# Decentralized Prediction Markets

Jianxiong Zhan (jz3030), Chi Wai Lau (cl3639), Lawan Rahim (lr2965), Quoc Le (qnl2000)

## Abstract

A decentralized prediction market (DPM) is a place where anyone can speculate on the outcomes of future events. In this project, we are going to explore key elements in DPMs and understand their inner workings. We are also going to dive into the actual data and observe how various situations have been handled in practice. Finally we hope to gain a good understanding of the current challenges in the DPM landscape and propose viable solutions in this project.

## 1 Introduction

Imagine a place that tells you about the future, e.g., *U.S have enough Covid-19 vaccines by mid-2021*? Imaging a place where you could google to get insight about what humanity's collective wisdom thinks is about to happen in the future. This would seem like science fiction, but thanks to the rise of blockchains and crypto, "google to the future" is a reality if we have truly open, transparent, and decentralized prediction markets (DPMs) to leverage *the wisdom of crowds* [1].

A prediction market is a place where participants can speculate on the outcomes of future events, such as election results, sporting events, business outcomes, and more which isn't different from traditional financial markets or betting markets. Like currency, prediction markets traditionally are managed by centralized entities. The issue with centralized entities is that they represent single points of failures; for example, governments have found a way to shut down centralized prediction markets such as InTrade over the past decade.

A decentralized prediction market (DPM) is inspired by Bitcoin such that its lack of a single point of power or failure makes it impossible to be controlled or shut down by anyone. Through here, anyone anywhere can create markets, bet on predictions and verify execution at any time, thus DPMs are ownerless and resistant to censorship. In addition, they have boundless ability to aggregate the world's information to make predictions.

Examples of DPM platforms that are open to the public include DPMs built on Ethereum such as Augur, Catnip built on Augur, Polymarket, Omen built on Gnosis, we also have PredIQt built on EOS, HiveMIND AGORA built on Bitcoin, etc. For example, with Omen anyone can create a prediction market for any topic, where they are able to specify the oracle that will resolve a market's outcome. One oracle choice would be picking Kleros as an arbitrator.

So far, DPMs have experienced successes as well as failures compared to other forecasting methods. For example, Polymarket has correctly predicted events such as Microsoft's failed acquisition of Tiktok. Other popular markets related the outcome of the U.S. election and the future of COVID-19. Furthermore, people have used DPMs to determine the accuracy of median claims. DPMs have also suffered from low liquidity and scaling challenges, which we will discuss in 6.1.

The evidence demonstrates that designing DMPs properly is essential for seeking truthful information with lower fees. This report surveys key elements in DMPs and explores real world data to see how they work. Therefore, our findings are valuable for moving forward for designing future DPMs.

The remainder of this report is structured as follows:
Section 2: Reviews the existing trading mechanisms.
Section 3: Discuss liquidity sources.
Section 4: Compare Augur and Gnosis protocols.
Section 5: Descripts life cycle of DPMs.
Section 6: Presents application of DPMs in the US-election.
Section 7: Address current challenges.
Section 8: Summarize conclusion and outlook.

## 2 Trading Mechanisms

The question “how supply and demand influence the price of an asset in the decentralized markets” has attracted a lot of attention in recent years. Trading mechanisms play a vital role in determining price. We describe the existing mechanisms below. The section reviews the existing trading mechanisms such as Continuous Double Auction (CDA), Logarithmic Market Scoring Rule (LMSR), and Constant Function (e.g., constant product/mean/sum).

### 2.1 Continuous Double Auction (CDA)

There are various forms of double auction trading mechanism such as p2p [2], MUDA: a truthful multi-unit double-auction mechanism [3], Huang: a multi-unit double auction e-market [4]. A CDA model is the most common and successful model in traditional centralized prediction markets, like the NASDAQ, the New York Stock Exchange (NYSE), the Iowa Election Markets, etc.

A CDA is a mechanism that matches buyers to sellers that was introduced by Daniels [5] under the assumption of random order flow at each time step. Buyers propose bids, and sellers submit asking prices. If the two sides of the market reach a mutual agreement in price, a trade is executed immediately. Trades are enforced at the highest price match, and trades can be executed on a continuous basis if there is enough liquidity.

For example, suppose we have the following bids (see Table 1) and the corresponding supply and demand curves are shown in Figure 1.

Table 1. Bids/takes dataframe

| user | quantity | price | buying | time | divisible |
|---|---|---|---|---|---|
| 2 | 1 | 1.2 | TRUE | 0 | TRUE |
| 1 | 2 | 1 | TRUE | 0 | TRUE |
| 3 | 2 | 1.3 | FALSE | 0 | TRUE |
| 4 | 2 | 1 | FALSE | 0 | TRUE |

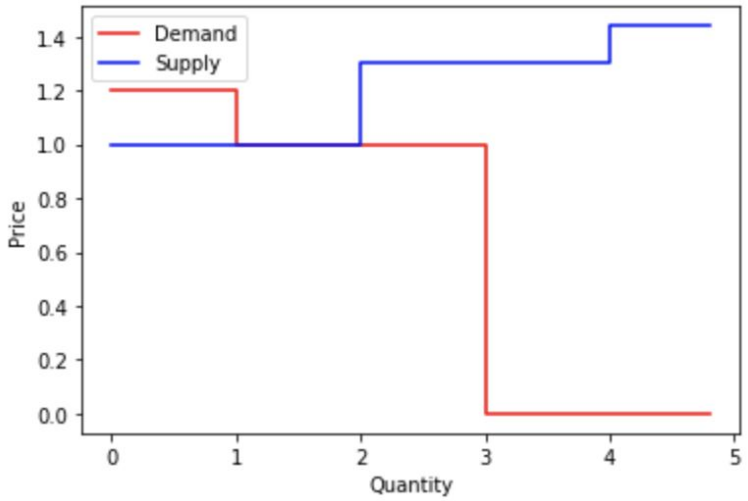


**Figure 1**. Original supply and demand curves

Analysis of the results are shown below:

Round 1

- user 2 trades with user 4, they trade one unit since CDA forces the higher price, and user 4 goes to next one with one remaining unit.
- user 1 and user 3 are not pared with anyone and continues to round 2.

Round 2

- user 1 trades with user 4, they trade one unit.
- Round 3 only user 3 remains, so no trade can occur, the algorithm ends.

We observe that in round 1, buyer 2 proposes for $1.2 and seller 4 asks for $1, so there are price differences. Thus buyer 2 and seller 4 should agree on a price. In the NYSE, the trade will be executed with the lower price and the broker would earn the difference. It is worth noting that no one takes the other side of seller 3 so that seller 3 can't make any trades, suggesting that a CDA can be a problem in a market with low liquidity.

### 2.2 Mathematical Models for Automated market makers (AMMs)

To alleviate the low liquidity problem, platforms use what is known as Automated Market Makers (AMMs). AMMs are algorithms (programs) saying "if this, then that", i.e. if the price of an asset moves up or down, then take action. Prediction markets use AMMs to set share prices. We compare mathematical models for AMMs including LMSR, Constant Function (e.g., constant product/mean/sum), and others below.

### 2.2.1 Logarithmic Market Scoring Rule (LMSR)

In recent years, the most popular scoring rule used in AMMs is Hanson's LMSR[6]. The LMSR market maker is designed specifically for the prediction market use case (e.g., Gnosis), and its properties have been well researched.

***LMSR***

*Let* $P_i(q)$ be the current market price of the $i_{th}$ security, $q_i$ is the number of outstanding shares for the $_i$th security in the market, and b is an arbitrary constant that is used to scale shares to make the price change reflecting real world probabilities.

**Cost function**. The cost of a given trade is defined as:

$$C(q) = b * \ln(\sum_{j=1}^{k} \exp(q_j/b))$$

***Price function***. *A price for* the $_i$th security is:

$$price: p_i(q) = \frac{\partial C(q)}{\partial q_i} = \exp(q_i/b) / \sum_{j=1}^{k} \exp(q_j/b)$$

**Fee**. A trade fee is determined by the changes in cost function before the trade and the after.

$$fee = |C(q)_{before} - C(q)_{after}|$$

Next let us look at an example to show how to design AMMs using LMSR. For example, consider a market with three securities and the numbers of outstanding shares are 10, 20, and 30 for security 1, 2, and 3, respectively. Let b be 10. The prices for security 1, 2 and 3 are:

$$price: p_1 = \frac{\exp(q_i/b)}{\sum_{j=1}^{k} \exp(q_j/b)} = 0.09$$

$$price: p_2 = \frac{\exp(q_2/b)}{\sum_{j=1}^{k} \exp(x_j)} = 0.24$$

$$price: p_3 = \frac{\exp(q_3/b)}{\sum_{j=1}^{k} \exp(q_j/b)} = 0.67$$

Suppose a trader wants to buy 5 shares of security 1, the trade fee would be 0.57.

$$cost_{before} = \mathrm{b} * \ln\left(\sum_{j=1}^{k} \exp(p_j/b)\right) = 34.08$$

$$cost_{after} = \mathrm{b} * \ln\left(\sum_{j=1}^{k} \exp(p_j/b)\right) = 34.64$$

$$fee = |34.08 - 34.64| = 0.57$$

Observe that in the LMSR equations, if b is small, the prices change fast, meaning purchasing a small number of shares increases the price a lot. Thus, choosing a good value of b can be tricky, which depends on the nature of the market.

It turns out that implementing LMSR or LS-LMSR type of functions can be quite expensive because its cost function uses a logarithm, especially when calculating price changes for many outcomes this can become quite gas costly. Augur originally was LMSR or LS-LMSR based but it is no longer using LMSR or LS-LMSR due to gas costs.

### 2.2.2 Constant Function

Another market maker formula is the constant function, which can be constant production/mean/sum. Automated market makers using constant function are often called Constant Function Market Makers (CFMMs). Currently, Gnosis offers smart contract implementations of two automated market makers for prediction markets: the LMSR market maker and the constant product market maker (CPMM).

**Constant Product** (CP). Market makers using constant product, aka, constant product market makers (CPMMs), or the fixed product market makers (**FPMMs**) in Gnosis' codebase, have been used in decentralized exchanges (DEXs) (e.g., *Uniswap*) to enable on-chain exchanges. CP is the first that incentivizes infinite liquidity by increasing slippage as large quantities of the pool are purchased. In the market using CPMMs, it keeps track of the cost function $C(q) = \prod_{i=1}^{n} q_i$ as a constant. So that the price function for *ith* security can be written as:

$$price: \; p_i(q) = \frac{\partial C(q)}{\partial q_i} = \prod_{j \neq i}^{n} q_j$$

The constant product function for XY=64 and XYZ=64 with two and three tokens, respectively is shown in Figure 2. Observe that the curve reflects behaviors in

competitive markets which echoes the combination of supply and demand curves. Moreover, the 'slippage' ensures that tokens never run out.

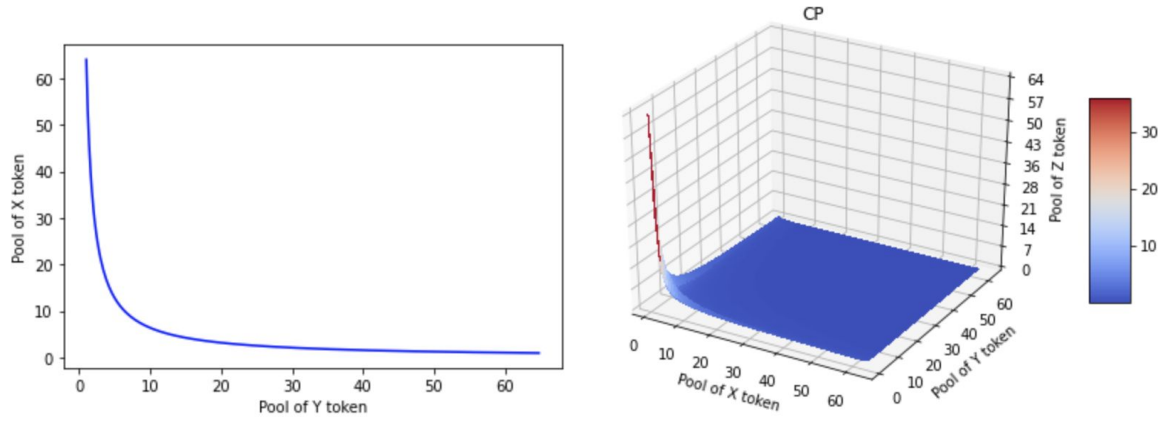


**Figure 2**. Constant Product Market Model. (a) xy=64[7]. (b) xyz=64.

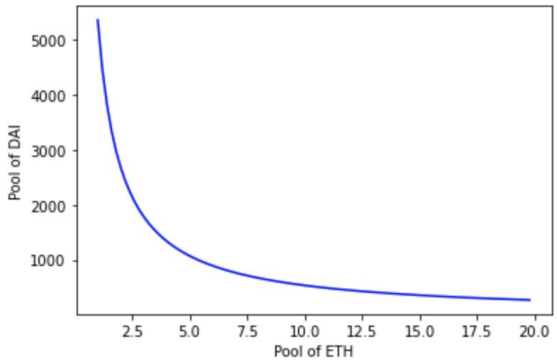


**Figure 3**. Constant Product Market Model. xy = k, where x = 3 ETH and y = 1,788 DAI loaded up initially.

For example, consider a pool which is loaded up with x = 3 ETH and y = 1,788 DAI to make equal values in the pool. The token exchange price is determined by the ratio of x and y so that the product xy = (x+Δx)(y-Δy) is constant. Observed that the price (Δx/Δy) is the function of x/y. Intuitively, if we draw the equation xy = 3* 1,788 as shown in Figure 3. Buying and selling ETH is basically moving up and down along the curve according to the ratio of the tokens, which is easy to implement. It is worth noting that the bigger k is, the smaller price changes. For example, if one bought or sold 1 ETH in the current market setting, it will end up having 33% changes in the price. But if the market is with x = 100 ETH and y = 596,000 DAI, then k is 100 * 596,00 which is much larger. As a result, if one bought or sold 1 ETH in the market, it only has a 1% change in the price. Actually, k here is related to what we call a 'thick' or 'thin' market. The bigger k is, the thicker the market is, and the less the price changes as one buys a certain number of tokens. The fee of this type of

market is about 0.1-0.3%. In general, the constant product function provides a simple approach but surprisingly effective for trading between pairs of tokens in a decentralized fashion.

However, CFMMs have their own drawback in the real applications, named 'impermanent loss' when the price of the deposited assets changes from the time of deposit. Observe the curve in Figure 2 and suppose that Alice has supplied $10 in Token X for $10 in Token Y so that they each make up 50% of Alice's share of the pool; however, as trading takes place, this composition may shift either way towards one or the other, and this means that Alice's relative weighting may change and hence Alice is not in control of her allocation. The example suggests that the smaller k is, the larger Alice is exposed to 'impermanent loss'.

**Constant mean (CM)**. The cost function with a weighted number of outstanding shares is also being used in DEXs (e.g., balancer). The cost function for *ith* security can be written as: $C(q) = \prod_{i=1}^{n} q_i^{w_i}$ . The price for *ith* security is:

$$p_i(q) = \frac{\partial C(q)}{\partial q_i} = w_i q_i^{w_i - 1} \prod_{j \neq i}^{n} q_j$$

Where $q_i$ is the quantity of *ith* security, $w_i$ is the corresponding weight for *ith* security, the sum of the weights equals to 1, that is: $\sum_{i=1}^{n} w_i$ , and k is the invariant. If each security has equal weight, then CM is equal to CP. For example, if there are three tokens x, y, and z with the equal weights, then the cost function is $(xyz)^{1/3} = k$, which can also be written as $xyz=k^3$ in the CM format.

It is worth noting that the equation can be written as log weighted average:

$$\sum w_i \ln x_i = \ln k \text{ or } exp(\sum w_i \ln x_i) = k \text{ or } exp(\frac{\sum w_i \ln x_i}{\sum w_i}) = k$$

This way is convenient because taking the weight average of the logarithms of the variables, it will raise an exception if any asset vanishes, but this is often good for our case because it reflects a desirable property of maintaining 'slippage'. The constant mean cost function for $xy^2=64$ and $xy^2z^3=64$ are shown in Figure 4. Similarly to CP, the function is still a convex function, which is desirable to derive prices. Observe that as weights change, the slope of the curves changes accordingly. This suggests that a CM model with changing weight might be able to mitigate 'impermanent loss'. It turned out this is exactly what Bancor v2 did. In Bancor V2, it adjusts weights dynamically according to the relative numbers of the tokens to maintain the balance of values for the tokens in the pool.

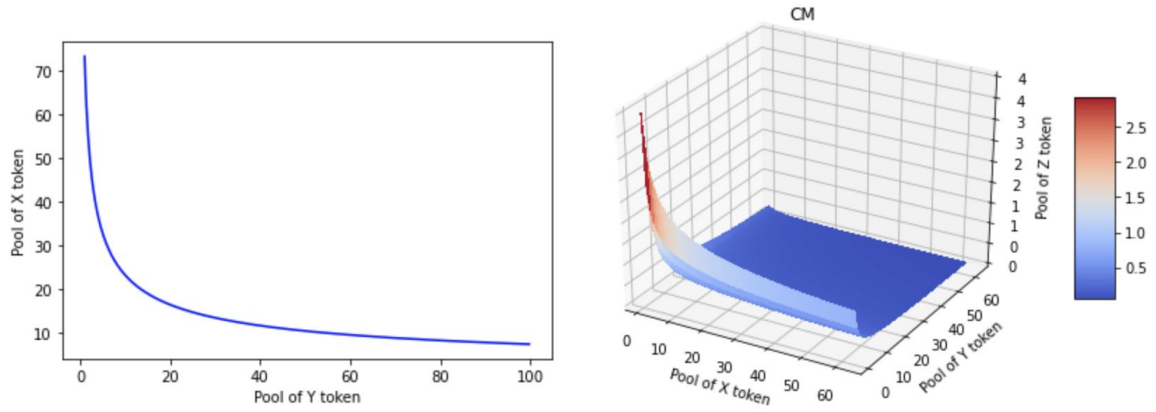


**Figure 4**. Constant Mean Market Model. (a) $xy^2=64$. (b) $xy^2z^3=64$.

**<u>Constant Sum</u> (CS)**. It is a natural tendency to use Constant Sum (cs) to design Constant Sum Market Makers (CSMMs), where the cost function can be written as: $C(q) = \sum_{i=1}^{n} q_i$ and the price for *ith* security is: $p_i(q) = \frac{\partial C(q)}{\partial q_i} = 1$ ,meaning one of ith security can always trade for one of jth security as long as there is any. Observe that A CF model isn't able to support 'slippage' as both X and Y reach 'zero' as shown in Figure 5 so that it doesn't fit DEXs usage cases without additional intervention.

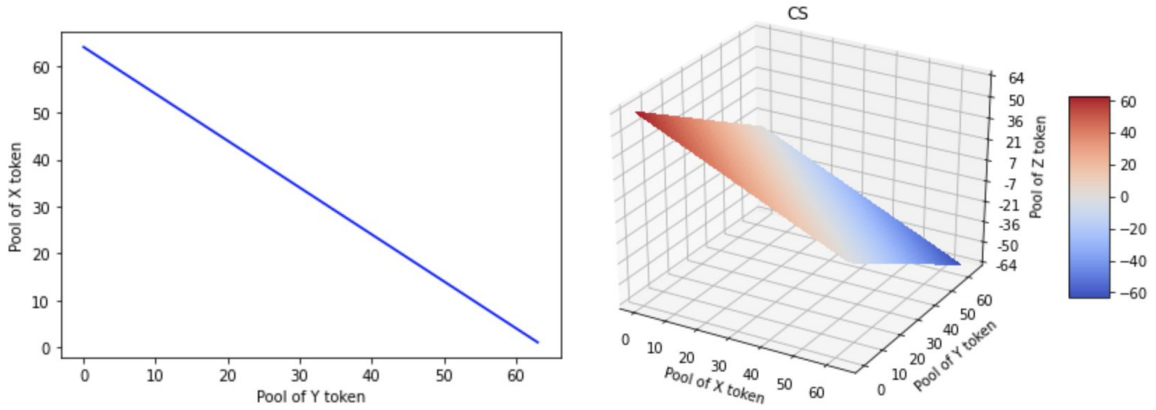


**Figure 5**. Constant Sum Market Model. (a) x+y=64. (b) x+y+z=64.

**<u>Hybrid CF</u>**. The cost function for CP and CM is well suitable for assets like ETH but if the price slippage for an asset is large and one should provide enormous funds to keep a meaningful liquidity. To minimize the price slippage problem, a hybrid function of a constant sum and a constant product was introduced by Egorov [6].

The cost function for ith security can be written as: $C(q) = \gamma \sum_{i=1}^{n} q_i + \prod_{i=1}^{n} q_i^{w_i}$ . The price for *ith* security is: $p_i(q) = \frac{\partial C(q)}{\partial q_i} = w_i q_i^{w_i - 1} \prod_{j \neq i}^{n} q_j + \gamma$ .

As we can see in Figure 1 in Egorov's paper [8], the price changes on the Hybrid CF model is less pronounced than it is on the CP model. Price 'slippage' is quite small near equilibrium point 1.0 which is a good property for stable coins, while it still provides liquidity for rare tokens as the price deviates from the equilibrium point. Due to its good nature, a few DEXs such as *Curve*, *Bancor V2*, are using the Hybrid CF. If exchanges are solely between stable coins, then market makers with Hybrid CF are a great choice.

In general, if we say that price in LMSR is determined by a SoftMax function, cost function in constant sum is the arithmetic mean function, in constant mean is a log weighted average function, and in hybrid model is a mixing model of an arithmetic mean and a log weighted average. These formulas are used to price discovery has its advantage in a sense that it can discover price without an order-book so that they are being widely used in the decentralized exchanges. It also makes manipulation less likely as you cannot see the order of other trades in these instances. Furthermore, AMM based DEXs have been proved to be path independent, meaning that if the market moves from one state to another state, the payment/cost is independent of the paths that it moves [9]. They also have other properties such as translation invariance, and liquidity sensitivity, which have been studied by Ganeris and Chitra [9]. Although it is easier to calculate and check trading rules with constant function models, it comes with risks such as high slippage for larger orders causing 'impermanent loss'.

## 3 Liquidity Sources

In this section, we review three models for designing decentralized exchanges: the order book based, auction based, and the AMM based liquidity pools.

### 3.1 Order Book based

Classical order book mechanisms have been widely used in the traditional prediction markets (i.e., IEM, Predictit) where there is consistently high liquidity. An *order book* is a list of buy and sell orders for a specific security at each price level. Decentralized exchanges have borrowed the orderbook concept from traditional

trades to decentralized exchanges. In the decentralized exchange, order books may exist *on-chain*, hosted on a distributed ledger, or *off-chain*, hosted by third parties.

**Off-chain order book**. Off-chain order books are order books that are hosted by a centralized entity outside of a distributed ledger. Exchanges based on off-chain order books are actually permissioned because there are permissioned exchange operators to actually collect and match signed orders off-chain on centralized servers and send the match to the blockchain and trades are settled on-chain in a non-custodial way. An off-chain order book with on-chain smart contract is a very successful market mechanism that has been implemented by etherdelta, Ethereum-0x, and others.

The biggest advantage of off-chain order books is that they are able to accommodate quick order turnover. Instead of waiting for a block to be mined and confirmed to update the order book on-chain, off-chain services can update ledgers almost instantaneously. It is cheap to adjust the actual order with bootstrapping exchanges; however, users must rely on the hosts of the off-chain order book. Thus, they are more likely to be more strictly regulated compared to a truly decentralized market.

**On-chain order book**. In contrast to off-chain order books, on chain order books are hosted directly on the distributed ledger, orders are submitted to the distribution ledger to match with previously posted orders on-chain. On-chain order books don't require a centralized server and thus have the benefit of censorship resistance, but users will need to create transactions and pay for gas to place limit orders and to cancel existing orders, thus they end up very expensive to implement.

In general, one of the advantages of order book is its underlying CDA trading mechanism which is very well known and has established itself in financial market (i.e. NYSE) and it works well in high volume markets such as Coinbase, but they are not an optimal model for low liquid markets (i.e., prediction markets) due to the shortcoming of CAD. In Ethereum, it brings loopring or ring trades to order-book to improve low liquidity problems. In Augur v1, it used 0x to boost off chain order book. In Augur V2, it was updated to an on-chain orderbook powered by Ox Mesh and in the meantime, it uses CF instead of CAD trading mechanism for price determination. One of the shortcomings in the order book-based trading is front running due to its racing condition between taker-orders. An unsavory off-chain operator or on-chain miner or on-chain validator who sees profitable transactions

submitted by traders or broadcasted on chain take advantage of their power to include their own orders before executing others to earn profits, which is called ‘front running’.

### 3.2 Auction based

To overcome low liquidity and ‘front running’ problems in the decentralized markets, Gnosis team developed *Dutch auction* in December 2017 and *batch auction combined ring trades* in 2019 to boost liquidity. In the auction-based exchange, the exchange is not continuous but split into discrete windows for traders to submit orders. In a Dutch auction, there are three ingredients: a predefined sell volume (window), a price (auction starts with high price and is decreasing) and bid volume that have been collected over time. Price will be defined as: *offering price = sell volume/bid volume*. The price with the highest bid volume is selected as the offering price. In a batch auction, a discrete window is defined by time and optimal price is calculated based on trading volume after the window closed. The advantage of batch auction is that it allows for ring trades by accumulating orders in the pool. In a sense that in each time window, a taker-order loop will not limit to two parties but many as long as they can form a circle. Both Dutch auction and batch auction are good for price discovery of illiquid tokens because of accumulating order. They both overcome front running occurred in the order book since optimal price is calculated on batch basis, gaming the system becomes hard. The advantage of batch auction is that it enables ring trades to further boost liquidity. The main drawback of dutch/batch auction is slower settlement, consequently higher arbitrage risk.

### 2.2.2 AMM based

AMM based DEXs that make use of liquidity pools have been widely used in the decentralized prediction market to solve the ‘liquidity problem’ where sparse order books struggled to guarantee liquidity to investors on both sides of the trade. Traders in the AMM based markets trade against a pool of assets rather than a specific counterparty which is less expensive than order books. Moreover, it provides an opportunity for everyone to earn income for providing liquidity to trading pools and thereby helping in the exchange of crypto currency, rewards is not only in the form of trading fees but also liquidity pool tokens which are paid out to those who are supplying the liquidity. These AMM based DEXs are designed for the community trade tokens without middlemen. All the benefits mean that the AMM based DEXs is becoming very popular. Currently, Uniswap is the world’s

largest liquidity pool with the market cap of $1.32B. The following three are Curve ($946.4M), Balancer ($395.6M), and Bancor v2($81.1M). We will briefly discuss these four below.

***Uniswap.*** Uniswap launched in late 2018. It is the first automated decentralized exchange at the moment. It is completely free of order books and it provides an on-chain liquidity pool. Trading fee in Uniswap is inflexible which is about 0.3% per trade that directly goes to liquidity providers. It uses CP model (see section 2.2.2) for price determination. Users are able to swap pair assets. The UI of Uniswap is very simple. Trading on Uniswap is pretty simple. When traders want to trade on uniswap, they will enter the amount that she would like to trade and Uniswap will provide a rate, this exchange is facilitated through the use of global liquidity pools for ERC-20 assets. It's with these pools that Uniswap is able to create a unique market for any two assets. As we mentioned in section 2.2.2, the price of tokens moves up or down along the curve in Figure 2 if trades occur and thereby Uniswap is able to provide liquidity by adjusting the price of the order up based on the size relative to the pool of liquidity. Of course, this will lead to significant price slippage on the order if there are smaller pools. For that, traders in Uniswap have the ability to specify a maximum slippage amount that they are willing to accept. There is also a time window setting option for traders to specify how long they are willing to wait for the transaction to execute. Another neat option on Uniswap is for users to earn passive income by providing liquidity to the pool so that other traders can use it in order to facilitate the transactions. The smart contracts will then use the pooled assets to swap the tokens that traders are looking to convert. Liquidity providers at Uniswap will get a share of those liquidity provider fees that the platform charges from traders and it is a neat way to earn additional returns on crypto holdings if there are a lot of trading activities. On Uniswap V2, it has its governance token UNI, and they have been distributed as a reward for providing liquidity so that everyone can provide liquidity and earn some passive income. Since the constant product model has its own drawback of ‘impermanent loss’ as we mentioned in section 2.2.2 in the sense that if the price of the deposited assets changes from the time of deposit, the liquidity provider may be exposed to impermanent loss. The more volatile the assets are in the pool, the more likely it is that you can be exposed to impermanent loss. Thus, users should keep an eye on the proportions in the pool if they do not want to fall victim to that impermanent loss. Also, it might be a strategy to start by depositing a small amount so that they can get a rough estimation of changes.

***Curve***. Curve launched in January 2020. It mitigates 'impermanent loss' by using hybrid CF (see section 2.2.2) to minimize slippage for stable coins [8]. It advocates itself to be an "exchange liquidity pool on Ethereum designed for extremely efficient stablecoin trading and low risk, supplemental fee income for liquidity providers, without an opportunity cost" [10]. It has a native governance token CRV. Curve's UI uses this old-fashioned web style which seems much more complicated at a glance than Uniswap, but it is just as sample as Uniswap to do a trade on Curve. On the homepage, it shows all of the tokens that users can swap between on the home page. Similar to Uniswap, traders can set maximum slippage, but here they will be able to customize gas price. The nice things about Curve are (1) fees on curve is about 0.04% which is lower than Uniswap; (2) Similar to Uniswap, curve liquidity providers also have the opportunity to earn rewards for generating tokens but with less risk from impermanent loss due to using hybrid CF.

***Balancer*** . Balancer launched in March 2020. It uses a constant mean model (see section 2.2.2) for determining price. Balancer has a native governance token called BAL. Similar to Uniswap, Balancer UI is clean and simple. Balancer distinguishes itself from others because it supports up to 8 different tokens, so it gives users a scope of options for different allocations. Additionally, Balancer allows users to customize trading fees they want (anywhere between 0.0001% and 10%) for each asset which allows liquidity providers to make the most of their funds. Furthermore, after selecting exchange rate and the maximum price slippage options on Balancer UI, Balancer also provides exactly how the order has been optimized using those balancer pools - this feature is pretty neat.

***Bancor***. Bancor is the first AMM on Ethereum in 2017. In Bancor V1, it used a fixed weight of CM model which is actually CP model and thereby liquidity provider was suffering from 'impermanent loss'. In July, 2020, Bancor V2 launched and it replaced CP function with unfixed weights CM to mitigate impermanent loss. In addition, Bancor V2 developers use an external oracle to feed external prices to the pool to further mitigate impermanent loss. Similar to uniswap, trading on Bancor V2 is simple and only works for two-tokens. The nice thing about Bancor V2 is that trading fees are flexible.

In general, AMMs are great for illiquid tokens but there are risks for liquidity providers who can be profitable when the token pair trades at around the same price ratio at which the liquidity provider supplies the token pair. This is not always the case because of impermanent loss when the price ratio changes. As we've discussed, some liquidity pools are much more exposed to impermanent loss than

others. As a simple rule, the more volatile the assets are in the pool, the more likely it is that you can be exposed to impermanent loss. In addition, currently, gas fees are making a lot of these protocols incredibly expensive to use.

## 4 Exemplary Protocols for Decentralized Prediction Markets

### 4.1 Augur

Augur is to date one of the frontrunner protocols for decentralized prediction markets [11], allowing anyone to create a market on anything by specifying the event end time, a designated reporter, a resolution source, a creator fee, a validity and a creator bond. After the event has ended, the designated reporter will declare the outcome of the event according to the options set by the market creator and using the specified resolution source. The validity bond serves as a security to safeguard the community from the possibility that the outcome cannot be resolved according to the specifications made by the creator. Similarly, the creator bond serves as a security, but here to incentivize the market creator to choose a reliable reporter. If the latter reports the outcome within 24 hours after the event end time and this reported outcome later turns out to reflect the final consensus, the creator bond is returned to the creator. Finally, traders are charged a creator fee for settling with the market [12, 13].

Assuming rational participants, the price of the shares ideally reflect the probability of occurrence which the community associates with each outcome and that probability in turn would ideally resemble the actual likelihood. The rationale behind this is that users are incentivized to place their bets on the most likely outcome. Trades of these shares are being managed by Augur's automated matching engine which keeps an order book and thereby, matches equalizing trades or creates new shares if need be. During consecutive periods of seven days, fees are being collected from trades and added to the reporting fee pool for that period. The latter is then used to incentivize and pay reporters for their service. In order to report on outcomes, the reporter has to own and stake Augur's native token, the REP (Reputation) token. If her report turns out to match the final outcome, she earns a reward in the form of fees which is proportional to the amount of REP she staked [12, 13, 14].

After the event ends, the market undergoes a number of steps from reporting to settlement which are highlighted in Figure 6.

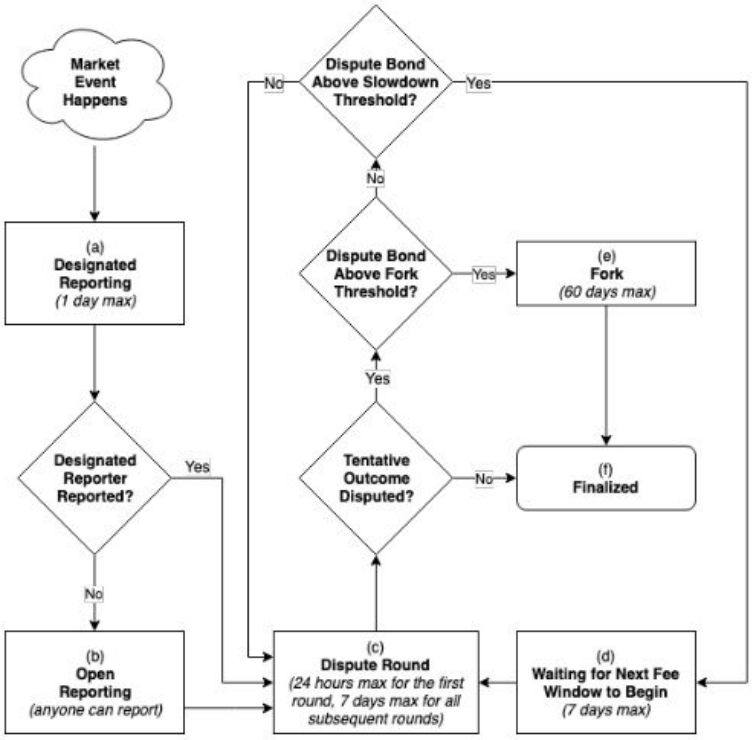


**Figure 6.** To reach consensus on the final outcome of an event, a reporting and appeal procedure is carried out over several stages (reprinted from [21]).

If the designated reporter fails to call the tentative outcome within the allotted time, everyone with REP tokens is allowed to determine it and obtain the creator bond without having to stake their REP tokens. During a dispute round, REP holders are able to dispute this tentative outcome by staking their tokens. A dispute is considered successful once the total amount of staked REP tokens reaches the so-called dispute bond size. Let $O(\beta, n)$ denote the total stake on the alternative outcome $\beta$ at the beginning of round $n$ and similarly, let $T(n)$ denote the total stake on all possible outcomes at the beginning of round $n$. Then the dispute bond size $B(\beta, n)$ necessary to overturn a report in favor of the outcome $\beta$ at the beginning of round $n$ is given by

$$B(\beta, n) = 2T(n) - 3O(\beta, n)$$

This is to ensure a constant return on investment for those who manage to successfully dispute a tentative outcome. If the tentative outcome is not being disputed, the market will get finalized. Otherwise, the market enters one of three possible states depending on the size of the dispute bond in relation to the size of all REP tokens. If the dispute bond size is less than 0.02% of all REP, a new

dispute round is initiated, but this time with $\beta$ as the tentative outcome. If it is between 0.02% and 2.5% of the same, the market will wait until a new fee period starts, again with $\beta$ as the tentative outcome. Finally, if the dispute bond size is greater than 2.5% of all REP, then a so-called fork is created. Forking creates new, so called *child universes*, one per outcome of the corresponding market and simultaneously stops reporting rewards from being paid out as well as prevents any markets in the *parent universe* from getting finalized. This should force participating reporters to solve the dispute corresponding to the fork by migrating their REP to one of the child universes [12].

The fork is being resolved once the 60 day period ends or at least 50% of the REP tokens have been redeployed. The outcome whose child universe contains the most REP tokens will be considered the final outcome. REP tokens which have not been moved to any child universe after the forking period will be permanently locked and those moved to universes not corresponding to the final outcome loose their economic value, creating a major incentive for the reporters to vote on one of the possible outcomes. After the final outcome has been determined, the participants can settle their shares with the market and thereby, ideally gain profit [12].

The backbone of Augur's protocol is the Ethereum blockchain as its functionality relies on smart contracts to orchestrate anything from market creation over trades to settlements [14]. Blockchains themselves however, are justifiably isolated in that only basic types of consensus can be reached, using only information which is available in their ledger. As a consequence, questions like "Who will win the US presidential election in 2020?" cannot be answered simply using a blockchain ecosystem. Instead, an oracle is needed as an interface between the on-chain (on the blockchain) and the off-chain world [15, 16]. This however, would reintroduce centralization and therefore runs counter to the critical intention of dispersing the power away from a central authority. The Augur protocol solves the aforementioned Oracle problem by providing a decentralized, automated and secure middleware bridging the blockchain to the outside world, allowing information flow in both directions [14, 15, 16].

The Augur protocol consists of a set of immutable smart contracts on the public blockchain of Ethereum, so that no single entity can make any changes to its functionality. These smart contracts manage the matching and settling of Orders [14]. Moreover, the REP token used for reporting on outcomes is based on ERC-20 which is a token standard of Ethereum and implements an API,

providing essential functionality for trading, monitoring and more [17]. In the initial version of Augur, Ethereum's native cryptocurrency token, the Ether, has been used as the currency for betting and trading. Ether's high volatility however, poses a problem to the prediction market as it may incurs losses which are not related to the events themselves. A user who may have betted on the correct outcome, might end up losing profits because of a drop in the value of Ether [18]. Stablecoins provide the natural opposite, a cryptocurrency specifically designed to minimize price volatility and peg their value to a reference asset, e.g. the US dollar [19]. As a result, the second rollout of Augur utilizes the Stablecoin DAI which proved to be resistant to censorship and showed low volatility in the long run [20].

### 4.2 Gnosis

In contrast to Augur, Gnosis' primary goal is to provide the foundational infrastructure for decentralized prediction markets and applications based on it, rather than function only as a portal for decentralized prediction markets. As a result, third parties can use their tools to build decentralized prediction markets for the use as an information system like Augur which has been successfully realized in the Omen prediction market [21] or innovate novel applications for example for decentralized governance. On a high level, the foundational building blocks for building decentralized prediction markets and applications on top of them are tools for market creation, asset management and trade and settlement orchestration. Gnosis' tools similarly to Augur rely on Ethereum as the backbone for decentralization, yet are different in some key ways [22, 11].

As one of the core building blocks of the Gnosis infrastructure, market creation requires the creation of two smart contracts which need to be specified by the market creator, but will be created and handled by Gnosis' smart contracts. One of these is called the event factory contract and is used to specify the oracle for resolving the market and the tokens, specifically the outcome and the collateral tokens. While outcome tokens are again shares of each possible outcome, collateral tokens specify the currency pegged to an outcome token. That is, users trade outcome shares of an event and these shares' values are denoted in the currency which the collateral token holds (for example DAI). A trader Alice can get a full set of all possible outcomes tokens of the event for one collateral token. In the case of two possible outcomes, this corresponds to two outcome tokens. She can then keep the outcome token for the event which she bets on, say event A, and sell the other one for its respective market price. If the market

attributes a probability of 1 with the occurrence of event B, then an outcome token of event B can be sold for 1 collateral token [22, 23].

Once event A occurs, Alice can redeem her outcome token for event A for a collateral token, yielding a total of two collateral tokens which conforms to a profit of one token. To allow for such trade, the second so called market factory contract has to be specified which involves the definition of an event, a market maker such as the logarithmic market scoring rule to support the trade management of outcome tokens and a market fee. The market maker helps determine the prices of outcome tokens in the wake of fluctuating demand. It therefore forms an integral part in the orchestration of the trades on decentralized prediction markets and hence, is a component of the decentralized exchange. A decentralized exchange in the context of decentralized prediction markets is a protocol which arranges the matching and settlement of trades in a market [22, 23, 24].

Gnosis' decentralized exchange protocol relies on a matching engine which is called *multi token batch auction with uniform clearing prices*. The rationale behind this matching mechanism is to provide consistent share prices within a batch of orders and increase market liquidity. As a result, trades can be matched in rings. A ring trade is a trade matching mechanism in which buys and sells are not just equalized between two participants, but may be matched in a cycle involving three or more traders as depicted in Figure 7. During a fixed time interval, submitted orders of market participants are aggregated in an order book. Participants will then be allowed to propose a matching plan which involves prices for each share to be traded and how these orders should be executed [24].

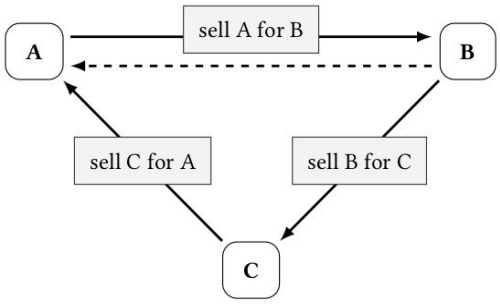


**Figure 7.** The participant of token C sells her token to the participant in possession of A, but buys token B from a different participant in what is called a ring trade (reprinted from [24]).

Amongst these proposals, the matching engine has to select the optimal one which requires defining an optimality criterion to begin with. The Gnosis matching engine

will choose the proposal which maximizes the so-called *total trading surplus*. Let $\delta_i$ denote the trading surplus for the $i$-th order and let $i = 1, ..., N$. The optimization problem can then be formulated as

$$\max \sum_{i=1}^{N} \delta_i$$

where the trading surplus for the $i$-th order is given by

$$\delta_i = (x_{eff} - x_{min}(y)) \cdot p_j .$$

Here, $x_{eff}$ is the effective amount of the token to be bought (belonging to the outcome $j$) and $x_{min}(y)$ is the least amount a trader would have accepted in exchange for $y$ tokens of another outcome. The difference is weighted by the price of the token $j$ with respect to a reference token, a common numeraire for all tokens. Consulting the total trade surplus as the optimization metric is considered fair by Gnosis' developers as it prioritizes orders where participants are willing to offer the most for the token they intend to buy [24]. The optimization problem is well-defined and the exchange rates between the tokens as well as the trade surplus generated, mimic the behavior of the same on the established prediction market Kraken closely as has been shown in [24].

Once the event has ended, the settlement process begins for which an oracle is needed. Gnosis allows developers to choose between connecting third party oracles to the core tools or using one of Gnosis' own oracles. The dominant decentralized oracle put forward by Gnosis is the "Ultimate Oracle" (Figure 8). Anyone can participate in settling the outcome of the event by voting for one outcome. Voting is done through placing ETH on an outcome. The outcome which holds the most value in ETH for a period of 24 hours will be considered the effective outcome. To prevent gamblers from placing bets on a winning outcome shortly before the end of the 24 hours period, the total amount placed on the winning outcome is limited by the amount on the remaining options. After the 24 hours period, the winning outcome can be challenged within a 12 hours window before it is deemed final. To incentivize truthful voting, all money bet on the losing outcome will be distributed towards those having voted for the winning outcome once the market is resolved [22].

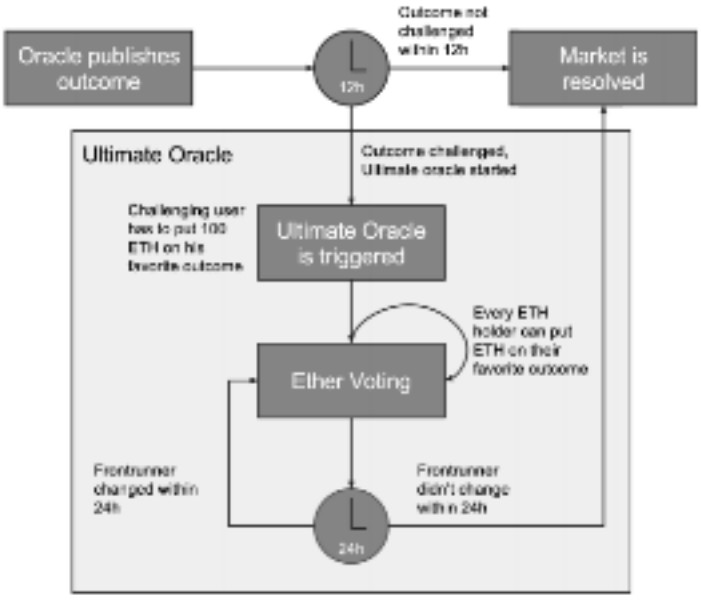


**Figure 8.** The “Ultimate Oracle” as one of the self-developed oracles provided by Gnosis for resolving markets (reprinted from [22]).

### 4.3 Key Differences Between Augur and Gnosis

Though Gnosis and Augur share main ideas, they differ in important ways, some of which will be highlighted in the following. For a more extensive comparison, the reader is directed to [25]. In contrast to Augur, Gnosis offers *conditional tokens*. In the previous subsection, it has been shown that market participants can buy a set of outcome tokens for a collateral token. In Gnosis, outcome tokens themselves follow the ERC-20 standard and can hence be utilized as collateral for outcomes of different events. If an outcome of the latter occurs, it can only be redeemed for the original collateral if the former outcome token is redeemable, i.e. the outcome it represents occurs. This chaining of events therefore allows the market to resemble conditional events and hence, conditional probabilities [22].

Another key difference between Augur and Gnosis relates to how they manage assets and fees. Ownership of outcome shares is recorded in the data structures of the smart contracts and a change in ownership can thus be followed easily. In consequence, Augur can oversee trades and collect fees accordingly. In Gnosis however, outcome tokens follow the ERC-20 standard themselves as discussed previously and hence, participants can freely dispose of these tokens. The latter has broad implications for various features of the Gnosis infrastructure as it allows users to untie the shares from the decentralized prediction market environment

and use them as ERC-20 tokens freely in any way and in any application which supports this standard. This unpredictable degree of freedom makes it difficult, if not impossible, for Gnosis to collect fees on trades of shares [25].

Despite many potential use cases, decentralized prediction markets are yet to attract broad public confidence and participation, which we analyze further in 6.1. The latter however, is a curse and a blessing at the same time since it entails the need for scalable processes. The quest for scalability of blockchain applications faces a number of challenges, from limited transaction throughput to long confirmation times or high fees. One promising solution is to outsource much of the operations off of the blockchain through the help of so called *state channels*. State channels are tools with which operations on states can be performed off the blockchain without abandoning much of the security provided by the latter [24, 26, 27, 28].

Two parties can initiate a state channel by blocking a blockchain state, for example their balances, through a smart contract. Having done that, the parties can exchange funds in multiple rounds before submitting the final state back to the blockchain. Instead of submitting each exchange as a transaction to the chain, the number of transactions is reduced to only a few transactions needed to configure and close the state channel. Both Augur and Gnosis are pushing research on state channels, though Gnosis is currently ahead, while Augur focused more on developing the oracle [24, 29].

## 5. Life Cycle of Markets on Augur in Practice

At a high level, the cycle starts with a user creating a market by defining the market question and rules of resolution. With the market created, any users can then trade by buying or selling outcome tokens. When the market reaches its close date, a designated reporter will determine the tentative winning outcome and submit the result, which will subsequently release proceeds to those who hold the winning outcome tokens. This section will cover the individual phases of the lifecycle in detail.

### 5.1 Creation

The prerequisite of creating a market is to have some cryptocurrencies ready. Market creation requires a validity bond, a no-show-bond and a transaction fee. The valid bond is payable in ETH or DAI. The market creators will be able to collect this fund back if the market resolves to anything other than invalid. The purpose of this bond is to prevent market creators from creating poorly defined markets. The no-show bond is payable in REP.  The market creators will collect this fund back when the designated reporter submits a report of the answer for the market in time. The current time limit is 24 hours after the market end time. Finally the transaction fee is nonrefundable and used for documenting the market in a blockchain.

The first step of creating a market is to figure out a question and all of its possible outcomes. A market on Augur is created to seek an objective answer to a question about a future event. Currently Augur supports 3 market types:

- YES/NO: "Will Joe Biden win the US Election in 2020?”
- Multiple Choice: "Which team will win the 2020 NBA Championship?”
- Scalar: "How many inches of snowfall in Sioux Falls in 2020?”

After setting up the question, the market creator needs to figure out the resolution information such as reporting start date and time, resolution rules and the designated reporter. The resolution information sets up a guideline on when and how the market will get resolved. If the designated reporter does not report within 24 hours of reporting start time, the market creator will lose the no-show bond and the market will also enter the Open Reporting phase, at which time anyone can report on it. Once a report is submitted, other Augur users will have the option of disputing it before the market resolves.

The market creator has the option to set up a market creation fee, which charges a percentage amount every time market shares are settled during trading or upon market resolution. It is recommended to keep the fee under 2% for the market to attract traders. On average, markets charge 1% fee.  To promote the market, the creator can set up an affiliate fee, which is the percentage of the market creator fee that affiliates will collect. This fee helps markets attract more traders by incentivizing affiliates to promote markets and collect fees every time someone follows that link and trades in a market.

The market liquidity has a direct influence into how visible the market is on Augur. A market must also have a spread of 15% or smaller, inclusive of the market creator fee, in order to be visible to traders. For instance, a Yes/No market with a .55 bid and .65 offer is discoverable because of its 10% spread, whereas a market with a .55 bid and .71 offer will not be visible due to its 16% spread. The calculation accounts for fees, so a market with a .30 bid and a .44 offer will not show up when the fees are over 1%. To ensure the newly created market being visible to users, the creator should establish an initial market liquidity by adding buy and sell offers in a tight spread with sizable volume on each side.

Before submission, the market creator should definitely review all the information carefully as Augur has a strict guideline on the validity of a market and it is quite easy for a market to be labelled as invalid. When a market is deemed invalid, the creator will lose the validity bond. According to the Augur's official guideline, a market is invalid if:

- The market question, resolution details or its outcomes are ambiguous, subjective or unknown.
- The result of the event was known at market creation time.
- The outcome was not known at event expiration time.
- It can resolve without at least one of the outcomes listed being the winner, unless it is explicitly stated how the market will otherwise resolve in the resolution details.
- The title, details and outcomes are in direct conflict with each other.
- The market can resolve with more than one winning outcome.
- Any of the outcomes don't answer the market question ONLY. (outcomes cannot introduce a secondary question)
- If using a resolution source (a source is a noun that reports on or decides the result of a market), the source's URL or full name is NOT in the Market Question, regardless of it being in the resolution details.

- If using a resolution source, it is not referenced consistently between the Market Question and Resolution Details e.g. as either a URL or its full name.
- Player or team is not in the correct league, division or conference, at the time the market was created, the market should resolve as invalid.

### 5.2 Prediction

To bet on a market, users can buy or sell shares that represent an event outcome. The price of a share is between 1 and 99 cents and roughly corresponds to the market's estimated probability of the outcome taking place; for instance, The "Joe Biden" outcome in the 2020 presidential election winner market trading at 70 cents indicates the market thinks Joe Biden has an approximately 70% chance of being the next president at that point of time.

There are multiple ways to make money as a trader. Real world catalysts may cause an event to be more or less likely to happen over time. With fluctuating share prices, it is possible to buy positions at a low cost and sell them higher as sentiment changes before the market closes. Users can also sell shares or "short" an outcome when they believe the market is too bullish on the outcome. When users hold onto their shares until the market closes, the shares representing the winning outcome as well as the sold "short" shares not representing the winning outcome will return $1.

Other than the outcomes listed by the market creator, Augur makes "Invalid Market" an outcome users can also bet on. This mechanism helps inform traders the likelihood of the market being invalid and regulate the visibility of a market based on its validity.  When the market is resolved as invalid, only the shares representing "Invalid Market" will return $1.

Since Augur operates on a peer-to-peer network, a transaction fee is required for each trade to pay for the "gas" to document the transaction on a blockchain. This fee is separated from the market creation fee and is paid to miners on Ethereum. Also this fee is only payable when the trade order is filled. In the original version of Augur, this fee to miners is paid in ETH by the users. To improve the user experience, Augur now allows traders to use DAI for both trading and paying gas fees.

In summary, the amount required for buying or selling n shares would be:

Total cost = Number of Shares * Share Price * (1 + Market Creation Fee %) + Gas

The amount of payout for a winning outcome:

Total payout = Number of Shares * (1 - Market Creation Fee %) + Gas

### 5.3 Resolution

The reporting phase begins when a market reaches its reporting start date. At a high level, the purpose of the reporting process is to have a group of participants agree on the final outcome so that the winners get paid and the market gets resolved. The process starts with designated reporting. The designated reporter chosen during market creation has 24 hours to respond with a report on the market outcome. To hold the designated reporter accountable, the reporter will have to stake REP on the reported outcome. The reported outcome is essentially the Tentative Winning Outcome. However it is open to be disputed after being submitted. When the market ends up resolving to a different outcome, the designated reporter will lose the staked REP.

If a Tentative Winning Outcome is not submitted by the designated reporter within the 24-hour time period, the market will enter the open reporting phase and the market creator will lose the No-Show Bond paid during creation. During the open reporting phase, any user with REP holding could report on the outcome and receive the forfeited No-Show Bond if the market ends up resolving to the reported outcome. Unlike the designated reporter, the open reporter is not required to stake REP on the reported outcome.

After the market receives the initial report, there will be a 24-hour time window to dispute the reported outcome. Any users can dispute by staking REP on an alternative outcome known as the Dispute Bond. However, when the market does not end up resolving to the alternative outcome, the Dispute Bond will be forfeited. On the contrary, if the dispute was successful, the users would receive a 40% return of investment on the staked REP. This mechanism is to hold disputing users accountable and also encourage users to report the truth.

If an initial report is undisputed, the market will resolve and finalize with the reported outcome. Otherwise, an alternative outcome will receive a full dispute bond and become the new Tentative Winning Outcome. This process repeats with each successive dispute round and a higher Dispute Bond is required to change the Tentative Winning Outcome. A user could submit the full Dispute Bond or fill it partially along with other users. This process can repeat up to 16 rounds until the current Tentative Winning Outcome is undisputed in a round and the REP stake required to fill the Dispute Bond doubles each round. When the current Tentative

Winning Outcome does not receive a successful challenge, the market will resolve and finalize with that outcome.

Users can provide extra support for a Tentative Winning Outcome by pre-staking REP for disputing in the favor of that outcome when it is no longer the Tentative Winning Outcome. This pre-filled stake helps accelerate a market's resolution. It also yields the 40% ROI when the market resolves to the staked-on outcome and the pre-filled stake is used in a dispute. However, if the market does not resolve to the staked-on outcome, the pre-filled stake will be forfeited.

If a dispute does not get settled at the end of 16 rounds or a dispute bond for an outcome is at least 2.5% of the total supply of REP, the market will go to the Forking phase. This very disruptive phase is designed to be a rare occurrence with at least millions of dollars at stake, which serves as the last resort to resolve a highly disputed market.

Forking creates a version or "fork" of Augur for each possible outcome of the forking market. When the Forking phase is initiated, almost everything on Augur is put on hold until the forking dispute gets resolved. This platform-wide event forces every engaging user on Augur to vote on an outcome of the disputed market within 60 days. The 60-day period is much longer than the usual dispute round because a fork's final outcome marks the end of the dispute and the platform needs to provide sufficient time for REP holders and service providers such as wallets and exchanges to prepare for forking. It is assumed that users will want to be in the version of Augur representing the truth and make it the main fork. After the 60 days, all the users would have made the decision on which fork they want to be in permanently and different forks will operate as separate entities (i.e. separate markets, users, REP, etc). Logically the fork representing the truth will become the main fork where most users will participate in.

## 6 Case Studies

### 6.1 Historical DPM Liquidity

DPMs are all relatively new, with Omen, PredIQt, Catnip, and Polymarket all launched in 2020. Only Augur has an actual history to speak of, which has been marked by initial hype, disappointment, and ups and downs as shown in 9 below. For all intents and purposes, Augur's historical liquidity represents all DPMs.

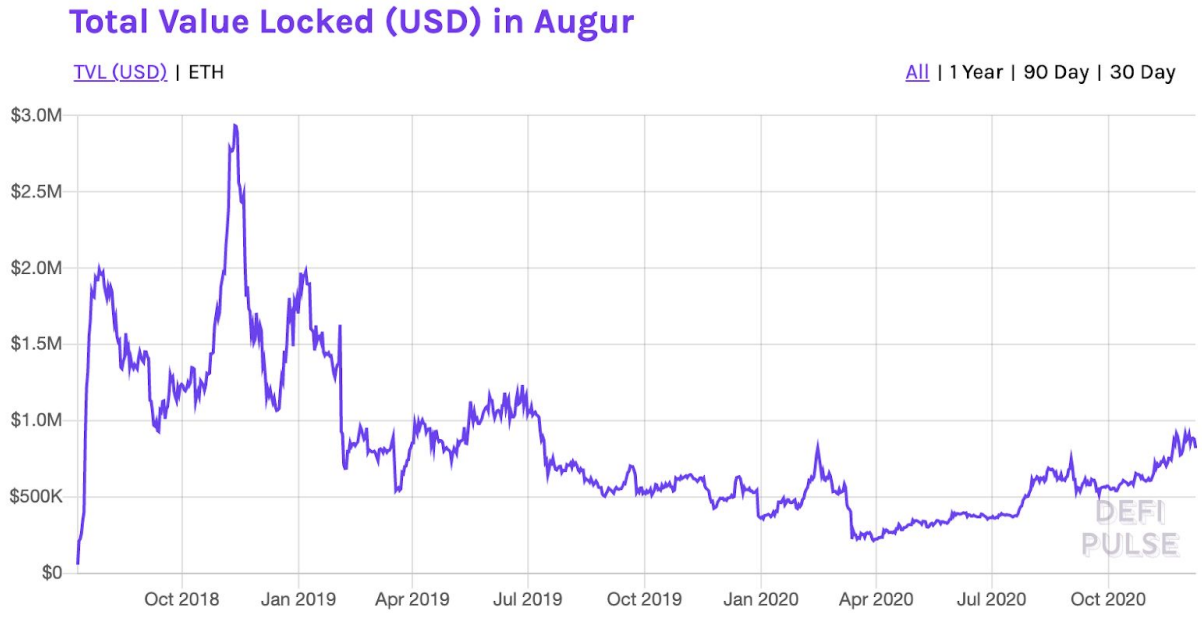


**Figure 9.** Total Value Locked in Augur from defipulse.com. Total Value Locked is the amount tied up in smart contracts for the platform.

To put Augur's size in the context of all betting markets, it is estimated that betting on U.S. elections in 2020 was worth $1 billion globally [30]. So while we do not have the equivalent of total value locked (TVL) for all betting markets over time, we can clearly see that Augur's peak TVL of $3 million USD is a small fraction of the overall $1 billion action happening in all betting markets.

### 6.2 U.S. Presidential Election (USPE) Market Size

The prediction markets related to the U.S. Presidential Election are some of the most popular markets on DPMs such as Augur, PredIQt, Catnip/Augur, and Omen. They are also big markets on centralized prediction platforms, such as PredictIt. In

this case study, we will compare the key metrics for the U.S. Presidential markets on each of these platforms, so that we can get an idea of not only how the DPMs compare to each other in terms of size, but how they compare to a centralized prediction market as well. In particular, we will focus on the Number of Markets and Total Volume.

We start off with definitions of Number of Markets and Total Volume:

- Number of Markets is the number of prediction markets. A market will have a prediction question, such as "Will Donald Trump win the 2020 U.S. Presidential election?", tradable shares representing outcomes for the question, and clearly-defined outcome criteria and deadline.
- Total Volume is the amount in USD of all shares traded to date in a market.

The data in Table 2 below is as of late November 2020, when almost all the USPE markets remained open and actively traded on the various platforms (these markets do not close until inauguration day January 21, 2021). We will discuss Table 2 further in the following sections.

**Table 2**. Total Volume for various DPMs, plus PredictIt (as a benchmark)

| Market Platform | Decentralized? | Number of USPE Markets | Estimated Total Volume in all USPE Markets (millions USD) |
|---|---|---|---|
| Augur (Native UI) | Yes | 4 | 3.038 |
| Catnip (via Augur) | Yes | 1 | 12.000 |
| PredIQt (*EOS.io blockchain*) | Yes | 2 | 0.004 |
| Omen (Gnosis) | Yes | 2 | 0.771 |
| Polymarket (Gnosis) | Yes | 2 | 5.100 |
| PredictIt | No | 47 | 302.500 |

### 6.2.2 Data Sources and Methodology

The data sources for Number of Markets was straightforward: we just relied on the advertised markets on each platform's website. For Total Volume we relied on either the websites of each platform or compiling it from blockchain transactions.

In the case of centralized platforms like PredictIt, relying on the website data was a necessity since there is no supporting blockchain transaction data. We also had to use estimation techniques for Total Volume since we didn't have detailed breakdowns on the number of shares traded for each market outcome (see A3 in Appendix). The main estimation technique was to multiply the number of shares traded in each market by $0.50, which would be the average share price assuming 2 outcomes and an even number of shares being traded on each outcome. While certainly not perfect, this technique was likely sufficient given that PredictIt Total Volume was at least an order of magnitude larger than all decentralized markets combined. In the example shown here, the market "Which party wins the presidency in 2020?" would get an estimate of 21.3 million x 0.5 = $10.65 million USD.

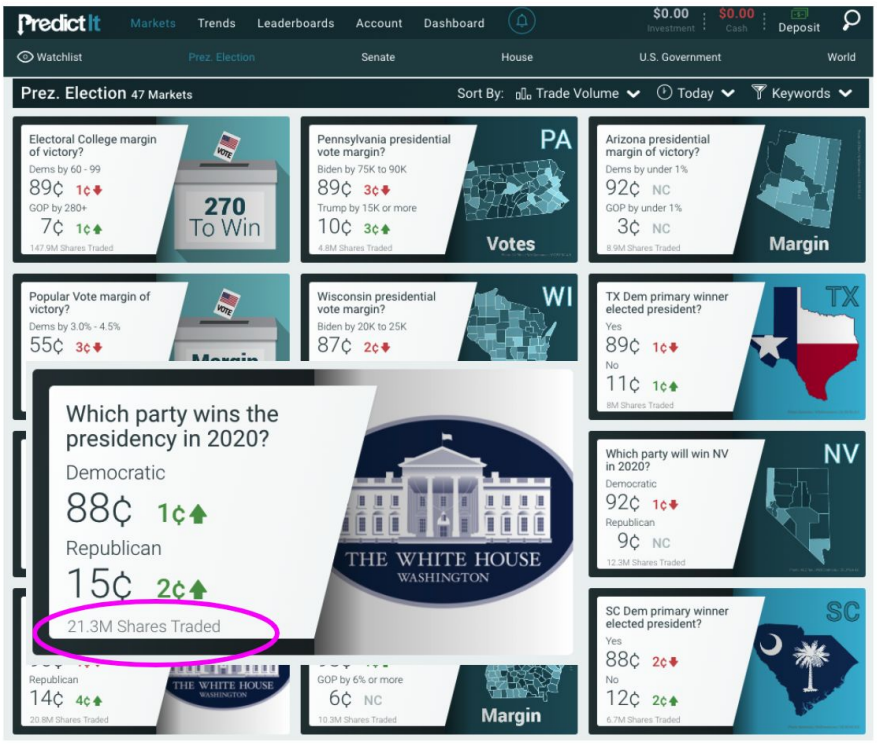


In the case of the decentralized platforms, we had the option of compiling the transactions from the Ethereum blockchain using tools like etherscan.io. We did this for Catnip successfully, by leveraging etherescan.io to search on "yTrump" and "nTrump" Augur tokens (Trump vs Biden winning the USPE, respectively). The

search as shown in Figure 10 returns transactions associated with yTrump Ethereum ERC-20 token contracts (which are wrappers for the Augur native ERC-1155 token contracts), which is most often a swap between yTrump and Dai.

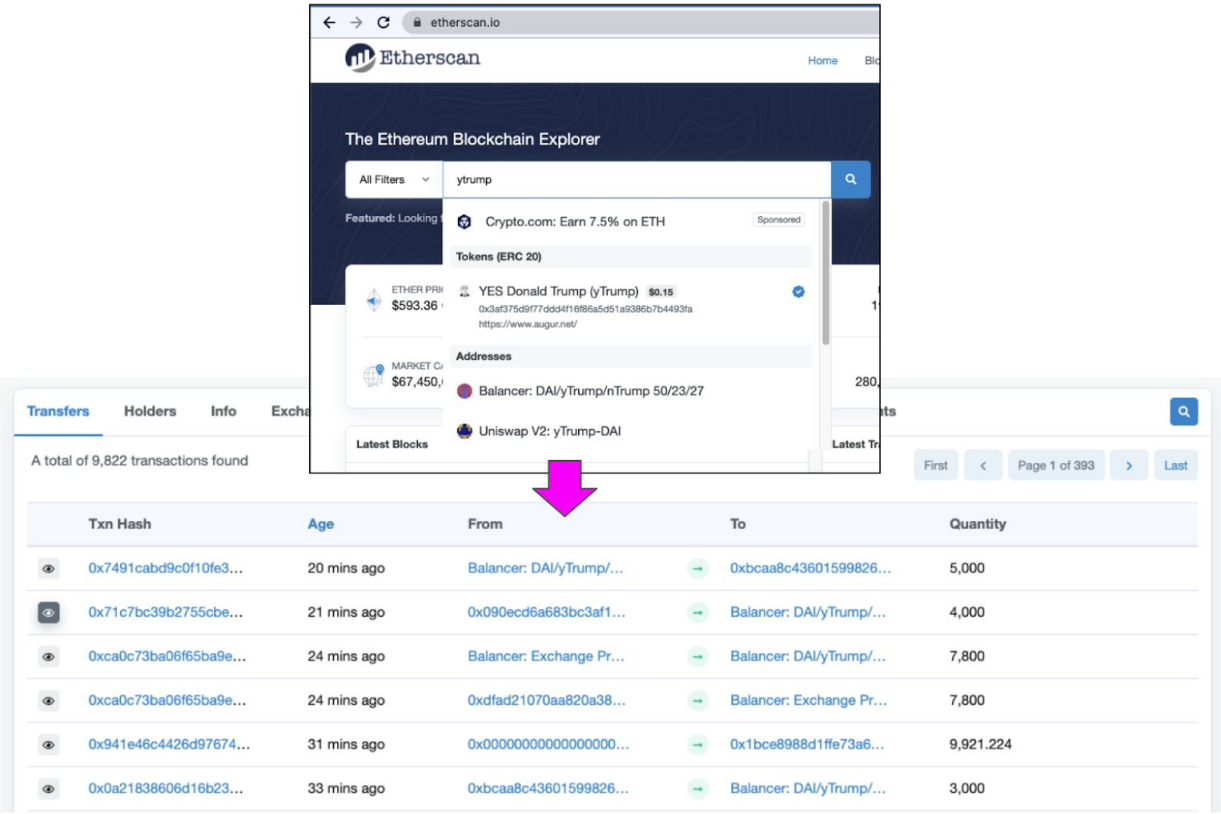


**Figure 10**. Querying the Ethereum blockchain for yTrump

After filtering all transactions found in the search, we arrived at a Total Volume for the Catnip presidential market of around $12 million USD, which was very close to the Total Volume number for Catnip shared by Augur on Twitter, that is $11.6 million USD [31].

While it is possible to query the blockchain for shares traded, we concluded that it's generally easier to use data provided on the websites of DPMs to estimate Total Volume, as we did with PredictIt. So for the other DPMs in general we used an estimation technique similar to PredictIt that estimates average share price and average token conversion to USD rates as needed. This was fine for our purposes because other than Catnip, the market sizes of the other DPMs were relatively low anyway.

### 6.2.3 Analysis of Market Sizes

The most striking conclusion from the market sizes in 6.1.1 was that as of late November, the PredictIt USPE markets were an order of magnitude larger in Total Volume than all the DPMs combined. Among the DPMs, Catnip-on-Augur was the most successful, with its Total Volume eclipsing even Augur's native interface and all other DPMs combined.

PredictIt, on the other hand, enjoyed not only superior Total Volume but also hosted 47 successful market variations on the USPE. The DPMs tended to just have a small handful of markets, and these markets were often confusing duplicates of each other. For example, the Augur native UI has these USPE market names which are confusingly similar:

1. Will Donald J. Trump win the 2020 U.S. Presidential election?
2. Will Donald Trump win the 2020 U.S. Presidential election?
3. Will Donald Trump win the 2020 U.S. Presidential election?
4. Who will win the 2020 U.S. Presidential election?

It turned out that the differences in these markets was due to differences in event expiration and not much more, so we could clearly see the fragmentation that can result from a decentralized process of creating markets. This sort of fragmentation surely hurts adoption.

PredictIt's markets were better differentiated, and they were able to generate lots of interest in outcomes like the Electoral College margin of victory, which candidate will win the popular vote in certain battleground states, and which candidate will win the popular vote. We provide a complete listing of their USPE markets in the Appendix A1, but the top PredictIt USPE markets are shown in Table 3.

It was notable that even with a large number of USPE market variations, PredictIt was able to attract large volumes of trades in dozens of these markets. In fact even if we combined the Total Volume for all of Augur Native it would only be the 24th largest PredictIt market with just over $3 million USD in Total Volume. Catnip would be the 4th largest PredictIt market (right after PredictIt's market "Which party wins the presidency in 2020?").

**Table 3.** Largest PredictIt USPE Markets.

| PredictIt Market | Shares Traded (millions) | Total Volume in USD (millions) |
|---|---|---|
| What will be the Electoral College margin in the 2020 presidential election? | 147.90 | 73.950 |
| 2020 presidential election winner? | 123.40 | 61.700 |
| Popular Vote margin of victory? | 57.10 | 28.550 |
| Which party wins the presidency in 2020? | 20.80 | 10.400 |
| Which party will win GA in 2020? | 19.10 | 9.550 |

## 6.3 Possible Inefficiencies in USPE Prediction Markets

### 6.3.1 "Irrationality" or Reality?


One of the truly fascinating aspects of the prediction markets for USPE is that even as of early December 2020 the price of Biden shares/tokens were below 90% while the price of Trump shares or tokens were above 12%. In a traditional election, it would seem irrational that prediction markets (both decentralized and centralized) attach a significant probability to Trump winning the election. In fact, shares of yTrump were as high as 0.20 on Nov 10, which was one week AFTER election day, when most media outlets including conservative-learning Fox News had called the USPE for Biden.


The apparent irrationality of high Trump share prices was a mystery to many pundits and forecasters. In late November, Nate Silver noted that the prediction markets were still assigning significant probabilities of Trump winning in states where results were already certified! Although this particular observation pertains to centralized markets (FTX, PredictIt, and BetFair), we observe that this pro-Trump bias exists in all predictions markets (both centralized and decentralized), as we will show in 6.3.2.

There are at least 2 theories that we can propose for the apparent strong support in prediction markets for Trump winning the election (against all odds). One theory is that perhaps nTrump token holders are locking in their gains by purchasing yTrump tokens, to protect against the possibility that Trump is able to "steal" the election and invalidate the widely-accepted Biden win. Another theory is that there is a Trump or Republican bias in the prediction markets, in other words there may be a disproportionate number of traders who truly believe the election is not over and are acting "rationally" based on the information they have (or believe). This latter theory gives rise to the question of whether there is enough participation in prediction markets to produce predictions of any value, and we will explore this next.

### 6.3.2 What level of liquidity might be necessary to get results that inspire any confidence?

Suppose we take as fact that Biden has won the USPE, and his elevation to the U.S. Presidency is a formality of the Electoral College. Under this belief, a rational market would place a very high share price on Biden tokens and a very low share price on Trump tokens. However, every market is a statistical sample of the beliefs of its participants, so if the market is small then it is easy to get biased results.

In Figure 11, we analyze Trump share prices in binary markets on different platforms, in an attempt to detect a relationship between Trump Share Price and Total Volume data as of early December 2020.

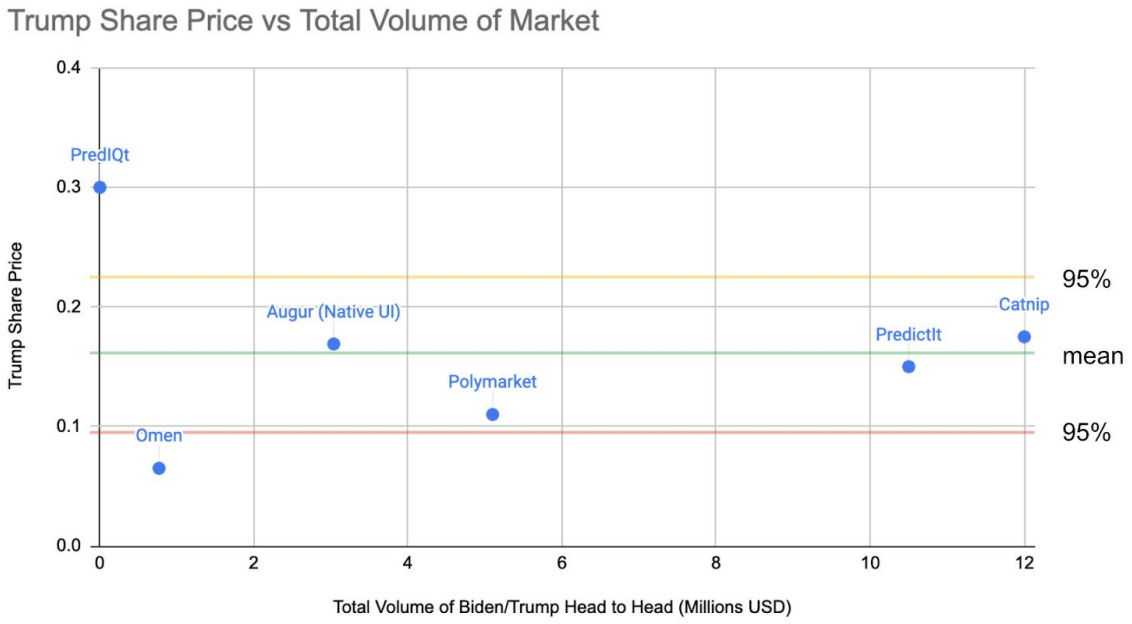


**Figure 11**. Total Trump Share vs Total Volume. For PredictIt, we just use the market whose outcome is most clearly about the head-to-head matchup between Biden and Trump, which was the market "Which party wins the presidency in 2020?" This market was around a $10.65 million market, compared to $302.5 million for all PredictIt USPE markets.

One conclusion from this analysis is that there's not a clear relationship between Trump Share Price and Total Volume. In other words, even the larger markets are displaying some apparent irrationality, in that they peg trump's chance of winning at over 15%. In fact, on the day of writing this (December 6, 2020), the yTrump token on Catnip jumped from 0.13 to 0.17. By our definition of rational, we can't say that larger markets are necessarily more rational, since smaller markets like Omen place a lower probability of a Trump win.

With such a small sample size it's hard to draw any confident conclusions, but it does appear that the smaller markets PredIQt and Omen also have the most variance, whereas the larger markets (PredictIt and Catnip) have less variance. This is what we would expect. We can only theorize what the "true" price of Trump shares would be in a very large market (say, $1 billion), but given the polarity and bifurcation of reality in U.S. politics, perhaps we're already seeing representative convergence to "true" prices on Catnip or PredictIt.

So in conclusion, it's clear from the chart above that any markets under $6 million USD in Total Volume is not enough to inspire confidence. At over $10 million each, PredictIt and Catnip's Trump share prices, while puzzling in that they seem to overprice Trump, are perhaps a reflection of what a larger population of participants might truly believe.

## 7 Challenges

### 7.1 Slow, Complex User Interfaces

For many new users of DPMs, their first experience will be the Augur Native UI. The Augur UI has gone through many iterations, from a desktop client 2 years ago to a fairly complex web user interface that launched recently in Augur v2 (Summer 2020). Ease of use has been a challenge for Augur user interfaces from its beginning [32].

Our experience with the new web user interface reflected the same experiences others had with the "slow, clunky" version 1 [33]. Our experience was that the website was often unavailable for hours at a time, and when it was available it was very slow. When the website was available, it more resembled the complex tools that perhaps a trader would use, with bids, asks, limits, order books, expirations, and other bells and whistles. To make a trade, one must have two types of currency (ETH and Dai) and be familiar with setting up a crypto wallet to connect to.

We were impressed with the newer user, simplified interfaces and experiences on Catnip, Omen, and PredictIt. As an extreme example, Catnip generated the highest Total Volume of all the DPM USPE markets by using a "less is more" strategy, where their user interface is a simple widget offering swaps between yTrump, nTrump, and Dai, and their corresponding prices. The style of Catnip's widget was apparently a familiar style employed by Uniswap [34]. Catnip was much easier to use than Augur Native UI because it was faster, more reliable, and easier to understand in our experience. In many ways, Catnip traded market breadth for ease-of-use, but given that its single USPE market exceeded all other DPM Total Volume combined, this appears to be a good tradeoff under the circumstances.

We believe that Augur will take a lot of learnings from the success of Catnip. We anticipate that future DPMs will begin converging on UI styles that are familiar and easy to use, and away from "power user" interfaces like Augur.

## 7.2 Transaction Fees

Transaction fees on Augur Native UI are known to be high [33]. Our own experience also confirms there is much room for improvement. In the scenario shown in the screenshot here from Augur Native UI, we were attempting to buy 10 Trump shares which works out to placing a bid price of $1.88 USD (1.88 Dai). While ETH gas prices can certainly fluctuate, the gas price for the $1.88 transactions was 0.0561 ETH which worked out to about $33! (Rest assured, we did not press "PLACE BUY ORDER".) While this is a fixed transaction cost, it is clear that only much larger transactions are worthwhile on Augur Native, limiting wider adoption among smaller players. We also note that there are exchange fees when buying both ETH and Dai to make this transaction, which also add a few more percentage points in fees.

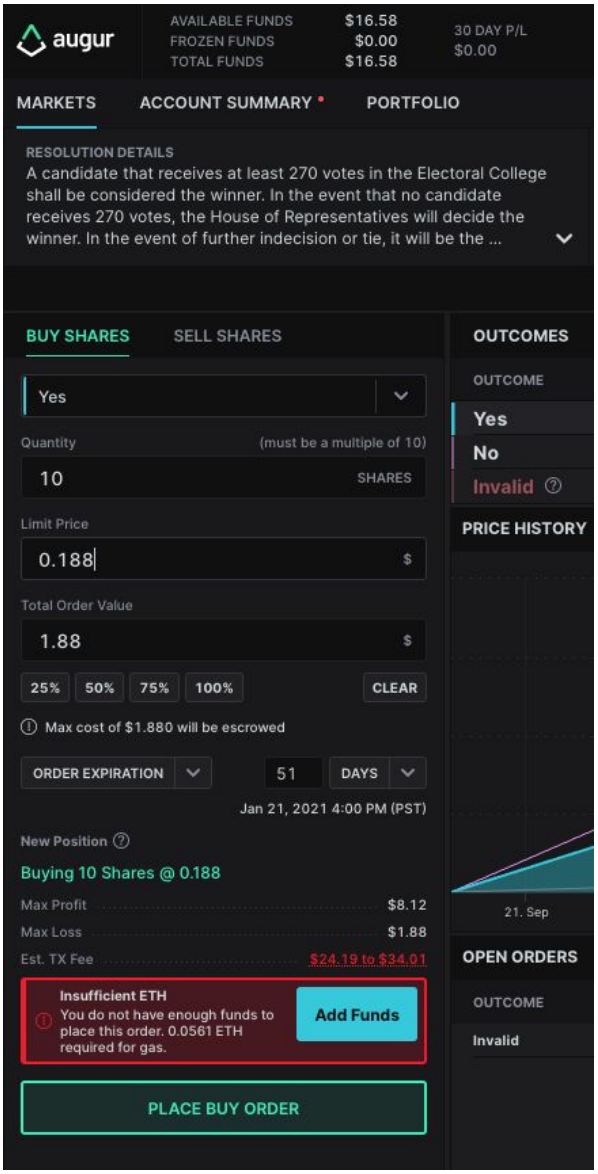


Catnip was perhaps the most successful DPM at lowering transaction fees. It accomplished this in part by modeling the transaction as a token swap [34]. The claim from [33] is that Catnip's transaction fees are 10x lower than Augur Native. We actually put this to the test, and our own similar transaction on Catnip was only 0.005 ETH or about $3.49. So the 10x lower claim was actually true in our case.

The main takeaway about transaction fees is it is no wonder that Augur Native uptake will be heavily muted when high $33+ transaction fees rule out many smaller transactions. Even on Catnip, we see fixed cost fees around $3 plus the cost of currency conversions. There is still much room to lower fees and this will go a long way to improving adoption of DPMs.

### 7.4 Other

One of the key challenges for implementing decentralized prediction markets is the verification of the outcome of events, thus putting external data on the blockchain. It is also known that the coupling of DPMs and cryptocurrency such as Ethereum has led to scalability issues. We plan to have a good understanding of the current challenges in the DPM landscape and propose viable solutions in this project.

## 8 Conclusion and Outlook

Today's DPMs, and Augur in particular, are going through adolescent growth and have not established consistent growth. In the recent USPE prediction markets, however, DPMs have been helped by more nimble newcomers such as Catnip. For all the areas where Augur struggles, in particular its UI instability/slowness, complex interface, lack of liquidity, very high transaction fees, Catnip was able to innovate and show a path forward for DPMs in general. On the share trading side, Catnip has addressed Augur's weaknesses with a simple, intuitive UI can do, greatly enhanced liquidity via the Balancer AMM, and low transaction fees through enabling predictions via simple low cost token swaps. The result was that Catnip, even as a single market, was easily larger than all DPMs in all markets combined ($13 million for Catnip versus less than $10 million for Augur Native, Omen, Polymarket, and PredIQt combined).

While Catnip's success provides a blueprint for Augur to address its adoption woes, the centralized prediction markets like PredictIt show what is possible if Augur or Gnosis if they can improve their UI, lower their fees, and/or increase their liquidity. PredictIt, both in its market variety and total volume, are up to an order of magnitude larger than all of the DPMs combined, using the USPE markets as a benchmark. Augur has many advantages over centralized markets, including greater transparency through the blockchain, higher or unlimited betting limits, and less regulation. Therefore, issues like UI, fees, and liquidity seem very solvable especially since Catnip has shown the way. So while DPMs are still getting their footing in terms of consistent growth, they have a blueprint to correct some of the obvious issues inhibiting their adoption. Our prediction (pun intended) is that DPMs will grow their volumes by an order of magnitude, to the size of PredictIt, by the next US Presidential Election.

Moreover, although it has been suggested that price determination using LMSR is expensive, the LMSR market maker is well studied and tested in the practice and it

is designed specifically for the prediction market use case (e.g., Gnosis). From a statistics perspective, conditional tokens were built for LMSR use and have a very nice property in a sense that the definition of price in LMSR is actually a softmax function that is used in various multiclass classification methods. In reinforcement learning, a softmax function can be used to convert values into action probabilities. As GPU/TPU/others and layer-2 developed, LMSR combined stochastic gradient descent to calculate price might be possible in the future.

**References**


[1] "Wisdom of the crowd", https://en.wikipedia.org/wiki/Wisdom_of_the_crowd. Last accessed 11 Dec 2020.

[2] M. R. Blouin and R. Serrano, "A Decentralized Market with Common Values Uncertainty: Non-Steady States," Rev. Econ. Stud., vol. 68, no. 2, pp. 323–346, 2001, doi: 10.1111/1467-937x.00171.

[3] E. Segal-Halevi, A. Hassidim, and Y. Aumann, "MUDA: A truthful multi-unit double-auction mechanism," arXiv. 2017.

[4] P. Huang, A. Scheller-Wolf, and K. Sycara, "Design of a multi-unit double auction e-market," Comput. Intell., vol. 18, no. 4, 2002, doi: 10.1111/1467-8640.t01-1-00206.

[5] M. G. Daniels, J. D. Farmer, G. Iori, and E. Smith, "How storing supply and demand affects price diffusion," arXiv Prepr. cond-mat/0112422, 2002.

[6] R. Hanson, "Logarithmic market scoring rules for modular combinatorial information aggregation," J. Predict. Mark., vol. 1, no. 1, 2012, doi: 10.5750/jpm.v1i1.417.

[7] Y. Zhang, X. Chen, and D. Park, "Formal specification of constant product (xy= k) market maker model and implementation. 2018," 2018.

[8] M. Egorov, "StableSwap-efficient mechanism for Stablecoin liquidity," 2019.

[9] G. Angeris and T. Chitra, "Improved Price Oracles: Constant Function Market Makers," arXiv Prepr. arXiv2003.10001, 2020.

[10] "Curve.fi", https://www.curve.fi. Last accessed 11 Dec 2020.

[11] E. Fröberg, G. Ingre, and S. Knudsen, "Blockchain and prediction markets: An analysis of three organizations implementing prediction markets using blockchain technology, and the future of blockchain prediction market." 2018.

[12] J. Peterson, J. Krug, M. Zoltu, A. K. Williams, and S. Alexander, "Augur: a decentralized oracle and prediction market platform," arXiv Prepr. arXiv1501.01042, 2015.

[13] J. Peterson, J. Krug, M. Zoltu, A. K. Williams, and S. Alexander, "Augur: a Decentralized Oracle and Prediction Market Platform (v2.0)." [Online]. Available: https://www.wunderground.com/history/airport/KSFO/2018/4/10/.

[14] "Augur Homepage", https://augur.net/faqs. Last accessed 30 Nov 2020. .
[15] G. Caldarelli, "Understanding the Blockchain Oracle Problem: A Call for Action," Information, vol. 11, no. 11, p. 509, 2020.
[16] A. Egberts, "The Oracle Problem-An Analysis of how Blockchain Oracles Undermine the Advantages of Decentralized Ledger Systems," Available SSRN 3382343, 2017.
[17] "Ethereum Homepage", https://ethereum.org/en/developers/docs/standards/tokens/erc-20/. Last accessed 30 Nov 2020..
[18] "Augur Homepage", https://augur.net/blog/augur-master-plan/. Last accessed 30 Nov 2020. .
[19] A. Berentsen and F. Schar, "Stablecoins: The quest for a low-volatility cryptocurrency," Fatas A.(a cura di), Econ. Fintech Digit. Currencies, pp. 65–71, 2019.
[20] "Augur Homepage", https://augur.net/blog/integrations-overview/, Last accessed 30 Nov 2020.
[21] "Omen", https://omen.eth.link/. Last accessed 03 Dec 2020.
[22] "gnosis_whitepaper", https://github.com/gnosis/research/blob/master/gnosis-whitepaper.pdf.
[23] "Gnosis Prediction Market Contracts Documentation", https://gnosis-pm-contracts.readthedocs.io/en/latest/index.html.Last accessed 03 Dec 2020 .
[24] "Gnosis Github Page", https://github.com/gnosis/dex-research/blob/master/dFusion/dfusion.v1.pdf, Last accessed 03 Dec 2020 .
[25] A. Akhunov, "Hopefully impartial comparison of Gnosis and Augur, https://medium.com/@akhounov/hopefully-impartial-comparison-of-gnosis-and-augur-f743d11d6d37," Apr. 22, 2017.
[26] P. McCorry, S. Bakshi, I. Bentov, S. Meiklejohn, and A. Miller, "Pisa: Arbitration outsourcing for state channels," in AFT 2019 - Proceedings of the 1st ACM Conference on Advances in Financial Technologies, Oct. 2019, pp. 16–30, doi: 10.1145/3318041.3355461.
[27] J. Coleman, L. Horne, and L. Xuanji, "Counterfactual: Generalized state channels," Acessed Nov, vol. 4, p. 2019, 2018.
[28] J. Stark, "Making Sense of Ethereum's Layer 2 Scaling Solutions: State Channels, Plasma, and Truebit, https://medium.com/l4-media/making-sense-of-ethereums-layer-2-scaling-solutions-state-channels-plasma-and-truebit-22cb40dcc2f4," Feb. 12, 2018.
[29] "State Channels", https://www.jeffcoleman.ca/state-channels/. Last accessed 03 Dec 2020.

[30] "Betting on US elections is worth $1 billion globally", https://www.trtworld.com/magazine/betting-on-us-elections-is-worth-1-billion-globally-41225. Last accessed 11 Dec 2020.
[31] "Twitter: Augur@AugurProject", https://twitter.com/AugurProject/status/1328530227421532161. Last accessed 11 Dec 2020.
[32] "Why is no one really using Augur?", https://medium.com/coinmonks/why-is-no-one-really-using-augur-161448a8e198. Last accessed 11 Dec 2020.

[33] "5 Years After Launch, Predictions Market Platform Augur Releases Version 2", https://www.coindesk.com/5-years-after-launch-predictions-market-platform-augur-releases-version-2, Jul. 29, 2020. Last accessed 11 Dec 2020.
[34] "Catnip Exchange Makes Prediction Markets Usable", https://dexplain.com/catnip-exchange-makes-prediction-markets-usable/, Oct. 12, 2020. Last accessed 11 Dec 2020.

## Appendix

### A1. Data Referenced

47 U.S. President related markets as of 11/24/2020 on PredictIt:

| PredictIt Market | Shares Traded |
|---|---|
| What will be the Electoral College margin in the 2020 presidential election? | 147.90 |
| What will be the vote margin in the 2020 presidential election in Pennsylvania? | 4.90 |
| Popular Vote margin of victory? | 57.10 |
| Arizona presidential margin of victory? | 8.90 |
| Wisconsin presidential vote margin | 2.30 |
| TX Dem primary winner elected president | 8.00 |
| Penn presidential MOV? | 12.00 |
| Will Trump win MI, WI, or NV? | 5.10 |
| Which party will win NV in 2020? | 12.30 |
| Which party will win PA in 2020? | 15.40 |
| Which party wins the presidency in 2020? | 20.80 |
| SC Dem primary winner elected president? | 6.70 |
| Georgina presidential margin of victory? | 10.30 |
| Will Trump win PA, AZ, or GA? | 5.90 |
| Which party will win GA in 2020? | 19.10 |
| MA Dem primary winner elected president? | 6.40 |
| Which party will win MI in 2020? | 16.30 |
| Woman VP in 2020? | 8.30 |
| Which party will win AZ in 2020? | 18.20 |
| Georgia presidential vote margin? | 9.80 |
| 2020 presidential election winner? | 123.40 |
| Election results versus the polls? | 8.60 |

| | |
|---|---|
| State wil the smallest MOV in 2020? | 16.90 |
| Will Fox un-call any jurisdiction? | 0.68 |
| Which party will win WI in 2020? | 14.70 |
| Texas presidential margin of victory? | 6.10 |
| Biden's margin in WI shrinks by 100+? | 0.05 |
| Trump loses any state he won in 2016? | 5.40 |
| Turnout in the presidential election? | 7.70 |
| Popular vote winner wins Electoral College? | 3.10 |
| Popular vote majority for president? | 0.75 |
| Trump files for president before 2022? | 0.03 |
| Will there be a recount in Wisconsin? | 1.10 |
| Tipping point jurisdiction in 2020? | 6.10 |
| Trump win any state he lost in 2016? | 3.50 |
| Will there be a recount in AZ? | 0.47 |
| Biden's margin in GA shrinks by 100+ | 0.04 |
| Clean sweep for Democrats in 2020? | 5.50 |
| Pelosi Becomes Acting President on 1/20? | 0.61 |
| Presidential vote % not for Dem/GOP? | 1.40 |
| House delegations won by GOP? | 0.48 |
| Woman president in 2020? | 0.79 |
| Which party wins the presidency in 2024? | 0.27 |
| Votes for Kanye in 2020? | 0.61 |
| Iowa presidential margin of victory? | 1.50 |
| Harris files for President before 2023? | 0.01 |
| Pence files for president before 2023? | 0.01 |

**A2. 4 non-zero U.S. President related markets as of 11/24/2020 on Augur**

| Market | Expiration | Total Volume | Open Interest | Y |
|---|---|---|---|---|
| Will Donald J. Trump win the 2020 U.S. Presidential election? | 20 Jan 2021 | $697,526 | $590,451 | |
| Will Donald Trump win the 2020 U.S. Presidential election? | 21 Jan 2021 | $1,642,498 | $6,779,096 | |
| Will Donald Trump win the 2020 U.S. Presidential election? | 7 Jan 2021 | $182,287 | $158,470 | |
| Who will win the 2020 U.S. Presidential election? | 20 Dec 2020 | $251,339 | $162,760 | |
| Who will win the 2020 U.S. Presidential election? | 14 Jan 2021 | 0 | 0 | |
| Will Donald J. Trump win the 2020 U.S. Presidential election? | 31 Dec 2020 | 0 | 0 | |
| Which party will win Nevada in the 2020 U.S. Presidential election? | 6 Jan 2021 | 0 | 0 | |
| Which party will win Florida in the 2020 U.S. Presidential election? | 7 Jan 2021 | 0 | 0 | |
| Lawful or unlawful, is Donald Trump the highest power of the United States of America, at 12:00 AM UTC on July 21st, 2021? | 21 Jul 2021 | 0 | 0 | |
| Which party will win the 2020 U.S. Presidential election? | 31 Dec 2020 | 0 | 0 | |
| Will Donald Trump win the popular vote in 2020? | 19 Jan 2021 | 0 | 0 | |
| Will Joe Biden be sworn in as the 46th President of the United States of America before January 21st, 2020? | 20 Jan 2021 | 0 | 0 | |
| Will Donald John Trump win the 2020 U.S. Presidential election? | 21 Jan 2021 | 0 | 0 | |
| Will the 2020 US Presidential Election Be Decided By December 1st 2020? | 1 Dec 2020 | 0 | 0 | |
| Which party will win Pennsylvania in the 2020 U.S. Presidential election? | 7 Jan 2021 | 0 | 0 | |
| Will either Joe Biden or Donald Trump concede defeat in the 2020 US elections by December 1st, 2020? | 1 Dec 2020 | 0 | 0 | |

## A3. Email correspondence with Parker Howell of PredictIt



Gmail

**Quoc Le <quocnle@gmail.com>**

**RE: becoming a research partner**

**Parker Howell** <Parker.Howell@aristotle.com> Mon, Nov 30, 2020 at 8:22 AM
To: "quoc.le@columbia.edu" <quoc.le@columbia.edu>

Hi Quoc,

I couldn't confirm or deny your estimate even if I knew the actual figure. Your assumption that on average shares will be worth $0.50 seems reasonable, so if you've added up the quantity of shares traded that you see on the site I think that is the best approach.

Best,

Parker

**From:** Quoc Le <qnl2000@columbia.edu>
**Sent:** Sunday, November 29, 2020 2:26 PM
**To:** Parker Howell <Parker.Howell@aristotle.com>
**Subject:** Re: becoming a research partner

Thank you Parker, that's reasonable. The main data point I was seeking is to estimate the size of the current U.S. Presidential Election markets in PredictIt. Just from the PredictIt UI, I estimated that PredictIt currently has over $300 million USD in shares traded for the 47 "Prez Election" markets. This is based on the rough assumption that shares are, on average, worth $0.50 in these markets (assuming binary outcomes - I know shares in some markets could be worth less on average due to some markets having more than 2 outcomes).

If you don't mind, I wanted to get your opinion if that would be a reasonable estimation (I won't quote you). This is just for a class paper for an Fundamentals of Blockchains course, and not for any kind of published paper.

Thanks,

Quoc