Questionnaire - Reserve/PWC

Questionnaire

Legal Disclaimer: The following content describes currently intended designs. These designs are subject to revision over the coming months and are intended functionality with no guarantee of that functionality.

a. General

i. Which blockchain / DLT are you building on top of?

We are currently building Reserve on Ethereum. That said, we intend to have the Reserve stablecoin available on any smart contract platform that can support it via interoperability solutions like two-way pegs.

ii. How does the stablecoin work?

Stablecoins work by balancing supply such that, when combined with demand, the price stays at the target price. In the Reserve stablecoin system, new Reserve stablecoins are minted and sold by the protocol smart contract when the price of the Reserve is above $1, which reflects an increase in demand. This increases the supply and reduces the market price. Newly minted Reserves are sold for assets which are put into a smart contract called the Vault. When the Reserve is below $1, assets from the Vault are used to repurchase

Reserves, which are then burned. This reduces the supply and increases the market price. The Vault holds a portfolio of assets diversified across asset class, jurisdiction, and token producer. Such a portfolio minimizes any counterparty risk that would typically be associated with a particular tokenized asset while simultaneously producing a stablecoin that is designed to be robustly stable and decentralized.

In addition, there is a non-pegged token in the system called the Reserve Share. This token entitles the holder to the proceeds from appreciation of assets held in the Vault as well as transaction fees. It is used to fund development and promotion of the system, and as a backup source of funds for stabilizing the Reserve.

iii. **What is the purpose of your coin? What does it aim to achieve, and which problems does it solve?**

Stablecoin designs balance three axes of features: stability, decentralization, and scalability. Stability refers to the stability of the value of the coin, specifically in terms of its average volatility and its robustness in worst-case market conditions. Decentralization is a measure of the degree you have to trust any entity to ensure the stable value of the stablecoin. Scalability, in turn, is the degree to which the supply of the stablecoin can increase to meet demand.

One can think of stablecoin protocols as making trade-offs across these. For example, MakerDAO primarily trades off against scalability, while Basis trades off against stability in bad market conditions. Our protocol design is unique in that it aims to balance all three design features and avoid uneven tradeoffs.

iv. **When we say something is stable what do you think it means? And when it comes to monetary policy specifically?**

When people talk about stability, they typically mean that something has stable purchasing power. This means that a given number of units of the asset can buy roughly the same amount of goods and services now as it can in the future. There are two common monetary policy approaches available for achieving this: exchange-rate pegging and inflation targeting.

We, like all stablecoins we are aware of, follow the exchange-rate peg approach. This is because it is simple enough that it can be implemented as a smart contract. Inflation targeting would plausibly allow for a more robust and independent system, but is currently prohibitively difficult to implement.

v. **What is your revenue model?**

The protocol can make money from transaction fees and the appreciation of assets held in the Vault. This topic will be expanded upon elsewhere.

b. Launch & marketing

i. **What does the market need to be confident in the stability of your token?**

The core feature to check when assessing the stability of a stablecoin is the amount of value backing the stablecoin relative to its market cap. In the short to medium-run, the switching costs between different stablecoins will be very low. This means that small changes in the market’s confidence in the stability of a given stablecoin can lead to large and rapid reductions in demand for that stablecoin. If the stablecoin system is unable to repurchase a large portion of the stablecoins it issues at the peg price, it is likely that its peg will break at some point. As such, it’s important for a stablecoin to have a large amount of collateral backing its value.

Our system is designed to maintain 100% collateralization of the Reserve by holding tokenized real assets in the Vault. The Reserve Share holders are also used as a lender-of-last-resort for worst case scenarios. This means that the market can be confident in the stability of our token as long as decreases in demand for the Reserve do not exceed the combined backing provided by the Vault and the Reserve Share holders.

ii. **How are you bootstrapping to that level of confidence?**

Given that the assets backing the value of minted Reserves are held in a public smart contract, anyone can audit and see that these assets are available and being used to back the Reserve. Also, the Vault will be capitalized over 100% for a period after the launch of the protocol to provide additional confidence to the market. This reduces uncertainty around the stability of the system, though time passing with the protocol operating correctly will also help build confidence.

iii. **What are your go-to-market strategies?**

For competitive reasons, we are keeping this information private for now.

c. Economics

i. **What is your coin stable with respect to?**

Reserve is designed to be stable against 1 USD.

ii. **How much volatility can this peg withstand? Is that the same for upwards and downwards pressure? How wide is the band of behavior it can support?**

Upwards volatility: When demand for Reserve increases, the price it trades for will be higher than usual. The protocol uses this as an indication of increased demand, minting and selling new Reserves in response. These Reserves are sold for assets which are put into the Vault, collateralizing the value of the minted Reserves, should demand fall in the future. With this design, the only limit on the growth of Reserve is the availability of suitable assets to store in the Vault. Given the pace at which real assets are being tokenized by companies like Securrency, Harbor, and Polymath, we are not expecting to be bottlenecked by the availability of assets.

Downwards volatility: By targeting 100% collateralization of all minted Reserves, the protocol is designed to support reductions in demand up to the full market cap of the stablecoin. This works as long as the value of the assets held in the Vault do not drop significantly below the initial amount invested in those assets. Due to assets generally appreciating over time, this is easy to do in the long-run but more challenging in the short-run given the assets that will be available on the blockchain in that timeframe. To make up for this, Reserve Share holders are used as a backup source of funds to protect the system from short-run volatility.

iii. **How easy is it to analyze the band of behavior from which it can recover?**

Our protocol fails when the size of the drop in demand for Reserve is greater than the combined value of the Vault and the market cap of Reserve Shares. It is difficult to anticipate the value of Reserve Shares and how large the largest drop in demand for Reserve will be. That said, we can make ourselves resilient to this by:

1. Holding a conservative, well-diversified portfolio of assets in the Vault that are robust to market downturns. Tokenized gold and tokenized debt are good examples of this.
2. Providing a compelling value proposition to Reserve Share holders so that they are willing to assist in defense of the peg when necessary.

iv. **How expensive is it to maintain the peg/stability mechanism?**

After the initial over-capitalization period after the launch of the system, all funds necessary to back the value of Reserves come from the sale of minted Reserves. As such, the only expenses are the operating costs that the protocol incurs. There are four types of operating costs: oracles, trading fees, slippage, and gas fees.

Oracles: An oracle is an agent that provides information to the blockchain such that it can be used in smart contracts. Oracles are necessary for providing price information to the protocol, information which is needed for the protocol to choose whether to mint or repurchase Reserves. The cost of oracles is primarily determined by the approach used, with centralized solutions generally being less expensive than more decentralized solutions. Current estimates for annual cost for the fully-scaled system range from $100k - $3m annually. These expenses can be paid by the protocol using revenue from the appreciation of Vault assets and transaction fees. Given that market information is one of the most commonly desired types of off-chain information, we expect this cost to fall with time as robust and established decentralized oracle solutions become mainstream.

Trading Fees: We use the 0x protocol to facilitate the sale and purchase of assets in our protocol. This incurs a transaction fee of roughly 0.1% per trade. In order to pay for fees associated with trades, we maintain what is called a seigniorage band. Newly minted Reserves are sold for slightly more than the price at which Reserves are pegged, and the small amount of revenue is used to pay for the fees associated with both the minting of those Reserves and, if necessary, their repurchase.

Slippage: Any time you trade an imperfectly liquid asset, you will move the spot price of that asset when you execute a trade. The additional cost of the trade incurred due to this is called slippage. The cost of slippage is greater if you are executing large trades, if you are trying to execute the trade quickly, or if the asset being traded is less liquid. Our protocol mitigates the cost of slippage by limiting the rate at which it trades and by strongly preferring more liquid assets as assets to be used in the Vault. Any remaining costs due to slippage are paid for using the seigniorage band mentioned above.

Gas Fees: The protocol needs to have up-to-date price information for Reserve and Vault assets brought on-chain. This requires our protocol to call a function roughly once every 30 minutes which incurs the usual Ethereum gas fees. We currently anticipate these to cost $200k - $500k per year at current gas prices. This, like the cost of oracles, can be paid for using protocol revenue from transaction fees and the appreciation of assets.

In sum, we anticipate that the combination of slippage and trading fees will cause an average per-transaction cost of ~0.5%, which can be entirely paid for with revenue from the seigniorage band. Our median estimate for the combined expenses of oracles and gas fees is $2m annually. Given a conservative expected annual return of 2% of the total value of the Vault, the break-even point for the operation of the protocol occurs between $100m and $200m Reserve market cap.

v. How transparently can traders observe the true market conditions?

Our protocol controls the supply of the Reserve such that the “true market” price is $1. As such, traders can observe true market conditions just as well as any other asset.

vi. Which monetary theory (theoretical) assumptions do you think are not true and how does your protocol account for that?

Many stablecoin protocols assume a strong version of the quantity theory of money. The quantity theory of money states that, in the long-run, any change in the total supply of money will cause a change in the value of that money such that the value of the total money supply remains the same. For example, if you doubled the nominal number of dollars that everyone holds without changing anything else, the quantity theory of money posits that this would eventually cause the value of the dollar to drop by half.

It is important to stress that the quantity theory of money only claims that this holds in the long-run. Many stablecoin projects assume that it holds in the short run. In contrast, we intentionally designed our system such that we did not need to assume that the quantity theory of money holds in the short-run.

When a Reserve is minted, it is sold for slightly less than the market price, incentivizing people to resell it on standard crypto exchanges for a small profit. This reduces the price observed on these exchanges by increasing the supply on the order book. This incentive means that the protocol can control the trading price of Reserve more mechanistically, as opposed to exclusively relying upon the market responding to a global change in supply.

vii. Does your stablecoin supply scale in response to demand? If so, how?

Yes, the supply of Reserves expands in response to increases in demand observed in the market. If demand for the Reserve increases, it will tend to trade for prices slightly above the pegged price. The protocol uses this as an indicator to mint more Reserves.

viii. **Who provides the capital to maintain exchange rate peg? How are they compensated / Why do you think they would continue to lock up capital, given other investment opps?**

After the initial period of over-capitalization, all of the capital necessary for maintaining the peg is provided by the sale of Reserves due to increased demand. This means that the Reserve holders are the ones that provide all necessary capital. Reserve holders will be willing to lock up capital, because they primary thing they are looking for is stability and liquidity. This is the same reason people are willing to hold money in USD despite it not having a positive return. The asset is *useful*.

That said, some capital does come from Reserve Share holders to pay for the development and promotion of the system, along with over-capitalizing the Vault. If the Reserve market cap and trading volume grows significantly in the future, Reserve Share holders can potentially make a return on this initial investment.

ix. **An eventuality plan in case of a "black swan" event.[1,2] The 1% case will happen eventually.**

The performance of a stablecoin during worst-case-scenario market circumstances is one of the most important components of a stablecoin system. While we think that our 100% collateralization design will protect against most worst-case scenarios, we are considering further designs to increase the protocol's robustness. Several ideas in this area include rate limiting protocol trades (similar to banks closing branches for a few days during a bank run), dynamic peg bands, and off-chain stabilization commitments.

This is a complex topic and one that will be expanded upon in the future.

d. Tech

i. **Are any novel consensus mechanisms used, over and above the underlying blockchain?**

No.

ii. **What transaction throughput can the blockchain currently handle and how does it plan to scale? Do its plans coincide with your plans for your estimated demand?**

Ethereum can currently handle between 10-50 transactions per second. Sharding, Casper, Plasma, and lightning networks like Raiden seem to be the best scaling proposals for Ethereum right now. While Ethereum's available transaction volume will be enough for the very short-term, it won't be sufficient for major commercial transaction throughput, even assuming scaling solutions do work out well.

This is why we intend to take advantage of 2-way pegs to other chains with greater throughput. This can allow for Reserves to be available on other chains despite the protocol smart contracts living on Ethereum. We generally expect Ethereum to be the best place for the stabilization protocol, but applications that require very high throughput would have to occur on a platform like EOS.

iii. **What tradeoffs does your protocol make and why did you make those tradeoffs? (supply/demand, temporarily peg breaking) (censorship resistance) (privacy tradeoffs) (accuracy of present market data and ease of manipulation of the data feed protocol uses (responsiveness of market and ease of manipulation)**

Collateralization: In previous iterations of our protocol, we considered being less than 100% collateralized. There is reason to expect that a design like this could work, as most pegged currencies around the world maintain significantly less than 100% collateral. And being under-collateralized would notably increase the potential profitability of the system.

That said, it seems the primary reason pegged currencies can support undercapitalization is due to their strong network effects. Means of exchange are typically monopolistic due to the fact that they become more valuable the more people use them and due to governments typically outlawing private citizens from creating currencies. This causes the switching cost for typical fiat currencies to be very high. Consider the difficulty of purchasing chips from an American corner store with rubles.

Stablecoins won't enjoy such strong network effects, especially in the early days after their launch. As such, we've concluded that anything less than 100% collateralization would be irresponsible and jeopardize the stability of the system.

Oracles: Many projects require off-chain information to be brought on-chain. Until a fully decentralized oracle solution reaches the mainstream, we will likely use a centralized oracle in its place. If projections indicate the fully decentralized oracle solution will not be available on a suitable timeline, we may develop a semi-centralized oracle solution by leveraging multiple oracles to perform price feed reporting.

Vault Portfolio Selection: The selection of the portfolio of assets held by the Vault will be centralized in the early days of the system. This is due to a combination of the portfolio construction being unusually complicated early on due to the limited set of available assets and due to decentralizing this component being a challenging governance problem. We intend to decentralize this component after launch in future updates to the system.

Protocol Rate Limiting: We may limit the rate at which the protocol purchases and sells Reserves in order to make the system more resistant to attacks. Specifically, this is designed to prevent attacks that operate by manipulating the price of Vault assets temporarily in order to purchase Reserves at an unfair exchange rate. This trades off against the maximum rate of growth of the Reserve market cap, but we believe there are plausible values for the rate limit that both allow the Reserve market cap to grow relatively unhindered while also making these sorts of attacks unprofitable.

iv. Are there any centralized components of your system? Would any of these be easy for govs to shut down?

Governance, oracles, and Vault asset selection will be centralized at launch and decentralized over time through future protocol updates. These components need to be centralized at launch so we can respond to potential unintended behavior in the protocol. Also, successful decentralization of these components is very challenging and would likely delay the launch of the protocol. One of the primary focuses of the Reserve protocol team will be to build effective decentralized versions of these components after the protocol's launch.

v. Does your protocol require information outside the blockchain such as a feed of price data? If so, how does this oracle work? Who manages it, what are the incentives for managing it, and what happens if the data they provide has a glitch?

Our protocol requires information on the current trading price of Reserve and the current trading price of assets held in the Vault. This requires an oracle to submit price feed data. There are three phases the oracle will move through.

First, a centralized phase: Initially, we will write our own software that pulls price data from exchange APIs and sends transactions to the Ethereum network, reporting the needed price information. This software needs to be highly secure and robust. Notably, it needs to run in a distributed, fault-tolerant, replicated environment that can safely store an ethereum private key. So the earliest oracle will likely be hosted on AWS, Google Cloud, or both. This is currently being designed by our engineering team.

Second, a semi-centralized phase: We will engage with a set of chosen entities to provide market information to us, leveraging monetary, legal, and reputational incentives to ensure trustworthy reporting. We are examining both on-chain and off-chain aggregation techniques for this data. This approach should be significantly more resilient to bad data than the centralized approach but will take longer to set up.

[1] https://en.wikipedia.org/wiki/Black_swan_theory

Finally, a decentralized phase: We will utilize something like ChainLink or other decentralized oracle solutions once they have a proven track record of use in high-stakes situations. Given the enormous need for off-chain information, especially price feed data, we expect that more decentralized solutions like ChainLink will emerge over time.

vi. **Which participants can see which transactions? What is the data and metadata available, and to whom? How does this impact privacy?**

Initially, everything is public, including Reserve transfers, and the operations, trading rules, and balances of the Vault. While the Vault should always remain public, we are considering eventually providing options for Reserve transfers that would allow the transfer to be done privately. However, offering private transactions has to be done carefully, as private, decentralized currencies like Monero seem to often be used to facilitate illegal transactions. While we are still in the design stage for this component, it seems likely that some sort of KYC process would be necessary for people wishing to access private transactions.

This is a complicated topic, though, both in terms of technical challenges and balancing tradeoffs. We will likely write about this at greater length in the future.

vii. **Are you doing anything with formal verification? Smart contracts used?**

Yes, we plan to formally verify components necessary for stability. This entails producing machine-checkable proofs that the protocol's rate limits, circuit breakers, and contract upgrades through governance will be functional in case they are needed after launch. These are all temporal properties and thus quite challenging to verify using the available verification platforms; as such, we will augment that verification effort with automated model-checking tools, in which temporal-logic formulae can be cleanly specified.

We also plan to do standard static analysis and exhaustive property-based testing, though these are not strictly formal verification.

viii. **What is the rebase period? (Length of time between currency adjustments.)**

30 minutes.

ix. **Can we make this automated?**

In the short-term we will call the Vault smart contract ourselves once every rebase period. In the long-term we can likely use the Ethereum alarm clock.

e. Regulation

i. **What are your perceptions of local and global regulation in supporting stable coin, asset backed token economies?**

In the US, we face an alphabet soup of regulation, predominantly that by the SEC (securities classification of the tokens, Reg M, private and public offering routes, and broker-dealer and ATS matters), but also FinCEN (MSB issues, AML/KYC requirements), and the IRS (tax implications of using securities as currency). In the background are CFTC regulations (if the stablecoin may not be a security), BSA issues (if the stablecoin is to be considered first and foremost a currency) and ICA issues (pertaining to certain asset-backing of tokens). Beyond the US, we must learn and comply with local law in each jurisdiction we offer in.

ii. **What could be done to improve regulation in terms of speed, quality, value for your company?**

The main way we'd like to see regulations improved is having them reconsidered and updated to reflect the realities of blockchain activities, including their technological limitations. For example, once a protocol is released and decentralized, the group of people that created it lose control over it and should not necessarily be legally equated with it.

This has potential implications for money services businesses (MSB) regulations, where if a protocol is engaging in MSB activities, certain AML/KYC checks on users are mandated but may be technically infeasible to run, such that regulators could not force any party to comply.

Similarly, once the creators of the protocol fall out of the picture, it makes much less sense in the crypto world than the normal securities world to impose securities regulations on the creators, such as ongoing financial and other reporting (the protocol is not the company), and it may not make sense to even consider the tokens securities under the Howey test, where the predominant forces determining the expectation of profit (if any) are not the managerial efforts of the creators. On a policy level, the reasons for security regulation at all (e.g. high informational asymmetry between companies and investors) don't perfectly extend to the decentralized world.

Laws from 1933 were written without any thought given to blockchain technology and are unsurprisingly a poor fit. Until the laws are revised or new regimes created, we are stuck speculating on how regulators will inevitably inelegantly apply old law to a new world.

f. Testing

i. **What kind of simulations have you done and what have they helped you learn? (simulating broad array of market conditions)**

One of the simplest approaches to modeling uncertainty in financial markets is a random walk or coin flip model. In this approach, the entire evolution of a security's return is captured in two variables: the expected return,, and the volatility,. The basic model also incorporates some notion of time, denoted as t, and generates possible paths in the form of a binary tree, where at each time step, there is a 50% chance of ending up at t + t and a 50% chance of ending up at t - t.

When it is assumed that each step is independent of the previous step, this approximates a normal distribution. While simple, this method of modeling can be powerful because the time step can be shrunk very small so as to mimic a continuous distribution, and it allows for mathematical treatments that have convenient analytic properties. Indeed, this technique can be used to derive famous neoclassical finance results such as the Capital Asset Pricing Model as well as the Black-Scholes-Merton option pricing formula.

An obvious and important limitation of the coin flip approach is that it does not capture two main empirical facts that are observed in the return distributions of almost all securities: negative skewness and fat tails. Common techniques for creating models with these observed characteristics involve using market data to fit a distribution that contains both skewness and fat tails, or taking the empirically observed distribution and drawing samples from it. While these approaches have the added benefit of creating more realistic probability assessments of the returns for an asset, they often assume the distribution is stationary and that each return is independent of the previous one. More advanced approaches in this area allow for distributions to change and for returns to be serially correlated.

Modeling the market dynamics after interventions is a primary goal, and necessary to accurately predict the efficacy of the system. But, in general, all of the previous approaches, which are either based on strong assumptions of normality and independence or based on historically observed statistical relationships, suffer from a fundamental limitation in that they cannot predict or describe how market dynamics will change in response to the protocol implementing a policy — or to any particular policy or rule change. This limitation, popularized by Robert Lucas, implies that if we wish to derive policy advice from a particular model or analysis, then the model must be motivated by lower level descriptions of agent behavior. Since the advent of "the Lucas critique," academic research in macroeconomics has shifted majorly towards incorporating such "microfoundations" into models.

It is because of this that we relied on Agent-based Modeling (ABM) instead. ABM is a technique which models systems as collections of autonomous interacting entities. It is primarily concerned with specifying the rules and behaviors of the individual entities and then observing the macroscopic outcomes of the whole system.

Taking an agent-based approach, we developed a modeling framework that simulates the behavior of individual trading agents who interact in a marketplace by placing limit orders into central limit order books for the various assets of interest for our system. This type of approach has been used successfully by proprietary trading firms and other institutions who are highly incentivized to have accurate models of trading markets. While requiring significant engineering and computational time, our simulation engine allows us to thoroughly probe scenarios where many profit-seeking traders with different strategies and trading styles react to different policy mechanisms enforced by the protocol. In contrast to the simpler statistical models, our ABM approach satisfies both the critique from academic economic theory as well as the finance industry's desire for practically useful models.