from bad state actors and the protection from the strategic behavior of economic competitors who may use unencrypted blockchain transactions as strategic information.

Additional institutions required for our MES are a double auction and coin wrapping institution. The Double Auction provides a centralized exchange in which individual users and miners can buy or sell a cryptocurrency for a fiat currency. The coin wrapping institution provides a mechanism for wrapping foreign coins on a primary blockchain, permitting the formal simulation study of Zcash Shielded Assets (ZSAs) on the privacy-enabled chain. Of course, adding such assets requires support by the blockchain, so blockchain institutions supporting ZSA are also modeled.

Experiment

The main criticism of simulation modeling has to do with the behavior of individuals – modeling human behavior is necessarily a sterile and axiomatic exercise. For this reason, results from MES simulation should be viewed as a theoretical prediction, so to test the validity of this theoretic prediction, an experiment will be run with the simulated agents and repeated with authentic human subjects. Individuals will either be provided with assets (native coins) by the experimenter throughout the experiment for no cost or be allowed to mine (pay) for the coin assets. Individuals will also have access to the double auction institution to trade coins for fiat and be informed of a fundamental value for the coins, which will be redeemed after a predetermined number of trading periods. The experiment is in the vein of asset market experiments pioneered by Vernon Smith in 1988, following the Smith 2000 model of flat asset values. It presents an improvement over a similar human-subject experiment that attempts to study mining behavior, Lambrecht et al. 2021.

The experiment will allow us to test if the consequences of our theory assumptions are correct and study the possible reasons for the behavior and outcomes we observe if the theory is incorrect.

Economically Motivating Wrapped Assets (ZSAs)

As part of our research and learning on blockchains, we were asked to consider various economic mechanisms for supporting ZSAs, which are wrapped versions of foreign assets which are subsequently represented on the Zcash blockchain. This research aims to explore the space of economic mechanisms for imposing fees on ZSA holders that are fair and consistent with those paid by ZEC holders. In evaluating a mechanism design, we study the property rights supported by the mechanism and the miner and user incentives created by the rules of the mechanism. We do this in the spirit of building research hypotheses that require experimental testing, both in the laboratory and in the field, as we learn more about blockchain economics and private transactions. The discussion below should not be used as consulting advice but rather as contributing to the academic discussion of mechanism designs in blockchain economies.

Zcash Background and Advantages

Given the transparency of blockchain transactions, most cryptocurrencies reveal your full transaction history and holdings to everyone. We use Ronald Coase's model of transaction costs to help understand why interested parties may or may not engage in trade and propose the following model of user transactions. A transaction on the blockchain is a triple (s_id, r_id, q) where s_id is the sender, r_id is the receiver, and q is the quantity of currency sent from s_id to r_id. Moving forward, the model looks at sender incentives to make the transaction. The asymmetrical model can also be written down for receiver incentives. We do not provide an index for users to keep the model notation simple. Let v_k be the value to the user of making a specific transaction, denoted k, and c_k the cost of making the transaction. We assume these values and costs are evaluated in a domain of considered transactions, and we call v_k - C_k the gains from trade in making transaction k. In making the transaction, there is some uncertainty that the transaction will be completed as intended. We will model this as the subjective probability p_k, resulting in the expected gains from trade $e_k = p_k(v_k - C_k)$. We assume that users form a transaction plan by ordering potential transactions from the highest positive expected gains denoted e_1, …, e_T over time. Note that this model does not include discount rates but does allow for the dynamic removal and arrival of (C_k, p_k) opportunities that we assume are automatically ordered into the transaction plan. Note that sender plans are likely to be myopic about future opportunities, resulting in potential time inconsistencies.

Our main hypothesis is that a person can be hurt (at least economically) by being made a target by others when information about a transaction is made public. This harm occurs because of the actions of business competitors or bad actors, who take actions based on the information they see to exploit the users making the transaction. If we denote that harm by the constant C^I, we can model a user's subjective expected harm as $q_{k.}$ This results in the expected cost equation $C_K + q_k C^I$ where c_k is now interpreted as the direct cost of making transaction k, and $q_k C^I$ as the indirect cost incurred when others act on the public nature of the transaction. Note, in general, q_k will depend on the type of transaction, the identity of the parties to the transaction, and attributes of third parties who may act on this information. Note, it is also likely that some transactions will offer reputational benefits, which are added to the model as $r_k B$ where B is a constant benefit from public information, and r_k is the subjective probability of receiving this benefit. This results in an expected gain from exchange equal to

$e_k = p_k(v_k + r_k B - C_k - q_k C^I)$.

In our analysis, we assume either r_k is positive, and q_k is zero or q_k is positive, and r_k is zero, but the overall effect on e_k could simply handle mixtures.

The Zcash blockchain went live on October 28th, 2016, to help reduce the transaction costs of exchange. Using zero-knowledge proofs, Zcash allows users to decide to make transactions without revealing information about the transaction's sender, receiver, or amount. As a result, Zcash added privacy as a valuable property right on the settlement assurances already provided by Bitcoin. This means users can avoid informational costs when q_k is positive but gain information benefits when r_k is positive. This new property right increases the gains from exchange to users.

The Economics of ZSAs

The Zcash protocol in its current state (without ZSA support) provides strong support for privacy rights using zero-knowledge proofs. A mechanism for ZSA support would also need to support these privacy rights. The current Zcash protocol compensates miners through block rewards and small transaction fees. This rewards mechanism secures the chain by providing an incentive to mine honestly and incentivizes the inclusion of transactions, which has a positive, but small computational cost. This mechanism is currently working well, but adding ZSA support may increase the value of the blockchain and require additional mining rewards to maintain mining incentives. Suppose ZSAs become a major source of transactions in the future. In that case, a good transaction fee mechanism can continue to provide mining incentives, and other sources of revenue as future block rewards decrease and as the asset value of the chain increases.

Introducing a ZSA mechanism increases computational costs and block congestion, proportional to the number of ZSA transactions. The computational costs of adding transactions are negligible compared to the costs of finding the block hash and the externality costs associated with running a node. Present externality and mining costs appear sufficiently compensated by block rewards. Additionally, block congestion is not an issue as blocks are not full, but it could become a problem in the future.

The total asset value of the chain would also increase, proportional to the value of the ZSA assets. While the asset value may not appear as a cost, it can pose a security risk to the chain, as it increases the incentive for various forms of opportunistic behavior. Thus, there will be a need to increase compensation for mining in proportion to the value of the ZSA assets. If such a rewards structure is not adopted, ZSAs could pose a security risk, as documented in a paper by Tim Swanson in 2014, discussing "colored coins" on the Bitcoin blockchain. Swanson discusses in greater detail the danger of an imbalance of chain asset value to block reward value, and his conclusion is that if sufficient rewards are not provided to incentivize honest miners to secure the chain, opportunists will necessarily emerge.

Another cost imposed on ZEC holders, and by proxy of the block reward Zcash miners, would be the loss of relative advantage that ZEC has over other blockchain assets. Presently, ZEC is the only asset which can be used with the Zcash zero knowledge protocols, and thus is the only one with this specific form of privacy protection. If ZSAs are allowed on the chain, these assets would also have the Zcash privacy benefits. This is a clear benefit to ZSA users, but it would decrease the demand for ZEC as the only asset with this form of privacy. However, as ZEC users presumably value this privacy aspect, they are also likely to be users of ZSAs. If, for instance, a privacy-conscious user wanted to become diversified across different assets, holding a basket of ZSAs on top of ZEC would allow greater diversification without sacrificing privacy. Such a use case could increase the attractiveness of the Zcash blockchain and draw more attention to it and ZEC as the native asset. Furthermore, if fees on ZSAs are charged in ZEC, ZEC would again have a unique value proposition above ZSAs.

It is unclear if adding ZSA support will increase or decrease the value of the Zcash protocol. Suppose a user needs to make a private transaction. In that case, they can do this by buying ZECs without wrapping other coins. Users are more likely to trust the protocol if it maintains stability in the property right structure and only makes incremental refinement of the software and protocol rules for purposes of efficiency. If ZSA support is added, it will likely require additional support for distributed exchanges (DEXs) to allow users to trade securities. Indeed, one of the transaction fee mechanisms proposed would require minimal DEX support. Adding DEX support will add additional property rights to the system but may undermine user confidence in the consistency and implementation of these property rights. This will require a well-thought-out, user-involved, and user-incentivized testing phase, which is likely to be both time-consuming and costly.

We can think of four reasons/scenarios for adding ZSA support. The first and best reason is that the addition of ZSA will foster innovations in the way the Zcash blockchain and protocol are used. A second reason may be an increased demand for Zcash transactions using wrapped versions of some other medium of exchange - stablecoins in particular. The third reason for ZSA support is as a stepping stone to future projects, including making available decentralized exchanges (DEXs). The fourth reason would be to provide a support floor for the Zcash ecosystem in the case that ZEC started losing significant value and users wanted to use alternative wrapped coins while making privacy transactions, supporting the core ZEC users with ZSA fees. There might be other scenarios that would emerge from innovative experiments with ZSAs if they were made available.

Assessment of Proposed Mechanisms

This section aims to assess economic models of transaction fees that may be employed for the ZSAs. In general, such fee mechanisms change the cost, C_k, of a transaction. ECC proposed the list of mechanisms we are looking at. We can look at each mechanism proposal using the following criteria:

1. Does the mechanism impose costs on ZSA holders that are fair and consistent to current ZEC holders?
2. Does the mechanism extend Zcash property rights on ZECs to ZSAs?
3. Does the mechanism reduce the security concerns that emerge by adding ZSA support to the blockchain?
4. Does the mechanism provide revenue to cover additional protocol costs and Zcash core support revenues?
5. How does the mechanism incentivize user trust in the blockchain and, subsequently, the transactions they make?
6. How does the mechanism incentivize miner behavior?

The analysis we provide below is preliminary based on our understanding of the microeconomic system underlying privacy transactions. We would be interested in feedback as we continue to study these interesting mechanism design problems.

Computational-Cost Transaction Fees

This mechanism would simply have ZSA transactions pay a fee in ZEC to compete with other pending transactions for inclusion in a block.

This mechanism is identical to the current Zcash protocol, so it would be nearly trivial to implement and the easiest to understand by users. It is therefore the most likely to foster the use of ZSA's on the blockchain. Whether or not the protocol is fair to current ZEC holders is less obvious. On the one hand ZSA users would have access to the same privacy capabilities of the Zcash blockchain, but on the other hand they would not face the inherent currency dilution from block reward issuance. Conversely, they would have to purchase ZEC to pay fees, driving up ZEC prices, but not experiencing capital gains from holding ZEC. This buying pressure would also increase the liquidity of ZEC. Additionally, expanding the capabilities of the Zcash chain could increase the amount of interest in the chain and therefore increase the value of ZEC. On balance, these benefits to ZEC holders may be enough to create fairness, but this is a subjective valuation. If ZSAs prove to be very popular and impose significant costs on the chain, it seems this mechanism would indeed be unfair.

Currently, there is an excess capacity in blocks that are limited in size and production rates. If these capacity constraints are reached, fees will likely be bid up by users. Still, there is no guarantee that the fee will cover escalating externality costs. In this case the protocol might have to price this additional externality using a Pigouvian tax or similar scheme. Such a fee levied on ZSA transactions will help to further incentivize miners to secure the network.

It is worth noting that this mechanism has potentially weak privacy features, since observers will be able to see the fee a transaction was willing to pay. While this is a feature common to current Zcash mainnet transactions, it suggests that the mechanism entails a tradeoff between privacy and expediency. A simple solution to this privacy problem would be to mandate the same fees across all transactions, but this would prohibit individuals from competing for execution speed.

Shielded Hosting Fees

The idea behind this mechanism is to have ZSA holders pay a fee equivalent to the dilution of ZEC holders' equity by mining rewards. In this mechanism, a cumulative hosting fee is set for each ZSA. Since the fee is cumulative across blocks, the longer a user waits to transact, the more is lost in cumulative fees. The advantage to this mechanism is that it provides strong fairness between ZEC holders and ZSA holders, assuming that the value of ZSAs can be well determined in terms of ZECs. This may not be the case since different ZSA's, such as wrapped stable coins, may trade at different valuations relative to their unwrapped counterparts.

This mechanism can better control the security risk of increased network value by directly tying the increase in value to miner rewards. It also succeeds in its efforts to maintain the privacy rights enjoyed by ZEC holders. However, this mechanism has two potential problems: it creates an incentive for users and miners to redeem ZSAs for the underlying assets as quickly as possible and it creates additional overhead for miners as they are now being paid in many different assets.

The incentive created to dispose of ZSA assets as quickly as possible is a direct result of the cumulative hosting fee relative to the amount of time a public key has owned the asset. Absent direct transaction fees or mint/redeem costs pertaining to the ZSAs, users, and miners have an incentive to quickly redeem their ZSAs. This incentive creates several issues for the mechanism. First, it lowers the miner incentive from ZSAs, as additional block rewards are based on the total value of ZSAs on-chain. Second, the increase in mint/redeem transactions would result in a huge amount of uncompensated computation costs, increasing block congestion. Finally, the quick redemption of ZSAs would decrease the privacy function of the blockchain as it makes it easier to link transactions based on timing to transactions in the mint/redeem mechanism, assumed to be public, revealing a large amount of the previously private information.

Another major issue with this mechanism is the incentive for miners to censor ZSA transactions – a miner should, in equilibrium, never include a ZSA transaction in a block unless the originator of the transaction is the miner itself. The reason why rational miners are expected to behave like this is simple – a miner receives block rewards relative to the total amount of ZSA on the chain. However, they do not receive any reward directly for including a ZSA transaction in the block, as the transaction's "fees" are being collected implicitly through the ZSA block rewards. There is an exception to this behavior – the miner wants itself to transact, especially redeem, ZSAs. If mining is sufficiently concentrated, this will simply incentivize the miner to censor all non-self ZSA transactions but include its transactions in its blocks, expecting to win blocks at a reasonable rate. If mining is sufficiently diffuse, a social norm may emerge, where miners do not censor ZSA transactions to induce other miners to not censor ZSA transactions.

In summary, the ZSA hosting fee cannot be so large as to disincentivize ZSA users but large enough to incentivize miners. As miners in this mechanism are paid in ZSAs directly as their mining block rewards, this creates additional overhead for miners. Even if all the management of these additional block rewards is handled by software, it creates additional computing costs to miners and additional cognitive costs as miners are forced to consider their strategy for every ZSA they are paid.

Shielded Hosting Fees + Mandatory DEX Market Buys of ZEC

This mechanism is similar to the Shielded Hosting Fees mechanism, but fees are automatically converted at some rate to ZEC using an internal Zcash DEX. This mechanism will extend the strong privacy rights enjoyed by ZEC holders to ZSA holders with many but not all the same incentive problems. The incentive to mint/redeem ZSA assets quickly will remain. The fact that miners will now be paid in ZEC will make miner costs more manageable and produce a potential revenue stream over time as new ZEC production slows down.

The use of a DEX introduces a new problem for ZEC to ZSA pricing. First, the DEX will have to set a willingness to pay ZSA for ZEC and a willingness to accept ZEC for ZSA for each ZSA minted in the market. From an economics viewpoint, this will result in an incomplete set of markets unless DEXs can be made that trade one type of ZSA for another. Given the incompleteness of the DEX markets, the ZEC to ZSA DEXs will have to rely on price oracles, ZEC escrow reserves, i.e., liquidity providers in standard AMMs, and by inviting arbitrage behavior. The results from arbitrage in incomplete markets is an empirical or possibly an experimental question, requiring further study. An argument often proposed why DEXs would present the proper price and thus be not gameable is the no-arbitrage principle, and it would apply here, too, if the DEX were to allow unlimited buying and selling. DEXs are constantly arbitraged by bots, especially in thinly traded pools, so one could expect a similar result here, with the exception that privacy would make arbitrage decisions more difficult. Over time, the price will therefore reflect the true value of the asset, on average and within the margin of error provided by transaction costs which will likely be higher with incomplete markets. Still, the fact remains that miners, having to sell into the DEX at a protocol-determined rate and time, will suffer from being too predictable and thus being on the wrong end of arbitrage on average.

A second issue with mandatory DEX exchange based on accumulated fees is that other DEX traders will anticipate these market events and gamed to cause lower ZEC outcomes for miners. This may be a reason to pay miners in ZSA fees and allow them to make their own DEX timing decision. It is unclear if a DEX could be implemented that allowed unlimited buying and selling and miner timing decisions while preserving the privacy property right of the Zcash protocol.

Private Transaction-Volume-Proportional Fees

In this mechanism, fees in ZEC are made proportional to the amount of ZSA being sent in a transaction. A price oracle is used to determine the exchange rate of ZSAs for ZECs. The sender would have to include both the ZEC fee and the quantity of ZSA being sent in the transaction. The proportionality satisfies a non-decomposability result making costs the same if a single ZSA transaction is broken up into smaller ZSA transactions. Since the ZEC are burned, this mechanism will extend the privacy rights of ZEC holders to ZSA holders.

On the transaction demand side, this strategy fails to account for the differential demand and differential cost of multiple transactions totaling the same volume of transacted ZSAs. If one sends two transactions of 10 ZSAs or one transaction of 20 ZSAs, they will pay the same amount. While the cost to process the transactions is negligible relative to the cost of hashing a block, this could lead to needless congestion. If the user gets more utility out of sending the two transactions as opposed to the one transaction, additional fees would be justified. This mechanism fails on the security compensation side as it only creates incentives for miners when ZSAs are transacted. If a user placed a huge volume of ZSAs into their wallet at some point and then never transacted these ZSAs, the increased risk to the system would still be present, but there would be no increased incentive to miners to provide security against this risk.

The mechanism as proposed does not compensate miners except through the indirect effect of removing ZEC from circulation. Since burning ZEC is likely to increase the value of the remaining ZEC, this mechanism constructs a public good inviting free-riding behaviors. Furthermore, as the transactions are shielded and the removal of ZEC from circulation is unknown, it will be harder for markets to determine the actual value of ZEC, thus making the price oracle less certain.

The introduction of an oracle into the system is highly complex from an economic perspective, as it creates a host of new problems. The main question is how the oracle will price assets? One solution to this problem is simply to use an external exchange, say a DEX like Uniswap, or a centralized exchange like Coinbase. As the price oracle must be specified in the code, this creates a new vector of attack, as it allows a way to manipulate the rewards mechanism by manipulating the source of the oracle's prices. This problem is especially pronounced when using a single-source oracle or an oracle that relies on information from a thinly traded asset. Moreover, suppose the asset traded on the oracle is the underlying asset and not the ZSA representation of the asset. In that case, the price will not reflect the true valuation of the ZSA.

Transparent Direct Transaction-Volume-Proportional Fees

A proportional fee, denominated in ZSA units, is paid and used as a block reward to miners. This mechanism fails to extend the privacy property right to ZSAs since the asset type is visible and the proportion is known. Thus, the amount being transferred is trivial to infer, even if otherwise obscured. In essence, it makes ZSAs entirely pointless outside of a Defi ecosystem on Zcash, as the lack of privacy makes the use of this mechanism equivalent to transacting on another chain and does not provide the unique value proposition that Zcash seeks to confer.

Furthermore, this mechanism suffers from the same "dust" problem as the "Shielded Hosting Fees" mechanism. It forces miners to accept payments in many different ZSAs, likely in small amounts per transaction. Miners may respond to this effect by censoring transactions using ZSAs if they are sufficiently small – made possible by the ZSA transaction fees being entirely transparent. This effect will aggregate up, with miners censoring any ZSA transaction with a fee less than that which can be unwrapped and spent as the underlying asset, factoring in all unwrapping and intermediary transaction fees. Thus, the minimum spendable cache of a ZSA would induce such a high transaction fee. This issue is a problem for all ZSA mechanisms which pay fees in ZSA.

Scarce Active Asset Type Auctions

In this mechanism, there is an interval during which a limited number of asset types are allowed to be “active”, i.e., allowed to transact. Parties interested in ensuring a given asset type will be active post-ZEC-denominated bids in support of activating a specific asset type in the upcoming interval. When the interval begins, the asset types with the largest aggregate posted bids become active, and the posted bid amounts are paid out in subsequent block rewards during the interval. Bids posted for asset types that did not become active are returned to bidders.

This auction is proposed as a first-price auction where the bidder's bids are added up to become a final bid, and the asset types with the N highest winning bids are then allowed to transact with the winning bidders paying what they bid. This amounts to paying for a public good giving bidders an incentive to free-ride on other bidders' behavior. From experiments, we know this problem can be solved in part by using a continuous contributions mechanism where bidders can see the total contributions to the mechanisms and make incremental bids to increase these contributions. Other pricing mechanisms based on the second-price or uniform-price auction can also be considered more difficult if winning bidders for a winning asset type have different bids.

A second problem for such a design is the interval for active assets and the number, N, of active assets. This is both an empirical question based on the demand for ZSAs, and a security question based on security audits of the blockchain.

Scarce Asset Type Creation Rate Auctions

In contrast to the previous mechanism, the permanence of ZSAs in this setting would allow their continual use and not incentivize their rapid removal from the chain. However, the mechanism does present other problems as it would create an ever-expanding set of possible ZSAs, leading to increased computation costs. At the same time, the payments for the inclusion of ZSAs would only occur once at the beginning. This will present a problem to miners, as they may not be fairly compensated for the increased computation cost, and there will be no lasting security incentives.

An alternative to this auction mechanism would allow assets to stay active until they are beaten by bids on another asset. In this case, the winning bids would be depreciated at a given rate (perhaps proportional to block payments). New assets will be added if they have higher aggregate bids than the depreciated bids of currently active assets. In this type of auction, bidders can add bids to the depreciated bid for a currently winning asset type to keep that asset type active.

Supply-Proportional Issuance Bonds

In this design, an issuer may not mint new units of a given ZSA unless they have posted a price-proportional amount of ZEC to a bonding system. Once the asset supply decreases due to redemption, the issuer can then withdraw the excess ZEC held in the bond. This rule scales the demand for ZEC with the demand for ZSAs. The bonding and issuance system is public, but transactions can be private. Bonded ZEC is not burned but temporarily removed from the supply, providing a transient benefit to holders of liquid ZEC. This will require an escrow mechanism.

Since the bonding is price-proportional, it introduces many of the same risks as using a price oracle, explored in the “Private Transaction-Volume-Proportional Fees." However, the value of bonds would be especially vulnerable to manipulations of oracle prices, as bond liquidation events may cause unintended losses. The bonds would need to be overcollateralized to account for the fluctuations in the value of the underlying assets.

Finally, the bonds themselves become assets and would carry specific property rights. Given the decentralized nature of the blockchain, these would need to be transferable claims on the underlying ZEC held in-bond unless the issuance of ZSA was permissioned and the bonds were therefore non-transferable. If the bonds were transferable, this would allow for greater flexibility for ZSA issuers but could create speculation. If the bonds were tradeable or not, they would either require a seigniorage or an interest rate charge, using the profits to fund security subsidies for miners. Otherwise, this mechanism lacks any incentive for miners, compensating them for the increased security risk to the chain.

Appendix - ZSA Fee Mechanisms

Design Constraints and Assumptions

These mechanisms all assume that the creation of a new asset type and issuance of new units of an asset is public. By implication an upper bound on the supply of each asset type is also public. This is an upper bound because users can always privately destroy units of an asset.

These mechanisms all assume that the authority to issue new units of a given asset type is constrained cryptographically to an abstract “issuer” of that type which may be an institutional issuer or some kind of decentralized issuer or bridge protocol.

A “price oracle” is some mechanism to provide exchange rate information between each ZSA asset type and ZEC.

A “DEX” in this document refers to an “Automated Market Maker Decentralized Exchange Protocol” which allows the equivalent of immediate market buys between two asset types based on previously provided liquidity. The assets providing liquidity and the exchange transactions must be directly managed by the protocol.

Mechanism Overview

Given the target use case of supporting only fungible assets, there are a few axes in technical design common to all of them:

- **Payers:**
    - Every fee mechanism is explicitly paid by some party.
- **Beneficiaries:**
    - Some mechanisms pay fees directly to recipients, whereas others have indirect beneficiaries. A common kind of indirect beneficiary arises when fees are either temporarily or permanently removed from the supply, which indirectly benefits the remaining holders of that supply.
- **Economic Proportionality:**
    - The costs and benefits to different parties of each mechanism are proportional to different aspects of the overall system
- **Privacy Spectrum:**
    - Different mechanisms may have different privacy impacts. The stronger the privacy of a mechanism, the more all participants benefit from the network effects of privacy. For example, if the type of assets involved in a transaction are publicly known, then users of

rarely used assets will have less privacy benefit. If, on the other hand, all users of any asset rarity contribute to the same privacy set, all users benefit from any usage.

- **Technical Requirements:**
  - Different mechanisms may require more or less technical sophistication or prerequisites which impact the design complexity.

These mechanisms are presented in a standalone fashion, while a system design in practice can also combine these.

Mechanisms

Computational-Cost Transaction Fees

Payer: transaction senders
Beneficiary: block producers (currently: miners, could be generalized to block reward recipients)
Proportional to: computation cost of transaction inclusion
Privacy: poor - the fee is public, so observers can distinguish the importance of transactions.
Technical Requirements: No new requirements; already active.

This is the mechanism currently active in Zcash.

In this mechanism, transaction senders pay an in-band ZEC denominated fee to compete with all other pending transactions for inclusion in new blocks. (Blocks append groups of transactions to the consensus ledger.) Adding a transaction into a block imposes costs on miners, who produce blocks, so a rational miner should only include transactions whose fees are above this cost.

Blocks are limited in size and production rate, so if more transactions are sent than the network can process, competition for inclusion climbs and transaction senders escalate the fees they pay according to their demand.

Since there's no new computation costs associated with ZSAs, the fees would be identical to a ZEC transfer.

Shielded Hosting Fees

Payer: asset holders
Beneficiary: block reward recipients (miners and Dev Fund recipients)
Proportional to: the value of assets held on Zcash over time
Privacy: strong - every holders pays the same rate, and there is no distinction between rates or asset types
Technical Requirements: Zero-knowledge circuit and surrounding consensus rule support

Hosting fees are charged at a near-continuous rate proportional to the amount held from all holders of all asset types to block reward recipients. The value is proportional to total assets held on the entire Zcash network.

The value is not correlated with the cost of including transactions in blocks, so users who transact frequently pay the same rate as users who rarely actively use the network. Additionally, block producers (miners) do not directly benefit from individual transactions, and have no direct incentive to include transactions.

Note: in the basic technical design, recipients receive fees denominated in every extant asset type.

Shielded Hosting Fees + Mandatory DEX Market Buys of ZEC

Payer: asset holders
Beneficiary: block reward recipients and ZEC holders
Proportional to: the value of assets held on Zcash over time
Privacy: strong - every holders pays the same rate, and there is no distinction between rates or asset types
Technical Requirements: all requirements of Shielded Hosting Fees, plus an automatic-market-maker Decentralized Exchange (DEX) protocol

This is an extension of basic Shielded Hosting Fees, except rather than delivering fees paid in each asset type directly to block reward recipients, those fees are used to make immediate market order purchases of ZEC on an exchange. The purchased ZEC is then delivered to block reward recipients.

Because of the constant buy pressure for ZEC on the exchange, ZEC holders benefit from an “always available buyer” as well as by the price impact of this stream of purchases.

Private Transaction-Volume-Proportional Fees

Payer: transaction senders
Beneficiary: ZEC holders
Proportional to: the amount transferred
Privacy: strong
Technical Requirements: A cross-asset price oracle + zero-knowledge circuit support

These ZEC-denominated fees are charged proportional to the amount being sent in a transaction of any asset type.

A price oracle defines the exchange rate between any ZSA asset type and ZEC. Privacy is achieved because the ZEC-denominated fee is privately permanently removed from the supply. The sender must provide both the ZSA assets to be sent as well as sufficient ZEC which is used to cover this fee.

This mechanism is proportional to the amounts transferred, not the rate of transfers or amounts held. (Example: sending 10 units at once or 1 unit 10 times has the same fee cost.)

This mechanism is not proportional to the value held on the network or to transaction processing costs. Since block producers (miners) are not direct recipients, they have no incentive to include transactions in blocks.

ZEC holders indirectly benefit due to a reduction in the supply. However, since the amount of these fees are private, no one will have an accurate view of the ZEC supply other than an upper limit.

Transparent Direct Transaction-Volume-Proportional Fees

Payer: transaction senders
Beneficiary: block reward recipients
Proportional to: the amount transferred
Privacy: very weak - asset types and amounts are fully public
Technical Requirements: consensus rule support

These directly-denominated fees are charged proportional to the amount being sent in a transaction of any asset type. The fee is paid in that asset type directly to block reward recipients.

This mechanism has very weak privacy, since the asset type is visible, and since the proportion is known, the amount being transferred is trivial to infer even if otherwise obscured.

Because this mechanism scales with the amount transferred, it has some similar properties to Private Transaction-Volume-Proportional Fees. However, since fees are paid explicitly to block rewards rather than privately burnt, miners have a direct incentive to include transactions, and would rationally prioritize transactions which transfer larger amounts.

Scarce Active Asset Type Auctions

Payer: asset type creators
Beneficiary: block reward recipients
Proportional to: the demand for distinct asset types
Privacy: strong for transactions (asset creation and issuance are assumed to be public in Design Constraints and Assumptions)
Technical Requirements: consensus rule support

In this design, there is a sequence of intervals during which an artificially limited number of asset types are allowed to be “active”. Leading up to each interval, parties interested in ensuring a given asset type will be active in that interval post ZEC-denominated bids in support of activating a specific asset type in the upcoming interval. When the interval begins, the asset types with the largest aggregate posted bids become active, and the posted bid amounts are paid out in subsequent block rewards during the interval. Bids posted for asset types which did not become active are returned to bidders.

Scarce Asset Type Creation Rate Auctions

Payer: asset type creators
Beneficiary: block reward recipients
Proportional to: the rate of demand for distinct asset types
Privacy: strong for transactions (asset creation and issuance are assumed to be public in Design Constraints and Assumptions)
Technical Requirements: consensus rule support

In this design, there is a sequence of intervals during which an artificially limited number of asset types can be created. An auction is performed in each interval to create asset types. Unlike the previous design, once an asset is created it remains active and operational.

Supply-Proportional Issuance Bonds

Payer: asset issuers
Beneficiary: ZEC holders
Proportional to: the amount issued across assets
Privacy: strong for transactions (asset creation and issuance are assumed to be public in Design Constraints and Assumptions)
Technical Requirements: a price oracle + consensus rule support

In this design, an issuer may not issue new units of a given asset type unless they have posted a price-proportional amount of ZEC to a bonding system. Once the supply of the asset decreases due to redemption, the issuer can then withdraw the excess ZEC held in the bond.

This scales demand for ZEC with demand to issue other assets. The bonding and issuance system are public without privacy. Bonded ZEC is temporarily removed from the supply, providing a transient benefit to holders of liquid ZEC.