GNOSIS

Whitepaper - 05.04.2017

Contents

1. Executive Summary

Prediction markets are poised to become one of the most disruptive innovations in capital markets and data science since the beginning of the Information Revolution. First proposed in the early 90s, prediction markets have yet to attract mass attention in the realm of forecasting and decision-making despite their documented efficacy for information aggregation. This is largely due to over-regulation in many of the world's leading financial sectors. With the invention of powerful, peer-to-peer computing technologies such as Ethereum and Bitcoin, the scientific exploration of market-based forecasting can proceed uninhibited and at a rate and scale previously unimaginable. Our team believes undoubtedly that prediction markets will disrupt some of the largest existing industries in the near term. Looking forward, we expect that the Gnosis prediction market platform will form the basis for machine information economies on a global scale.

In order for a prediction market platform to become truly disruptive, it must be universal and draw from a global liquidity pool. The platform must be decentralized, permissionless, and trustless for such a liquidity pool to exist. With these requirements in mind, the Gnosis team has selected the Ethereum network as the core protocol upon which the platform will be built.

1.1 Problem Overview

Generally speaking, the Information Revolution has made it easier for individuals to quickly retrieve data about any topic. Despite the ease of access we enjoy today, this form of information aggregation still requires a great deal of coordination to be effective. More often than not, the data is severely lacking in context and objectivity and requires heavier lifting to produce actionable information for use in decision-making processes. The reason for this is straightforward: written information is inextricably linked to the writer's individual biases and agenda, making it difficult to delineate useful information from opinions or intentional misinformation. In other words, it's easy to find what people have said but hard to ascertain what they actually believe.

Financial markets are particularly interesting in this regard in that the act of speculation elicits a highly effective form of information aggregation that requires no coordination (i.e. the "invisible hand") and more closely mirrors individual beliefs. Principally, market speculators who believe they

have superior information buy shares when they believe a company is undervalued and sell shares when they believe the company is overvalued. A monetary incentive exists to "update" a common data point (i.e. share price) when there is profit potential, and there is a disincentive to misreport in the form of financial loss. The resulting equilibrium share price reflects the prevailing market-wide sentiment about a company's value at any given time. In summary, information aggregation occurs with *skin in the game* - a characteristic that:

1. effectively glues an individual's action to their privately held beliefs and
2. is absent from other methods for information aggregation such as polling.

This is vital for understanding the principal function of prediction markets.

A prediction market, in essence, is a vehicle for aggregating information about the expected outcome of a future event.[1] Unlike a traditional financial market, prediction markets frame themselves as questions about the future, typically with binary outcomes. For example: Which presidential candidate will win the 2016 election? Shares are divided among pre-defined options (e.g. Hillary Clinton, Donald Trump, Other) with corresponding share prices equaling $1. Each option's share price reflects its probability of occurrence. So long as an individual believes they have superior information about the event in question, s/he has an incentive to purchase shares that reflect his/her beliefs about the outcome, thereby updating information captured by the prediction market. At the market's conclusion, the winning option's shares become redeemable for $1, while all other shares become worthless. Individual actors who purchased the winning shares receive profit equal to $1 - purchase price x number of shares.

Ex: Which presidential candidate will win the 2016 election?

(A) Hillary Clinton ($0.40) = 40% chance of winning
(B) Donald Trump ($0.50) = 50% chance of winning
(C) Other ($0.10) = 10% chance of winning

(A) + (B) + (C) = $1.00 (100%)

Over the last several decades, prediction markets have seen a surge in use due to their superior ability to effectively aggregate all available information relevant to an event's outcome. Prediction markets have already been implemented with success for a variety of applications. Initially, these markets were limited to academic purposes, the first of which being the Foresight Exchange. Later, prediction markets were tested by the intelligence industry through projects such as DARPA's FutureMAP, which was, "an experiment to see whether market-generated predictions could improve upon conventional approaches to forecasting." Perhaps most straightforward and general was Intrade's prediction market, which allowed for event futures on the outcome of decidable events such as elections, current events, and sports. During its existence, Intrade showed that such markets could garner significant volume and estimate the likelihood of potential outcomes with greater accuracy than traditional polling methods. More recently, prediction markets have found use as internal tools to inform organizations, such as large corporations and nonprofits.

Prediction markets have limited accessibility, liquidity, and market variety resulting from strict financial and gaming regulation in the U.S. that extends well-beyond the country's borders. In 2012, Ireland-based Intrade was sued by the CFTC in violation of its ban on off-exchange options trading, leading to the site's ultimate closure in March 2013.[2] Per the U.S. Commodity Exchange Act, it is illegal to solicit U.S. based customers to buy and sell commodity options unless classified as exempt or listed on a CFTC-regulated exchange. Intrade's sudden closure raised suspicions about market

[1] Also known as predictive markets, information markets, decision markets, idea futures, event derivatives, or virtual markets

[2] http://www.cftc.gov/PressRoom/PressReleases/pr6423-12

manipulation and around claims about its markets' liquidity.[3] Similarly, prediction markets have been met with resistance within corporate environments. Despite evidence of early successes for corporate decision making in forward-thinking organizations such as Google, prediction markets have failed to garner popular support from product/project managers for a variety of reasons, including manager bias, organizational friction, and a lack of well-established guidelines for facilitating successful internal markets.

Despite the aforementioned challenges, prediction markets have seen a resurgence of interest both in academia and the private sector. More recently, New Zealand-based prediction market PredictIt has been able to operate under the benediction of the CFTC, albeit with severe limitations. PredictIt positions itself as "an experimental research and educational facility of Victoria University of Wellington," rather than a betting site. Pursuant to a letter of non-action received from the CFTC in October 2014, the platform's markets are restricted to 5,000 total traders per contract, with a limit on individual investment of $850.[4] Furthermore, use of the platform is illegal in the states of Washington and Nevada.

We must face a hard truth that prediction markets, despite being rigorously tested, cutting-edge instruments for forecasting the future, will never reach their full potential if built on 21st century database technologies. Traditional prediction market applications that operate on centralized platforms will tend towards proprietary designs, siloing data and reducing overall liquidity - a recipe for impotent markets that leave much to be desired in terms of accuracy and precision. Furthermore, these databases lack the resilience necessary to resist censorship and reach untapped liquidity pools across the globe, the effect of which suffocates any prediction market's viability as a platform and stunts its growth as a means of information exchange. Because these efficient, flexible, and scalable markets for specialized information exchange do not yet exist, investors and speculators are forced to accept the high costs associated with constructing custom financial instruments in highly regulated environments.

1.2 Mission Statement

> "Our mission is to build an accessible prediction market platform enabling the free flow of useful information."

Gnosis will be a disruptive force driving change in a number of systemically important global markets, including finance, gambling, insurance, and information. Gnosis prediction markets will also find applications in new forms of distributed, market-based governance protocols, and will provide unique incentivization opportunities for both local and global economies.

Gnosis is well-positioned as a medium for a long-term shift toward information arbitrage economies that will power the Internet of Things, as well as more advanced forms of artificial intelligence. We believe that we are on the cusp of a Cambrian explosion of machine intelligence that will leverage a global liquidity pool of information for decision-making and will be deeply interwoven on a shared blockchain fabric such as Ethereum. Decentralized prediction markets seeded on Gnosis will be the ideal medium of exchange for these intelligent agents.

[3] `http://www.newyorker.com/news/john-cassidy/what-killed-intrade`

[4] `https://www.predictit.org/Home/TermsAndConditions`

1.3 Core Objectives

1.3.1 Build the World's Most Efficient Forecasting Tool

Prediction markets can enable a more efficient and informed world. Prediction markets and oracles will bridge real world events to the blockchain, thereby strengthening its value as an authoritative source of truth about the world.

1.3.2 Create the "Google" of Customized Information Searching

Gnosis enables anyone to ask a question and fund the search for answers. This creates new economic opportunities for subject matter experts and more efficient avenues for crowdsourcing and aggregating information. The power of "search" is decentralized and inclusive.

1.3.3 Become the Standard for Predictive Assets

Gnosis seeks to establish a global, open prediction market platform with a single liquidity pool. This limitless resource enables the simple creation of custom prediction market applications and embodies a flexible marketplace for blockchain oracle services.

2. Token Mechanism

The token sold during the token launch is known as the Gnosis Token, or GNO. This is the only time that these tokens can be created, and therefore the total supply of GNO is fixed.

Fees, similar to those of a trading market, will be charged to participants on the Gnosis Services and Applications layers (but as a reminder, not the bare bones Core layer). These fees will initially be denominated in cryptocurrency, namely BTC or ETH. Gnosis seeks to not only create interesting software, but also a community of those interested in sharing their wisdom on Gnosis markets. To do this, we needed to create a model that lowers the barrier to entry for repeat users (e.g. having to pay BTC/ETH repeatedly). Therefore, in addition to paying this fee in BTC or ETH, Gnosis ecosystem participants will be able to pay the fee in Wisdom, or WIZ, tokens.

Gnosis Wisdom (WIZ) can be used to pay platform fees on the Services layer, subsidize the fees of other participants, provide initial subsidies for markets, or for market trading. WIZ will be pegged to $1 USD worth of fees. In this way, WIZ acts as a coupon for $1 of use within Gnosis.

Gnosis tokens (GNO) are the generator for Wisdom token (WIZ) creation. WIZ can only be created via activating the utility of the Gnosis (GNO) tokens. This is done via a smart contract system. The smart contract works as follows: GNO token holders agree to "lock" their tokens in a smart contract (30-365 days). A multiplier is added for longer lock durations. The smart contract determines the user selected lock duration and applies that duration to a formula that is designed to regulate the supply of WIZ tokens currently in use. Prior to locking their GNO tokens in the smart contract, users will be able to see exactly how much WIZ they will receive as a result of executing the smart contract. Once users execute the contract, 30% of their WIZ will be distributed for immediate use, and the remaining 70% will be distributed proportionally over the locked duration. Once the lock duration expires, the locked GNO ceases to generate WIZ and the GNO becomes freely transferable by the holder. There is no limit (other than duration) for how many times GNO tokens may be used to create WIZ.

How Can Gnosis Remain Viable if Participants Choose Not to Pay in WIZ?

A core value proposition of Gnosis (and decentralization) is to guarantee future characteristics of platforms to both users and developers without relying on the trustworthiness of an operating company. In order to do this, elements including fee rates, must be codified into the software itself. It

is expected that WIZ will be the overwhelmingly predominant method for paying fees in the Gnosis ecosystem. In the unexpected event that this is not true, and users are paying in BTC or /ETH, the platform may become vulnerable to low-fee copycats or potentially even illegal forks of the Gnosis codebase.

These alternative platforms may logically cause erosion of the Gnosis userbase, subsequently triggering justified loss of developer confidence that their created markets and applications will remain viable on Gnosis. In order to avoid this scenario, we designed a fee-reduction mechanism to bolster competitiveness of the Gnosis platform. The result is added confidence for developers and partners that Gnosis is the infrastructure they should be building markets on.

NOTE: It is unlikely that this mechanism will be used as game theory and expectations point to users predominantly paying fees in WIZ. In the event this mechanism is triggered, we expect the occurrence to be extremely rare.

Two core requirements for the fee reduction mechanism is that it is both decentralized and costly. The mechanism must be costly in order to eliminate spam or manipulation. The core functionality of the mechanism is as follows: All fees paid in BTC/ETH/Tokens go to an auction contract outside the control of the Gnosis team. If fees exist in the auction contract, any GNO token holder can submit a bid, bidding their held GNO against some amount of fees contained in the auction contract. If the bid is accepted, the GNO will then enter the auction contract and the user will receive the fees specified. When the user's GNO enters the auction contract, the fee reduction mechanism will be triggered causing a reduction in fees on Gnosis proportional to the total amount of GNO held in this auction contract. The auction contract is one-way and GNO cannot leave this wallet.

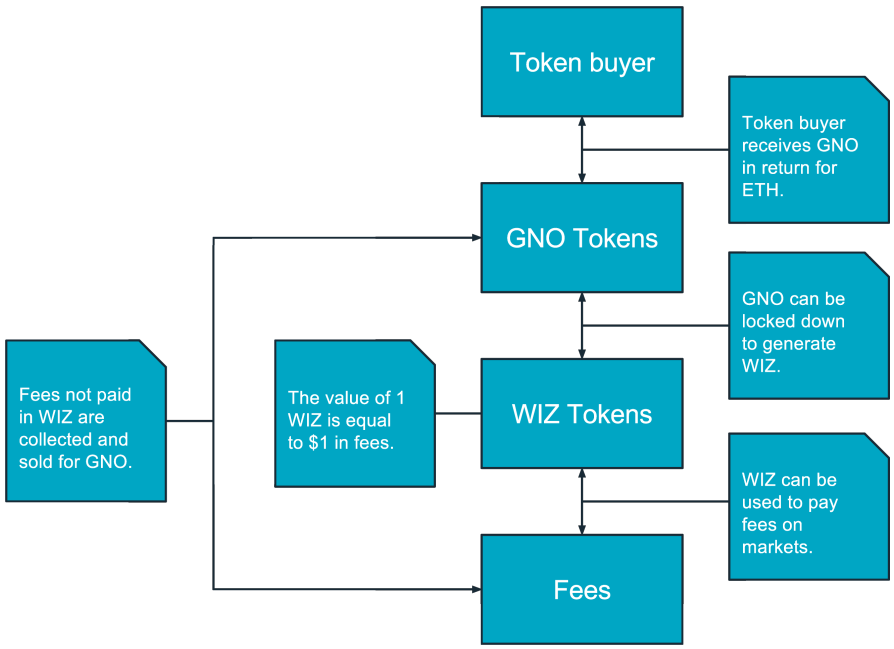

Figure 2.1: Examples of GNO and WIZ Utility

3. Platform Model

The Gnosis platform will be composed of three primary layers: Core, Services, and Application.

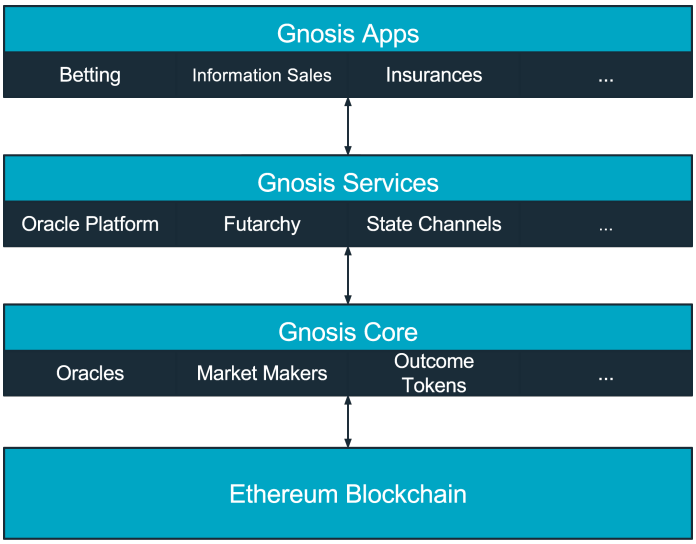

Figure 3.1: Primary layers of the Gnosis platform

3.1 Gnosis Layers

3.1.1 Layer One: Gnosis Core

The Core layer provides the foundational smart contracts for Gnosis use: event token creation and settlement, a market mechanism, oracle, and a management interface. This layer is and always will be free and open to use. Creating new markets is near zero marginal cost, and to remain competitive fees will have to approach zero. Instead of grasping at the maximum possible fees while remaining competitive, we feel that it is prudent to eliminate fees at the most basic contract level. It should be in

every party's best interest to use the existing open source and feeless contracts instead of deploying their own version.

3.1.2 Layer Two: Gnosis Services

The Gnosis Services layer will offer additional services on top of Gnosis Core and will use a trading fee model. These services will include a state channel implementation[1], new market mechanisms, stablecoin and payment processor integrations, open source template applications, application customization tools, and the oracle marketplace. More features may be introduced as deemed useful. These components are necessary for most consumer applications building on Gnosis.

State channels are a prerequisite for betting and financial applications requiring thousands or more transactions per second. Without stablecoins, market participants are subject to the volatility of the cryptocurrency which the market is denominated in and the event outcome that they are predicting. Application templates, customization tools, and advanced oracle selection will allow us to execute on our vision of lowering the barrier to entry for new prediction market based applications by at least two orders of magnitude. While some applications and participants will interact with Gnosis on the Core level, we are confident that these services will provide a compelling reason for Services level use.

3.1.3 Layer Three: Gnosis Applications

On top of the Services layer (or in some cases, just Gnosis Core) is the Gnosis application layer. These applications are primarily front-ends that target a particular prediction market use case and or customer segment. Some of these applications may be built by Gnosis, while others will be built by third parties. Our vision for Gnosis is to have a wide variety of prediction market applications built atop the same platform and liquidity pool. These applications will likely charge additional fees or use alternative business models such as market making, information selling, or advertising. As we'll see in the next section on tokens, many Gnosis applications may include token holding as a core component of their business model.

[1] `https://media.consensys.net/state-channels-ethereum-is-open-for-business-5b7cd4d7506c#.d95j8n6gh`

4. Gnosis Applications

With the Gnosis prediction market platform serving as a global liquidity hub, decentralized application developers will be able to create new classes of predictive assets that can be used in any number of simple or complex applications. The following section will introduce a set of innovations that are readily implementable in existing markets with the use of Gnosis prediction markets. It will also attempt to define entirely new verticals that are made possible through the use of predictive assets.

4.1 Financial Instruments

Prediction markets can enable the creation of financial instruments that track stock price or commodity value with greater specificity than existing derivatives. If we conceptualize traditional financial instruments as expressions of economic value, one could argue that the "expressiveness" of current market offerings is limited to statements of ownership in an asset (e.g. currencies, equities), of financial relationships between economic entities (e.g. bonds), and meta-statements about value relative to an instrument (e.g. derivatives). Prediction markets enable more nuanced and specific expressions about economic events, which in turn signal value more explicitly (along with risk) at both the macro and micro-economic level.

For example, markets can be created asking, "What will this corporation's Q4 gross revenue be on a specified date?" One might imagine that information drawn from this and many adjacent markets could be used as an input to inform more advanced predictive/decision-making models in finance, government, insurance, and beyond. Such a market could inform analysts' earnings per share (EPS) estimates for the quarter; alternatively, markets can be constructed to predict EPS itself, sidestepping the need for analyst forecasts altogether. More readily available and reliable information in this area can lead to better price discovery, minimizing both short-term volatility and long-term risk.

Alternatively, prediction markets can be used to create pegged and stable currencies. For example, an EtherUSD currency can be implemented by creating a scalar market which asks, "What will the Ether/USD exchange rate be at a future date?" Liquidity from this market can then be used to offer EtherUSD tokens which are pegged to the USD value. These tokens would be sold at a small

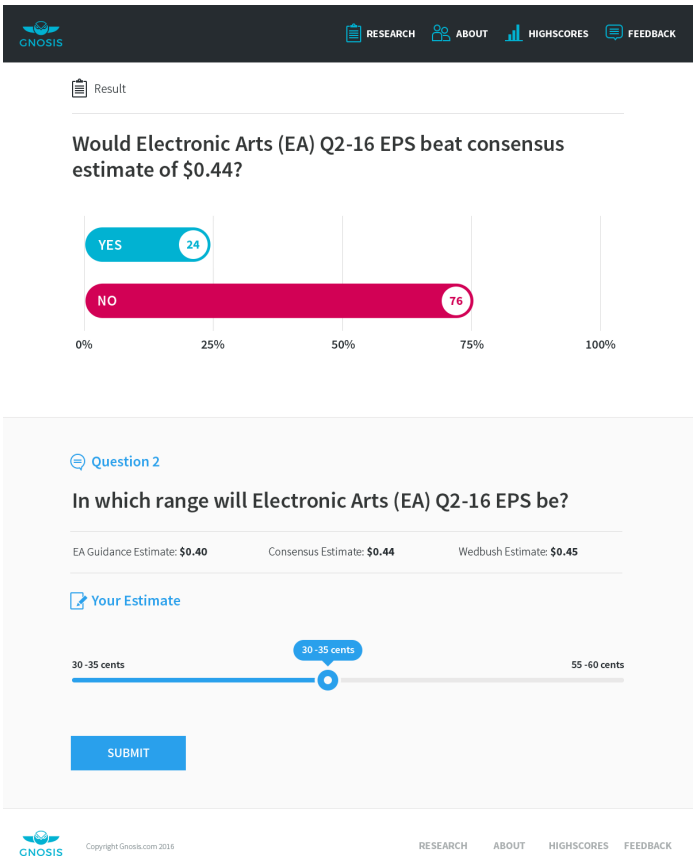

Figure 4.1: Financial Instruments: Predicting EPS

markup dependent on the cost of the market to provide liquidity. In another case, stable currencies (currencies designed to have stable value) can be constructed by taking a basket of positions in both sides of many markets.

While these synthetic instruments seem somewhat complex in nature, they illustrate the diverse applicability of prediction markets in finance and represent a tremendous economic opportunity for a diverse set of local experts in a truly global environment.

4.2 Insurance & Hedging Instruments

As stated previously, highly liquid prediction markets are remarkably accurate in assessing the likelihood of future events and therefore signaling associated risks. In the context of insurance, a prediction market could be used to estimate the likelihood of an insured event and may serve as an input to or even replace certain actuarial models. For example, a home insurance policy could create a market asking, “Will this area flood in the next year?” or “Will an earthquake over 5.0 magnitude occur within 50 miles of this location?” Probabilities drawn from these markets can effectively approximate more sophisticated actuarial estimates that require highly specialized (i.e. expensive) training. In our case, risk measurement becomes democratized, creating new economic opportunities for any participants with valuable localized knowledge.

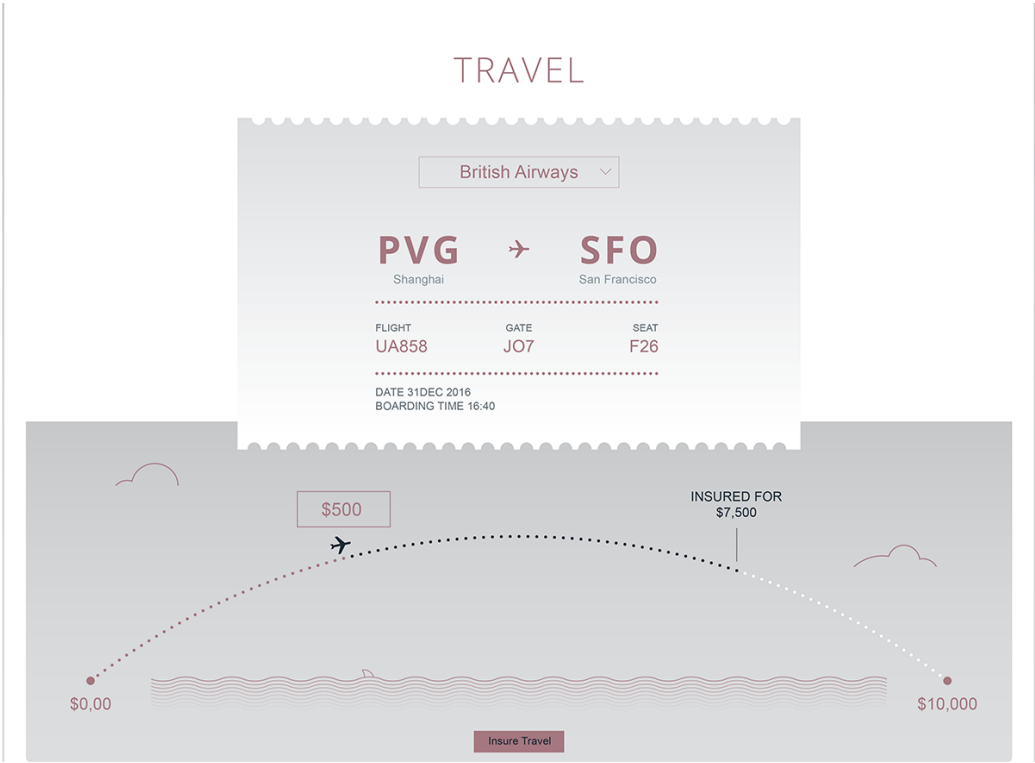

Figure 4.2: Hedging risks: Travel insurance

Furthermore, an insurance policy can be automated by creating smart contracts which buy "Yes" positions in each market affecting policy risk. In theory, shares in "Yes" positions should pay out an amount equal to or greater than the projected damage if an event occurs. The cost to the consumer for these contracts should be equal to the cost of all shares for their policy, plus some markup so that the insurance decentralized autonomous organization (DAO) may profit. The Gnosis Oracle Market serves as a hub for claim inspectors to offer their inspection services to settle outstanding claims. This insurance method should result in lower costs to the consumer as it automates business and accounting functions and crowdsources the actuarial process, as well as claims inspection.

4.3 Information

Prediction market applications for information sales can be broken into two major categories: *sales of insider information*, and *sales of device data*.

The sale of insider information has incredibly high utility but is ethically and legally questionable. For example, a market could be created asking, "Will corporation A acquire corporation B within set time frame?" In an anonymous prediction market, traders with insider information (e.g. high level executives in corporation A) could participate in this market for an almost guaranteed profit. The benefit of these types of markets is that they offer real-time access to all information relevant to a corporation's value, leading to more accurate pricing and strong-form efficient markets. The downside, however, is that such markets can sidestep existing SEC regulation that prevent insiders from individually profiting on material, nonpublic information. In the coming years, we will observe the interplay between new information markets and the incumbent financial system. We expect

regulators to continue to adapt its restrictive regime around insider information for prediction markets but with much difficulty given the decentralized nature of competing, blockchain-based platforms. What we may witness by way of response from the market is a stronger trend towards financial transparency by default, which arguably shrinks the pool of profitable insider information and introduces greater efficiency into global markets.

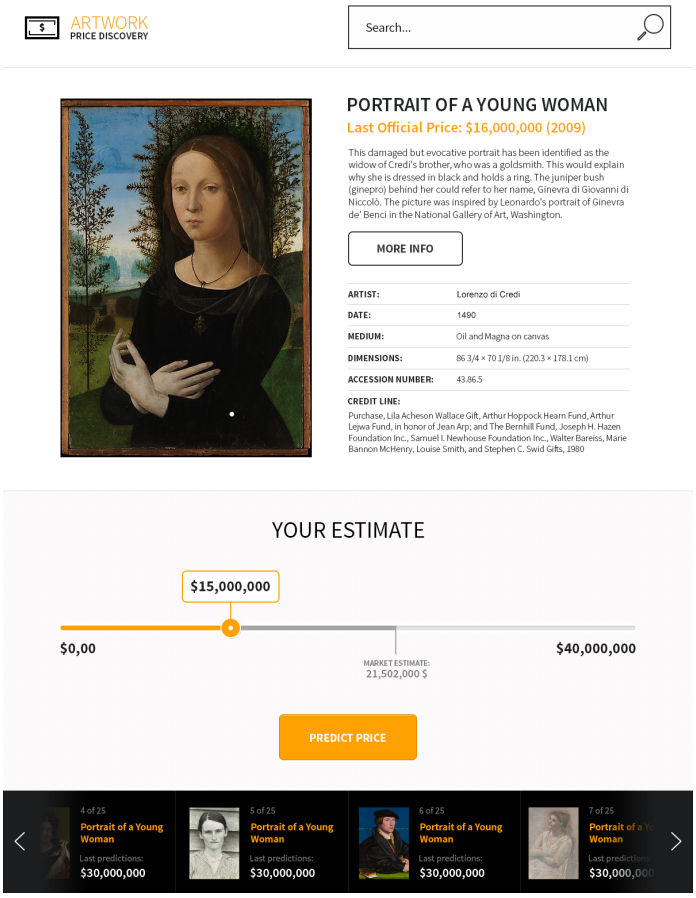

Figure 4.3: Information gathering: Gauging the price of a piece of art before the auction is vital for auction houses. Prediction markets can garner the expertise of specialists.

The latter example of device data sales may result in more efficient Internet of Things (IoT) optimizations, but it requires prediction market scaling tools (e.g. state channels) to achieve high frequency participation with low transaction cost. An example of this application could be smart cars participating in a market asking, “What is the traffic speed at this location?” Such a market resembles something like Waze, a consumer GPS navigation application, but for machine-to-machine interactions, providing a profit opportunity for sensor devices reporting probability estimates used to directly optimize IoT operations. In the prior example, markets on traffic speed would be used for driverless car routing or by third party maps services.

4.4 Governance

Decision making is at the core of all governance models. Organizations must make decisions on which policies to implement in order to maximize future welfare. For a government, this could mean deciding how to budget annual tax revenues among competing policy initiatives. Take, for example, the following questions:

"Should a portion of the budget be allocated for infrastructure projects or for education? Which option will result in greater GDP?"

Within a corporation, many decisions must be made over the course of a fiscal year, with variations in frequency, time sensitivity, and value to the firm. Disputes can arise within a broad class of decisions at all levels of a company, many of which can be catastrophic if poorly executed, such as a decision about whether or not to acquire a competitor. In most cases of governance, such decisions are made using a hybrid of democratic and autocratic processes. The former involves a voting process in which members of an organization or government cast votes (allocated through an egalitarian or proportional representation) where a plurality, majority, or supermajority is required to implement a decision. The latter involves a hierarchical model in which designated individuals make absolute decisions over their domains of control. Metrics on share price or future revenue may be the deciding factors for evaluating the resulting performance at a later date. Both of these models suffer from information and coordination inefficiencies, often resulting in the implementation of policies and actions that poorly optimize organizational welfare.

Futarchy, coined by Robin Hanson of George Mason University, offers an alternative, market-based approach to governance. In futarchy, markets are used to decide on and implement policies. These markets follow a general form of:

"What will a future welfare metric be if a policy is implemented?"

For example, a corporation could ask, "What will our Q4 revenue be if we fire our CEO?" (*Market A*) and conversely, "What will our Q4 revenue be if we don't fire our CEO?" (*Market B*). We can refer to these bi-directional markets as *decision markets*. Speculators who believe they hold unique insights into the outcome of firing or keeping the CEO are incentivized to participate in both markets. If the speculator believes that revenue will be maximized by firing the CEO, then he/she will buy long shares in the company's expected revenue [E(r)] if the CEO is fired and short shares in the E(r) if the CEO is not fired. Upon market closure, a decision is made corresponding to the greater expected outcome. In our CEO example, if the market value for *Market A* E(r) is greater than *Market B* E(r), then the organization fires the CEO. Market participants are then rewarded depending on their accuracy in predicting future revenue.

In this model, governance is both marketized and automated. Policies are determined by values found on an open market and implemented either through bonded delegates or an automated process. Prediction markets have shown to be the most efficient information aggregation tool, lending credence to the hypothesis that futarchy can more accurately identify policies that will optimize outcomes while also lowering bureaucratic overhead.

4.5 Incentivization

Apart from estimating event probabilities, prediction markets can be used to incentivize actions in the real world. An agent wishing to incentivize an action creates a market asking "Has a particular event occurred?" or "Will this event occur at a particular date?" Following the creation of these markets, the creator aggressively buys shares in the "No" outcome. Assume a participant observes

the prediction market and believes he/she can implement the action described. By buying "Yes" shares at low cost in the market, this participant is effectively guaranteeing payment if they can complete the action, resolving the market in their favor.

An example of this applied mechanism is a market for zero-day software security exploits. In this case, a company could ask, "Will a zero-day exploit be discovered in our software?" A penetration tester would attempt to discover exploits in the software. If an exploit is found, the tester would buy shares in the market for "Yes, an exploit will be discovered," and then reveal the exploit, either directly to the company by openly releasing the exploit or by using it in practice. For this particular application, the company would likely benefit from creating a secondary smart contract which pays the tester an additional amount for revealing the exploit privately to them.

This mechanism of incentivizing actions extends similarly to any event whose outcome can be heavily influenced by market participants. A use case which we find particularly interesting is the automated political action committee (PAC). A PAC raises money for the purpose of influencing an election or legislation. Through a combination of prediction markets and smart contracts, we can automate this process. First, a smart contract is created which is designated to impact a particular political outcome. Actors may donate money to this contract.

Following this, a market is created asking, "Will this legislation pass?" The PAC smart contract (smartPAC) spends its funds to buy shares in "No" (this will not pass). This effectively creates an incentive for actors to buy shares in "Yes" with the intent to impact the legislation positively and profit. If the legislation passes, then the smartPAC has achieved its purpose. If the legislation does not pass, then the contract profits from its positions in the market. This profit is then retained to be used for a future market created the next time similar legislation is proposed. Funds are rolled over similarly until the legislation passes. Effectively, this creates an ever increasing incentive to pass the legislation.

4.6 Sports Betting

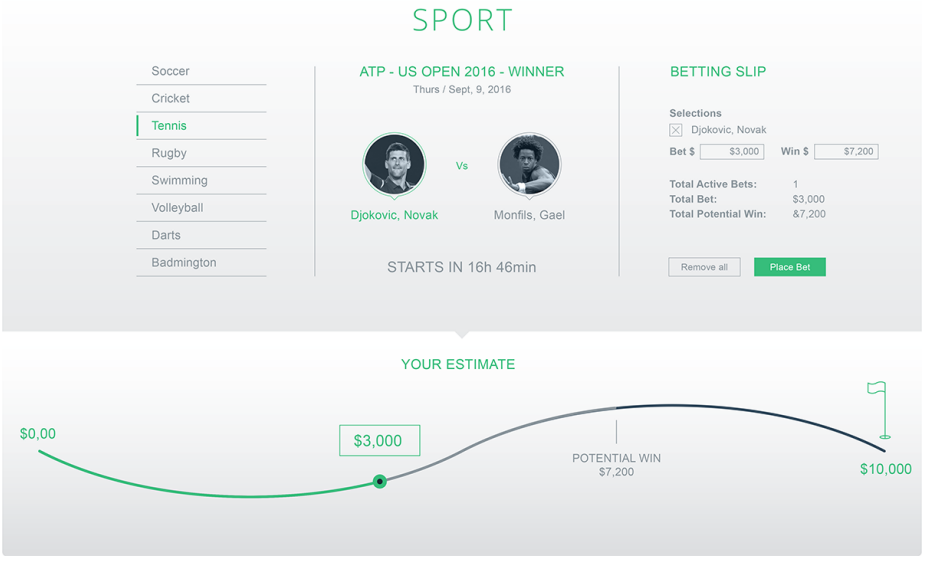

Figure 4.4: Using prediction markets for sports betting

Global online sports gaming is a massive market, with at least high 100s of billions wagered through regulated markets, and some estimate shows that 10 times that volume goes through unregulated markets. Companies such as William Hill make over 1B in revenue yearly. Despite this massive opportunity, companies and governments have been slow to innovate on existing models. Existing applications operate on siloed data and liquidity pools, have limited accessibility, and are slow to bring new products to market. Additionally, with centralized services, the user incurs additional risk such as theft or other failure, and unexpected issues with payment processors.

Two major impediments contribute to this situation. The first is that becoming a new entrant to the gaming market is a costly endeavor, with startup costs for a new company being at least in the several millions. Another major roadblock is the aforementioned siloing and accessibility issues. Without open, equitable, and transparent access to markets, products can not offer a platform model promoting innovation, and odds suffer.

Gnosis brings innovations that aim to solve these problems. Gnosis operates as an open platform, where access is unbiased and transparent. Thanks to this, incumbents and new participants can safely reap the rewards of operating on the same platform, such as added liquidity translating to better odds. Gnosis can be accessed anywhere and provides the same markets to all parties. Additionally, our platform provides the majority of back-end logic necessary for a new application, lowering the barrier of entry for new entrants by at least 2 orders of magnitude.

5. Roadmap

5.1 Key Activities & Partnerships

- Gnosis has been live on the Ethereum Mainnet for over 18 months. We updated to the next major software release in August. Gnosis was described by Vitalik as most advanced DApp currently live on main net in early 2016 interview
- Substantial work has been performed toward state channels and scaling. By implementing state channels into our software architecture, Gnosis is pioneering one of the most important tools for scaling DApps on the blockchain
- Ramping up community engagement efforts. Massive growth in social media following. Surpassed 1000 slack members and 2500 Twitter followers. Gnosis is actively pursuing high profile partnerships within Ethereum ecosystem, financial sector, and other fields.

5.2 Competitive Analysis

Gnosis competes directly with blockchain-based prediction market platform, Augur. Table 5.1 highlights the differences between the two platforms at a high level. The following section seeks to expound the differences between the two offerings in greater detail. It is also worth noting that Gnosis competes indirectly with sports betting exchanges and option markets such as William Hill and Nadex, respectively.

Gnosis differentiates itself from Augur in a few important ways. First and foremost, Gnosis is positioned as a platform for building decentralized prediction market applications as opposed to simply being a hub for market-making activity. The distinction is subtle but represents a wholly different strategic approach to capturing the burgeoning predictive asset market; the value proposition of the former centers itself on the production of actionable information, while the latter centers itself on generating data. Although Gnosis invests heavily in the development of its trading engine, crucial to Gnosis' mission and thesis is the transformation of this large set of prediction market data points into meaningful information that can be utilized by human and AI decision-making agents. In order to make this a reality, Gnosis plans on focusing a significant amount of resources towards enabling the creation of decentralized prediction market applications that make it simpler and easier for users

to gain insights into complex topics and make more well-informed decisions.

The platforms also differ in their mechanisms for resolving market outcomes. Augur's design is fully decentralized and requires *Reputation* (REP) token holders to allocate a certain portion of time per month for reporting event outcomes. The result is a resolution process that is theoretically more difficult to game but suffers from slow payouts. Taking this cost as a significant point of friction that may hamper the speed of prediction market adoption, Gnosis made an important design decision to create an oracle solution that is centralized by default and features a truly decentralized failsafe. This mechanism enables quicker resolutions via specialized, reputation-backed oracle services offered through Gnosis' oracle marketplace without sacrificing the benefits of decentralization.

Comparison of Gnosis and Augur Platforms		
Feature	**Gnosis**	**Augur**
Decentralized oracle	✓	✓
All funds held by contracts	✓	✓
Markets resolve quickly	✓	×
Active token holder requirements	×	✓
Scalable	✓	×
Application ecosystem	✓	×
Market-based governance protocol research	✓	×
Cross-compatibility as a standard	✓	×

Figure 5.1: Comparison between Gnosis and Augur

The trajectories of each platform also differ, at least in part, when considering the ongoing research that is being conducted around long-term scalability and market-based governance. With respect to scalability, blockchain-based prediction markets that aim to serve data-hungry IoT applications will suffer from prohibitive transaction costs due to the high frequency of requests. While protocol-level scaling solutions (i.e. sharding) are still in development, Gnosis has committed resources towards investigating off-chain scaling solutions such as *state channels*, which enable verifiable, on-chain settlement of large volume transactional events that is cost-effective. Gnosis identifies its active leadership in state channel research as a significant competitive advantage in the decentralized prediction market space precisely because it enables decentralization at scale. Any first mover advantage captured by Augur in the Ethereum space will be quickly outpaced once platform scaling becomes the primary concern. In addition, Gnosis was awarded a grant by the Ethereum Foundation to conduct research on market-based governance, known as *futarchy* (detailed in *Ongoing Research* below). Selection by leading scientists in the blockchain field further signals the overall strength of the Gnosis leadership team.

5.3 Finances

We are raising approximately $12.5M USD denominated in ETH in our token launch. Fund usage will be split approximately evenly between platform and application development.

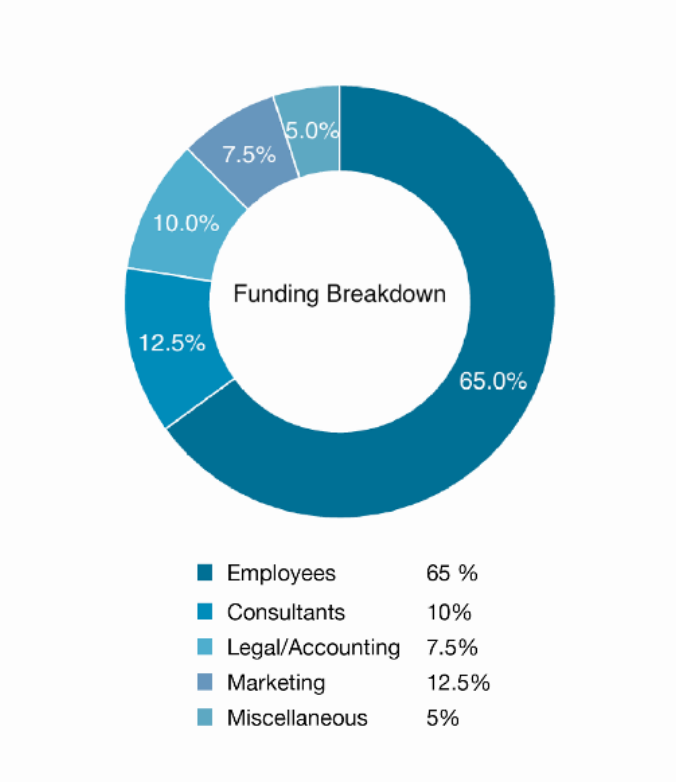

Figure 5.2: Funding Breakdown: Use of sale proceeds

5.3.1 Use of Sale Proceeds

Platform Development

Platform development will include building upon and securing core smart contracts, additional frameworks such as a comprehensive oracle market, trading and management interfaces, service level app templates and customization tools, and integrations with future Ethereum infrastructure such as state channels and stablecoins.

Legal Costs

Legal requirements include corporate setups in at least 3 locations for crowdsale, operations, and gaming licenses. Work has been done prior to token launch with a US law firm to develop a legal opinion of the interpretation with US law. Ongoing resources will be required for gaming and possible financial use case legal work. A legal contingency fund will be reserved in case of future issues.

Marketing and Business Development

Business development efforts will be focused on identifying and forming relationships with new projects and existing partnerships which can be built on Gnosis. Marketing will be focused on marketing Gnosis applications to their potential customer segment. Additional efforts will be spent on increasing awareness and knowledge of the Gnosis platform and what can be built with it.

5.4 Development Roadmap

5.4.1 Current State

Gnosis has been under development for over two years starting at the beginning of January 2015. Since then there were multiple iterations over the core parts of Gnosis. This includes the smart contracts powering the framework as well as the general web interface and the gnosis.js library. Several integrations of Gnosis with other projects like uPort, MetaMask or RealityKeys have been tested. In addition to developing Gnosis core infrastructure Gnosis has been developing other products like the Gnosis multisig wallet (`https://wallet.gnosis.pm`), which is used by several projects (Golem, Weifund, Stable, ...) setting a standard for secure fund management. Gnosis is also involved in development of developer tools and developed the first version of a program, which turned into the widely used TestRPC.

5.4.2 Future Development

Throughout the preparation of the token auction, development of core Gnosis technology continued making sure that next iterations can be published in time.

2017 Q2:

- **Additional office in Berlin**
  Parts of the team will join the vibrant Ethereum community in Berlin, working out of the ETHDEV office.
- **gnosis.js beta & smart contracts beta**
  gnosis.js is an easy-to-use tool to enable every website developer to build applications on top of Gnosis. gnosis.js will integrate with redesigned, more modular smart contracts and GnosisDB as only backend.
- **Gnosis management interface beta**
  A new interface to create and maintain markets is in the making, giving Gnosis a modern look and intuitive and productive user experience. The interface will offer a light version providing all functionalities needed to get started as well as an advanced interface to get the most out of the functionalities offered by Gnosis.
- **GnosisDB alpha**
  Storage of data on a blockchain is very expensive. The blockchain should only be used to verify the accuracy of information. GnosisDB is a generic database layer combining cheap storage with the advantage of fast document search and retrieval.

2017 Q3:

- **Gnosis Hackathon**
  Gnosis will host a hackathon to support developers building apps on top of Gnosis.
- **RealityKeys oracle integration**
  RealityKeys will be the first external oracle provider integrated into Gnosis. The RealityKeys

service provides automated and human-verified data designed to enable a new generation of automated information services.

- **CryptoEconomic experiments**
  Thanks to a generous grant from the Ethereum foundation, Gnosis will be running a series of experiments to test the viability of Futarchy.
- **Oraclize integration**
  Oraclize will be the second external oracle provider to be integrated into Gnosis. Oraclize aims to be the privileged data gateway between blockchain protocols and the world wide web. Oraclize's main goal is to provide a way for smart contracts to break free of their constraints and provide them with the ability to access all the data they need from the web without compromising their trustless nature.

2017 Q4:

- **Gnosis AMA DApp with Twitter integration**
  The first end user facing application will target the blockchain audience. Gnosis AMA (ask me anything) DApp allows anyone to ask the market for an estimate: "When will Ethereum switch to Proof of Stake", "Will Bitcoin hardfork?". A Twitter bot integration will make it possible to participate by tweeting.
- **GNOWIZ functionality**
  Together with the first end user facing application we will enable the WIZ functionality for paying fees.
- **On-chain exchange as price oracle**
  As part of offering WIZ functionality, a newly designed market maker will be deployed. This market maker can serve as a universal price oracle for tokens and will enable using WIZ to pay fees for markets traded in any ERC20 token.

2018 Q1:

- **GNODEX alpha**
  GNODEX will be the decentralized, scalable state channel based exchange to trade predictive assets. The development already started, however, due to its complexity it is expected that an MVP will take at least a year of development. Gnosis seeks to combine forces with other teams to build the basic layers required to offer a decentralized exchange, with Gnosis adding layers required to offer trading of prediction market outcomes.
- **Futarchy DAO alpha**
  Gnosis will build a tool to create DAOs that only use prediction markets to make decisions. The concept was popularized under the name Futarchy by our advisor Robin Hanson. Futarchy DAOs will measure the market reaction of a decision. The DAO will only make decisions if the market would react friendly to it.
- **Gnosis trading interface alpha**
  The generic Gnosis interface does not meet the requirements for a trading application. For this reason Gnosis is working on an additional interface dedicated to traders. Gnosis will combine the best parts of existent trading interfaces from Kraken, Poloniex and other trading platforms to provide a seamless user experience.

2018 Q2:

- **Gnosis desktop standalone client alpha**
  Gnosis will offer a standalone application based on Electron, which will come with IPFS

and an Ethereum light client integration. It will allow a fully decentralized prediction market service ensuring that Gnosis prediction markets can be used from anywhere in the world.

- **One other Gnosis DApp**
  It is the core concept of Gnosis to enable third parties to develop prediction market applications on top of Gnosis. To accelerate platform adoption, the Gnosis team will develop DApps directly as well. Which application areas we choose to work on will be based on our own business development research and coordinated with outside teams that decide to build on Gnosis.

2018 Q3:

- **Gnosis mobile standalone client alpha**
  As mobile becomes the dominating platform for application users, Gnosis seeks to have dedicated applications for Android and iOS. Based on the popular Status.IM client, Gnosis will provide fully decentralized prediction markets on mobile.
- **Gnosis DApp Store**
  Gnosis will integrate a DApp store into the standalone apps for desktop and mobile to promote apps.

5.5 Ongoing Research

Thanks to a generous grant from the Ethereum foundation, Gnosis will be running a series of experiments to test the viability of Futarchy. Over the next two months we will be running at least three experiments that will test foundational assumptions toward the successful implementation of Futarchy. The first two experiments will test the ability of actors to manipulate the outcome of markets when incentive is provided. In the first of these experiments, a market will be created which is resolved by a smart contract verified to output 5 at a particular time. In the second experiment, the smart contract will resolve to either 0 or 10 dependent on a random seed. In each of these markets, an incentive will be provided to manipulators if they are able to push market values away from their expected outcomes. By providing this reward we simulate market participants who have incentive (likely by being shareholders in an organization) to misinform a Futarchy market and will be able to provide quantitative information regarding the viability of such schemes. In the third market, we will simulate Futarchy more directly with markets determined by Ethereum blockchain difficulty at a time in the future. Through these experiments and general tools which we plan to create on Gnosis, we aim to forge a solid platform for DAOs (and other types of organizations) to use Futarchy to inform and automate their decision making.

6. Token Auction

Platform, or "app" tokens in decentralized networks can be distributed in a variety of ways. In the Bitcoin and altcoin model these tokens are distributed gradually via a proof of work or proof of stake mechanism. There are variations within these mechanisms which are primarily dependent on the supply curve. This curve determines the rate of dispersion of the total coin supply.

Over the last two years, an informal standard for the token purchase mechanism has emerged. In this mechanism, cryptocurrency is sent to an address in exchange for some amount of app tokens. These launches typically run for approximately a thirty day period. The rate of conversion between the sent cryptocurrency and the received app tokens decreases over this thirty day period (i.e. fewer tokens per unit of cryptocurrency). For example, the rate during the first two weeks may remain at 1 ETH for 200 Tokens and then decrease linearly to 1.5 ETH for 200 Tokens and below over the open period. The amount of tokens created is often wholly dependent on the amount of cryptocurrency sent in. If the rate is 1 ETH for 200 tokens and 1M ETH is sent, 200M tokens will be created.

This model causes friction for purchasers in that they are compelled to purchase app tokens earlier than they feel is justified for fear of price increases. The uncertainty in token creation may also potentially lessen the utility of the app tokens themselves.

Our Innovations

In the Gnosis launch, 10 million Gnosis tokens (GNO) will be created and a percentage of them will be distributed through the token launch. The launch period will end when either of the following ending criterion is met: approximately $12.5M USD denominated in ETH worth of GNO is sold, or 9 million GNO tokens are sold.

During the launch, participants will be able to send Ether to a token launch address, committing to buy GNO at or below the current price at the time of their purchase. The price of GNO will be determined by a falling, as compared to the current trend of rising, price specification. The price of GNO will decrease every block that elapses during the launch. The price per GNO sold in the final block, when either ending criterion is satisfied, is the price that will be applied to all preceding sales during the launch period. Therefore, token launch participants are committing to a maximum price per Gnosis token and will receive tokens at this rate or lower.

Example user experience

For example, Gnosis creates 10M tokens and begins the token launch. Alice sends 1 ETH to the token launch address while the rate is at 1 GNO for 1 ETH. The token launch process continues, with the price per GNO lowering each block. approximately $12.5M USD denominated in ETH worth of GNO is sold on the 7th day of the launch. On the ending block, tokens were sold at a rate of 1 GNO for 0.5 ETH. The token launch concludes, and every participant gets tokens equal to the amount of Ether that they sent, at the rate of 1 GNO for 0.5 ETH (the price when the ending criterion triggered). Alice, who contributed 1 ETH on day one, would therefore receive 2 GNO, applying the final sale price to her 1 ETH purchase. Again, participants declare the maximum price they are willing to pay for GNO, but ultimately receive the lowest price that any purchaser pays for GNO as the final sale price is applied to all purchasers.

Participation

Economic theory dictates that purchasers should participate only when GNO tokens reach a price they feel is representative of their utility in the Gnosis ecosystem. Tokens will sell for a very high price on the initial blocks, representing high demand for the product. If a buyer believes the price is fair, economic theory would encourage participation. If not, economic theory dictates buyers should wait until GNO reaches a price they feel is warranted by GNO's functionality in the platform. It is important to emphasize that our token launch model is significantly different from previous models. Participating early provides no special benefits.

Origin of this mechanism

Alex van de Sande, seeking out a model that better enabled participation and eliminated some of the aforementioned friction created by existing mechanisms, originally proposed this auction mechanism. We agreed with his suggestion and following the initial idea, we worked with Alex and Vitalik Buterin, who also supported this design, to create the mechanism that we have now settled on.

Code

Stefan George, technical co-founder of Gnosis has released our smart contract source code along with the code for our token launch mechanisms. Additionally, his multisig wallet, recently used by Golem, is also available for review.

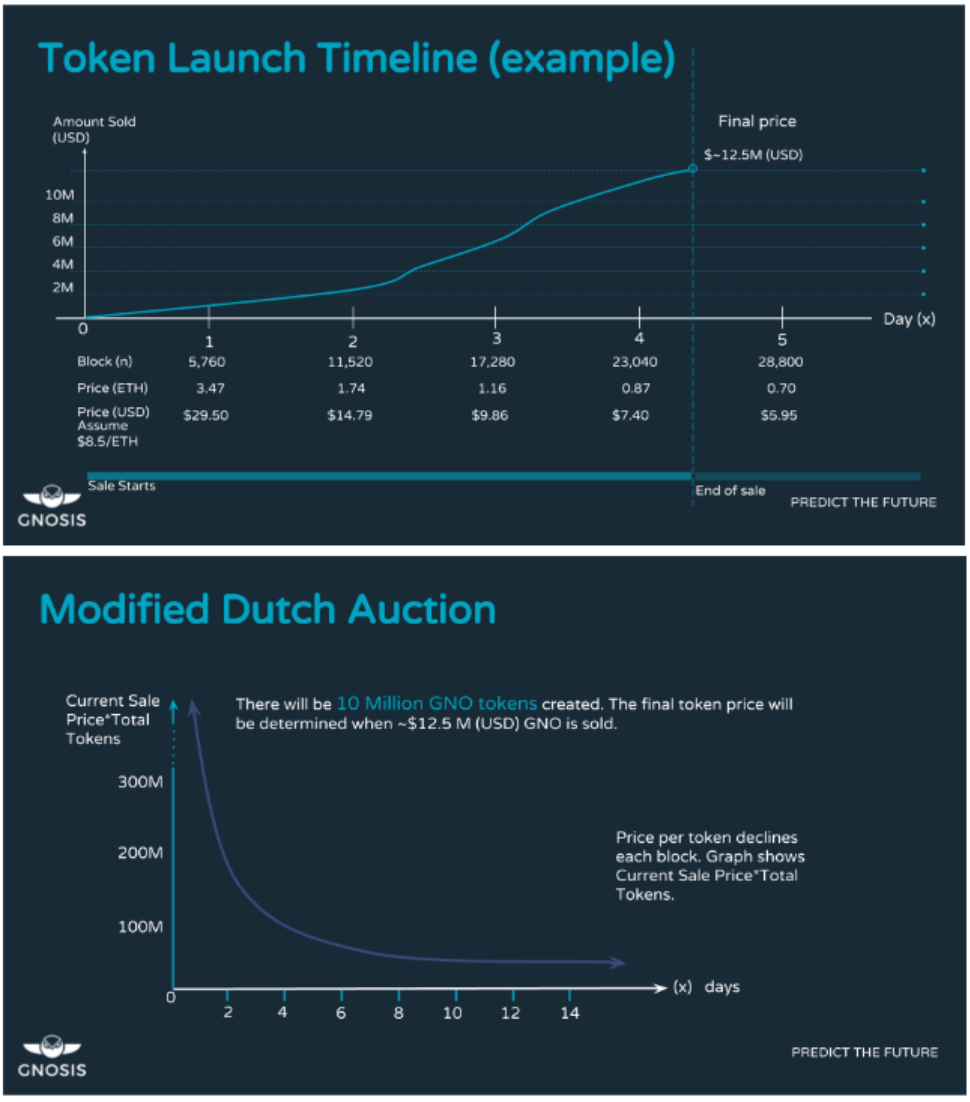

Figure 6.1: The Gnosis Token Auction

7. Leadership

7.1 Core Team

Martin Koppelmann, CEO

Martin Köppelmann has been an entrepreneur and thought leader in the blockchain space for more than 3 years. He is the founder of the biggest Bitcoin prediction market, Fairlay, but has shifted his focus to Ethereum and co-founded the decentralized Gnosis prediction market - the first bigger DAPP that went live on Ethereum. Closely related to prediction markets is his work on decentralized market driven governance mechanisms: Futarchy. Beyond the entrepreneurial activity Martin has done research on the economic incentive structure of different consensus mechanisms and scalability solutions via state channels. Martin co-hosts the Ethereum meetup groups in the Silicon Valley and San Francisco. Finally, Martin is well known for his work and research on "basic income on the blockchain: Circles" - a new currency built on top of Ethereum that aims to implement a basic income as monetary policy.

Stefan George, CTO

Stefan is an entrepreneur and developer who became interested in Bitcoin in 2013. He started his own Bitcoin startup Fairlay, one of the leading prediction markets using Bitcoin today. Previously Stefan worked at tech companies in Silicon Valley and at Berlin-based startups. After finishing his Master's in CS he decided to travel Asia for a year in 2014 and started Gnosis afterwards working from Berlin. The first alpha version of Gnosis was released just one week after the launch of Ethereum. Stefan leads the development at Gnosis and implemented the smart contracts behind the prediction market platform. In addition, Stefan wrote the crowdfunding campaigns for Gnosis and SingularDTV and the multisig-wallet used by Gnosis and Golem to store their funds and tokens.

Matt Liston, Strategist

Matt dove into the cryptocurrency space in 2013 after becoming fascinated by the potential for smart contracts and blockchains to revolutionize information economies. Currently he works on strategy and communications for Gnosis. Prior to joining ConsenSys Matt founded Augur, consulted for EthDev, and designed an IoT blockchain solution for Enlighted. When not DAO-whispering he

composes algorithmic music and dreams of a future Skynet-on-a-blockchain.

Dr. Friederike Ernst, COO

Friederike is a physicist by training and after obtaining her PhD from the Free University of Berlin worked at Columbia University and Stanford/ SLAC before becoming a guest professor at the University of Hamburg. Friederike has been interested in cryptotechnology for many years and has switched gears from basic science: She now structures, organizes and directs company operations.

Denis Granha, Full stack engineer

Denis is a software developer who started working with Cryptocurrencies at the end of 2015. He found in Ethereum the potential of growth and freedom where creativity it's a value. In Gnosis Denis works both on FrontEnd and BackEnd development. Before joining Gnosis, Denis worked as software consultant in projects for important companies at Spain. Besides his technology passion, Denis is a musician and a singer who loves to play live shows with his band.

Giacomo Licari, Full stack engineer

Giacomo is an enthusiastic software developer. He got in touch with cryptocurrencies and blockchain in late 2016. Prior to Gnosis, he worked as software developer for an international consulting firm and for a company developing ERP services. He's currently working on Gnosis services along with the Development Team.

Alan Lu, Mathematician, cryptographer

Alan was born in Shanghai, but grew up in Dallas, Texas, where he really got into technical stuff. He studied math at UT Dallas and worked as a programmer. He enjoys gaming, training, and pursuing his interests in physics. You may find him hanging out at hackerspaces, reading at a bookstore, or on an adventure somewhere. At Gnosis Alan finds suitable models and algorithm implementations and works on the distributed key generation used by GNODEX.

Rami Khalil, Information security

Rami developed an interest in decentralized currencies and infrastructure recently in 2016. While currently pursuing a Master's in Information Security at the Swiss Federal Institute of Technology in Zurich, Rami researches payment channels as part of his studies and designs and implements core Gnosis functionality. Previously, he interned in Silicon Valley, and was an ACM ICPC World Finalist.

7.2 Board & Advisors

Joseph Lubin (Board member)

Co-founder of Ethereum and founder of ConsenSys. An academic background in Electrical Engineering and Computer Science from Princeton University and research experience in the field of Robotics Learning. Former VP of Technology at Goldman Sachs in the Private Wealth Management Division.

Jeremy Millar (Board member)

Chief of staff at ConsenSys. As Chief of Staff, Jeremy oversees many of the Enterprise activities and strategic initiatives of the firm. Previously, Jeremy Millar was founder and managing partner of Ledger Partners. Ledger Partners developed out of Jeremy's increasing focus and passion for the blockchain and bitcoin ecosystem. This began with what was supposed to be a blog post that

became 'arguably the most comprehensive report to date on what is happening in the world of bitcoin and blockchain startups' which you can see here: `http://bit.ly/1Zq2Pvy`. Jeremy began his career as one of the first Java architects at Oracle, before moving into sales management and strategy roles, both within Oracle and at a number of start-ups. He went on to complete his MBA at Oxford University before joining the M&A team at Goldman Sachs. Jeremy was a founding partner at Magister Advisors, advising fintech and SaaS companies across Europe. He is also an active angel investor and mentor with the Barclays Accelerator powered by Techstars.

James Slazas

20 years of capital markets experience, initially on the futures' exchanges of the CME and La Matif. Managed a proprietary derivative arbitrage and structured products book for Lehman Brothers. Also, held $1B in emerging market credit risk for Lehman's London, Swiss and Hong Kong banks for HNW clients. James managed a life settlement hedge fund uniquely acquiring longevity risk for limited partnership units.

Robin Hanson

Robin Hanson is an associate professor of economics at George Mason University and a research associate at the Future of Humanity Institute of Oxford University. He is known as an expert on idea futures and markets, and he was involved in the creation of the Foresight Institute's Foresight Exchange and DARPA's FutureMAP project. He invented market scoring rules like LMSR (Logarithmic Market Scoring Rule) used by prediction markets such as Gnosis, and has conducted research on signaling.

Jason Trost

Founder and CEO of Smarkets (>$1.5B betting exchange). Prior to founding Smarkets, Jason was an application developer at UBS's Global Asset Management (New York) where he focused on innovative web technologies. Jason founded internet startup Descipher, a consumer medical website and has also been an equities trader at Great Point Capital (Chicago).

Vitalik Buterin

Founder of Ethereum, Ethereum Chief Scientist. Vitalik Buterin is a Russian born programmer and writer primarily known as a co-founder of Ethereum and as a co-founder of Bitcoin Magazine. Vitalik helped to develop Gnosis auction mechanism and is involved in the crypto-economic experiments conducted by Gnosis.

8. Legal Considerations

Due to our aspirations for what Gnosis may one day become, the Gnosis team exercised extreme legal diligence in the lead-up to our launch. This diligence includes significant expenditures on several law firms around the globe to evaluate the implications of our structure, token launch, and operations. In the United States, we've worked closely with Perkins Coie. In our home jurisdiction of Gibraltar, we've worked closely with Isolas. Due to the retrospective nature of regulatory action, the Gnosis team can make no guarantees regarding the legality of the platform or launch in any given jurisdiction. Regardless, we are confident in, and proud of, the work we've done to shape Gnosis into what we hope is a model of regulatory compliance for decentralized applications and token launches. We will be responsive and collaborative with any regulators as necessary going forward.

8.1 Legal Implications of Token Launches

GNO tokens are functional utility tokens within the Gnosis platform. GNO tokens are not securities. GNO tokens are non-refundable. GNO tokens are not for speculative investment. No promises of future performance or value are or will be made with respect to GNO, including no promise of inherent value, no promise of continuing payments, and no guarantee that GNO will hold any particular value. GNO tokens are not participation in the Company and GNO tokens hold no rights in said company. GNO tokens are sold as a functional good and all proceeds received by Company may be spent freely by Company absent any conditions. GNO tokens are intended for experts in dealing with cryptographic tokens and blockchain-based software systems.

8.2 Legal Landscape for Prediction Markets

As discussed herein, prediction markets are an area of interest for many regulators around the globe, including those within the United States. Though we feel decentralization holds great promise, we must, and intend to, operate our business in accordance with the laws of relevant jurisdictions. As such, Gnosis may not be immediately available in certain jurisdictions. The Gnosis team and our advisors are aggressively pursuing strategies to bring the benefits of Gnosis and the information

sharing economy to the globe as quickly as possible. First steps may include obtaining financial or gaming licenses as required by law.

9. Gnosis Architecture

Gnosis is still under heavy development and will be for a while. Parts in this section are subject to change.

9.1 Systems Architecture

Gnosis aims to be a fully decentralized, serverless prediction market application framework. Gnosis achieves this goal by utilizing existing technologies like Ethereum and IPFS advancing their features with new solutions like state-channels for scalable applications.

9.1.1 Core Components

The Gnosis core components consist mainly of four technologies, which are combined in a easy-to-use javascript library gnosis.js. Used technologies include:

1. **Ethereum** allows to run decentralized code with smart contracts, making transfer and settlement of assets simple and censorship resistant. Gnosis is using Ethereum for creation of events and markets. Every prediction market trade and every settlement is ultimately done on this layer.
2. **IPFS** allows to store static files in a distributed file system. It is using distributed hash tables to distribute files. Gnosis is using IPFS to store all static files gnosis.js or any UI element. In addition meta information of events is stored in IPFS.
3. **GnosisDB** tries to compensate the shortcomings of Ethereum and IPFS, which come without search capabilities. GnosisDB offers a distributed, scalable search combining on-chain with off-chain data and will be used to query event descriptions of prediction markets.
4. **GNODEX** is a decentralized exchange to trade predictive assets (prediction market outcome tokens) using state channels. State channels allow to scaleout transaction capacity of Ethereum, which is needed to make Gnosis prediction markets accessible to a large scale of users.

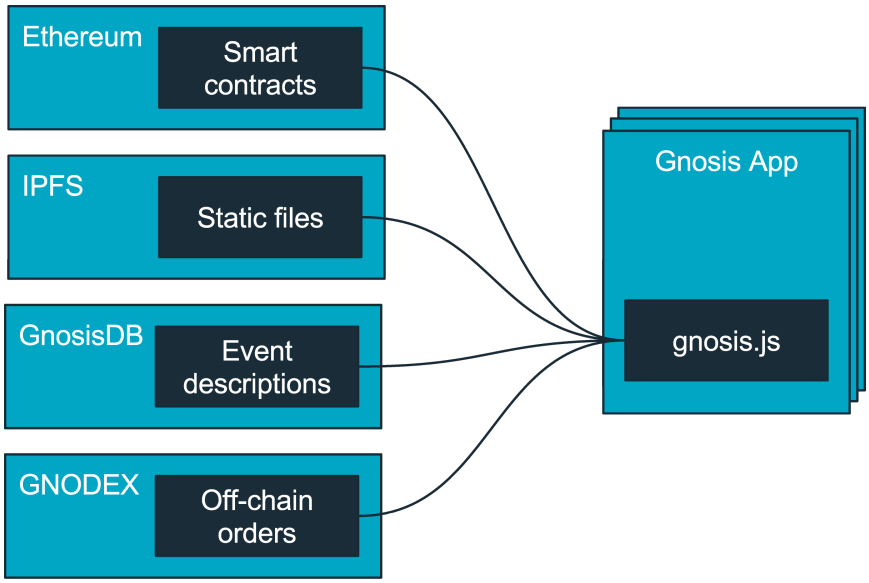

Figure 9.1: Gnosis architecture

9.2 Contract Architecture

The Gnosis smart contract design follows a very modular contract structure making it easy to split functionalities to upgrade or reuse parts. The current Gnosis implementation consists of over 25 smart contracts ranging from market makers to different oracle solutions. All of them have been written in Solidity.

For every prediction market two objects have to be created: An event object referencing a real world event and a market object, which connects the market maker with the event. Once the event occurred, the event can be resolved and winnings can be redeemed.

9.2.1 Event Factory

The event factory contract allows to create new event contracts, which can be used to resolve markets. Every event contract has the following (main) properties:

1. Oracle
2. Outcome tokens
3. Collateral token

Oracle

The event references an oracle contract, which will resolve the event. Gnosis is agnostic towards oracles and allows to use any contract as oracle, which is implementing the oracle interface.

Outcome Tokens

Events can resolve to a number in a range (ranged event) or an outcome out of a list of outcomes. An example for a ranged event is the Apple stock price on date X. An example for a non-ranged event would be the winner of the World Cup on date X, which has a list of teams as outcomes. For every outcome, an ERC20 compatible outcome token is created. Ranged events are represented with two outcome tokens for long and short positions. Non-ranged events have an outcome token for every defined outcome. Every outcome has a probability between 0% and 100%. All outcomes together

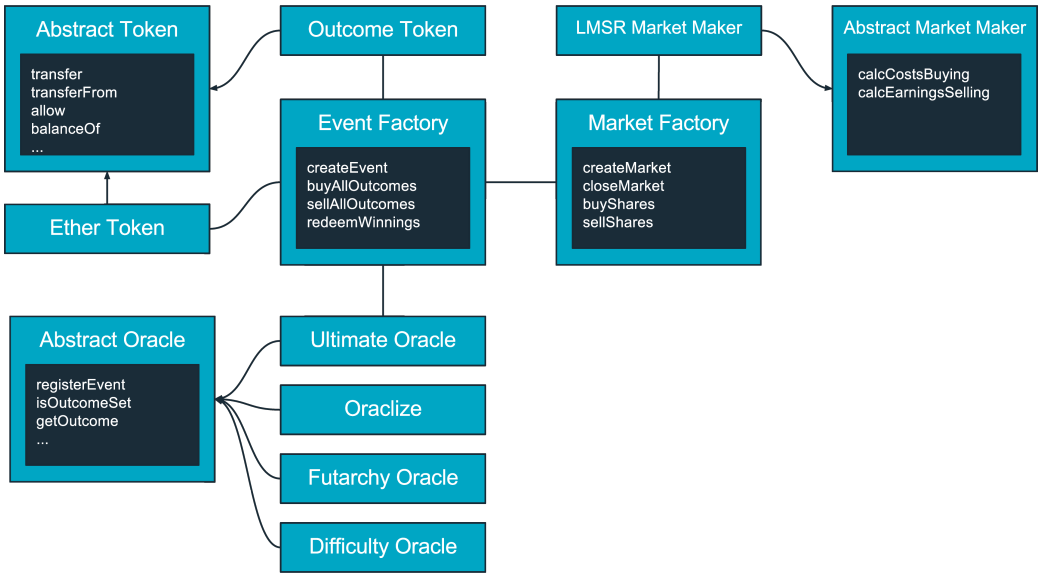

Figure 9.2: Contract architecture

add up to the probability of 100%. After an event is resolved the outcome is known. The winning outcome has now the probability of 100%. All other losing outcomes have the probability of 0%.

Collateral Token

The ERC20 compatible collateral token defines the currency in which the event is traded. It is possible to generate and redeem a full set of outcome tokens for 1 collateral token, because all outcomes together always have the same probability of 100%.

Assuming there is an event with two outcomes and a user invests 10 collateral tokens he will receive 10 outcome tokens for every outcome. The user decides to bet on outcome 1 and sells all outcome tokens for outcome 2. The price he can ask for outcome 2 tokens, depends on the market's estimate on outcome 2 will be the winning outcome. Assuming the market believes that outcome 2 will happen with 70% probability, 10 outcome 2 tokens can be sold for 7 collateral tokens. Hence the user invested 10 - 7 = 3 collateral tokens for 10 outcome 1 tokens. When the event occurred and outcome 1 was the winning outcome, the 10 outcome 1 tokens can be redeemed for 10 collateral tokens and the user earned a profit of 7 collateral tokens.

Common collateral tokens will be ether tokens or stable coins, which allow to trade without additional currency fluctuation risk.

Outcome Tokens as Collateral Token - Conditional Markets

Outcome tokens are ERC20 compatible and can be used as collateral tokens for other events. The interesting feature of outcome tokens is, that they only have a value when the outcome they represent occurred. Trading an event using an outcome token as collateral implies that the event is only relevant under the assumption that the outcome occurred. This allows to create events with conditional probabilities.

Assuming we want to predict, how the potential change of the Microsoft CEO affects the Microsoft stock price, we create two events:

1. Will Steve Ballmer be CEO of Microsoft end of 2014?

Outcomes: Yes, No
2. What is Microsoft stock price end of 2014?
Outcome: Any number

The first event can use Ether as collateral token but for the second market, we use the No outcome token representing the outcome "Steve Ballmer is **not** CEO of Microsoft end of 2014". Any market using the second event is predicting the stock price of Microsoft end 2014 under the assumption that Steve Ballmer is no longer CEO end 2017.

9.2.2 Market Factory

The market factory allows to create new market contracts, which allow to trade outcome tokens on markets with a market maker. Every market contract has the following (main) properties:

1. Event
2. Market maker
3. Market fee

Event

A market is always associated to an event and has to be funded in the event's denomination. The market maker buys and sells event outcome tokens via the event contract.

Market Maker

A market maker contract like the LMSR market maker allows trading of outcome tokens. The LMSR calculates prices for event outcome tokens based on demand using a logarithmic market scoring rule. The higher the demand the higher the price for the outcome token is. Assuming an event was resolved the demand for the winning outcome token will be so high that the price converges towards 1.

Market Fee

If a market maker is created the creator can define an optional fee - essentially a spread between bid and ask. This collected fee can compensate the market creator for the the initial funding he spent on the market maker to give the initial liquidity.

9.3 Oracle Architecture

While all other parts of a prediction market can be automatically executed by smart contracts on the blockchain, the somehow need to get the information about what happened in the real world to decide which outcome tokens have to be paid out. The Gnosis platform is oracle agnostic. Any contract can serve as an oracle so any oracle solution that will be development on Ethereum might be used by Gnosis. However, we do not only rely on external oracles, we developed a few our self. Different applications will have different oracle needs regarding resolution speed, decentralization or trust needs and costs. We anticipate 4 categories of oracles:

1. On chain oracles
2. Centralized oracles
3. Decentralized oracles
4. Hybrid oracles (centralized first + dec. as a backstop)

9.3.1 On-chain Oracles

We can directly provide an oracle about on-chain information. All that is needed is a simple interface contract that makes sure that the data is available in the right format. A few examples for this kind of data:

1. What will the difficulty be at block x?
2. What will be the price of a specific token on a decentralized exchange in one month?
3. Futarchy - will a DAO accept a specific proposal?

While the number of those use cases is limited today it will only grow. We can see Gnosis markets complementing almost every DApp project on Ethereum.

9.3.2 Centralized Oracle

For some applications it will just be fine to rely on a single data provider. In some cases there is just a single source of truth - e.g. a sensor measurement by a specific sensor, or game results from the NFL. If those data are available in a structured format only a signature of the issuer is necessary to make them smart contract compatible. We know that more and more institutions are looking into making their data smart contract compatible. However - in the meantime a trusted data signer might as well be for some applications the best solution. It is important to note that the signer (oracle) does not need to know about the prediction market or even Ethereum. Once the right outcome is signed anyone can submit this data to the blockchain and the smart contract will evaluate the signature. Centralised oracle providers that we have already integrated or are planning to include `https://www.realitykeys.com/`, `http://www.oraclize.it/` and `https://smartcontract.com/`.

9.3.3 Decentralized Oracle

Many proposal have been made for how to implement decentralized oracles. The common scheme is that the truth is a schelling point. Different decentralized oracle mechanisms usually set up various coordination games where participants are voting in an economically incentivised way. If they vote along with the majority they gain some value, if they vote against the majority they lose value. Projects implementing those schemes include Augur, Aeternity, Reality Token. All of these projects could (maybe slightly modified) be used on Gnosis. We tried to use the essence of those concepts and came up with a solution that does not have any operating costs. In contrast the Augur concept expects individuals to actively and regularly report on the outcome of events. Gnosis in contrast will use a concept called the **"Ultimate Oracle"** with the big advantage of not having any operation costs.

In the concept anyone can make a claim that a specific event had a specific outcome by putting money on the line. This claim can be disputed by anyone by putting more money on the line on an alternative outcome. The Ultimate Oracle will decide on an outcome if it continues to be the frontrunner (outcome with most money behind it) for a specific period of time (e.g. 24h). In this case all money that is put on losing outcomes will be distributed to those putting money on the winning outcome. To avoid that people will put money on the winning outcome shortly before it is confirmed the amount is capped at a proportion of the amount of money on the other outcomes. A concrete implementation can be found here: `https://github.com/ConsenSys/gnosis-contracts/blob/master/contracts/solidity/Oracles/UltimateOracle.sol`. A discussion about the game theoretical reasoning behind this concept can be found here.

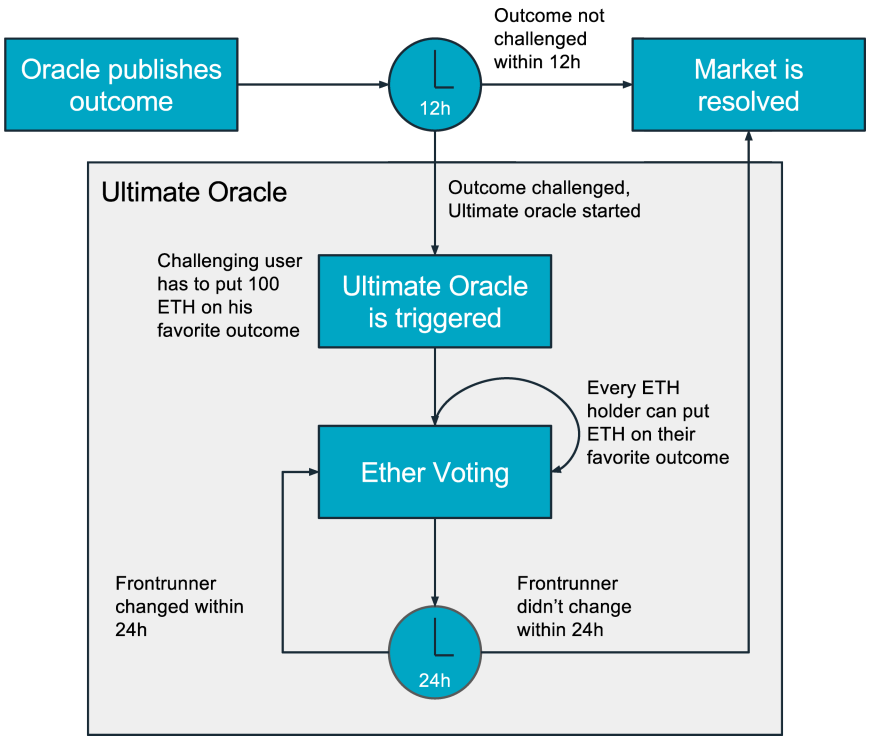

Figure 9.3: The "Ultimate Oracle"

9.3.4 Hybrid Oracles

In practice mainly combinations of those models will be used. For cost and resolution time reasons a centralized oracle will be used first. The oracle will publish a result as fast as possible (could be within seconds after an event occurs). This will trigger a dispute period. The length of the dispute period could be determined by the latest trading prices of the prediction market. If the suggested outcome traded close to 1 the market already agreed on this outcome and the dispute period can be very short. During this period anyone can trigger a dispute at some costs and give the responsibility for the outcome to a more decentralized or more secure oracle - e.g. a 3 out of 5 oracle. If the outcome of this oracle will still be disputed a fully decentralized oracle like the "Ultimate Oracle" can be used as a mechanism of last resort. This approach combines the low cost, high speed advantages of centralized oracles without compromising with the security that comes from a fully decentralized oracle which is used as a backstop. Find a implementation of this hybrid model here: `https://github.com/ConsenSys/gnosis-contracts/tree/master/contracts/solidity/Oracles`.

9.3.5 Oracle Standard for Event Descriptions

It does not make sense to store the full plain text event description on the blockchain. For this reasons we have in coordination with other oracle providers defined a standard format to store event information in a JSON file together with a deterministic process to generate an unique identifier for such an event description.

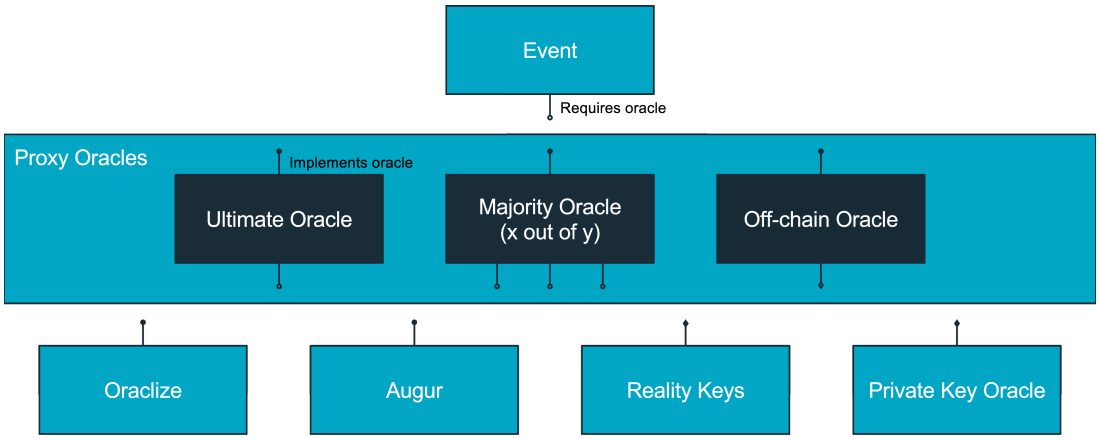

Figure 9.4: Hybrid Oracles

9.4 gnosis.js

gnosis.js is a free, easy-to-use javascript library, which can be used by anyone to create applications on top of Gnosis. gnosis.js combines the interfaces exposed by all four core components adding wrapping functions to allow for easy trading and creation of events, markets and oracles. An alpha version can be found here.

gnosis.js is used by the Gnosis administration tool offering a standard generic interface to manage prediction markets. An alpha version of the admin interface can be found here: `https://admin.gnosis.pm/`

9.5 GNODEX

9.5.1 State Channels

State channels allow transactions to occur without requiring them to be sent to the blockchain immediately. Transactions are exchanged in the form of signed messages which guarantee the transaction's future execution on the blockchain. The speed for such transactions is not limited by the block time, but by the speed of the channel used to exchange the signed messages. Additionally, the number of transactions are not limited by the block size or gas limit of the blockchain, but by the bandwidth of the exchange channel. In short, state channels allow an almost unlimited number of transactions with almost instant finality.

The difference between regular blockchain transactions and state channels is similar to the difference between a regular bank transaction and a bank check. Instead of sending a transaction via a bank and waiting for their confirmation, one can sign a check independently and exchange it with someone else. However, unlike the bank check, which might be not covered, state-channel transactions are cryptographically signed and can be verified by anyone.

9.5.2 Off-chain Order Books

The use of state channels is crucial to the development of an order book for a decentralized exchange, as the speed and volume of transactions in an exchange factors critically into its usability and cost effectiveness. An off-chain order book that operates with minimal trust and can be audited or

disputed by anyone through submitting cryptographically signed orders to the blockchain, if the order book misbehaves, is the goal of our implementation.

The Open Order Problem

Although cryptography renders the orders themselves tamperproof, it does not resolve a form of market manipulation here on referred to as the open order problem. Market makers stand to lose on markets if enough winning shares are bought from them, so they are incentivized not to include unfavorable orders come settlement. Similarly, market participants may manipulate the market by taking advantage of better network access to front run regular participants, degrading the user experience for those traders without the resources to conduct high frequency trading.

To solve the open order problem, we use blinded batch auctions. Traders submit their orders encrypted with a public encryption key, the orders are aggregated by the exchange, and a corresponding private key is published after the batch is resolved. The provider of the key pair will be a group of independent servers running a distributed key generation protocol. That way, nobody is able to manipulate the market as indicated unless they somehow corrupt the distributed key generation service past a threshold of servers.

It should be noted that although orders may be blinded by the use of hash submissions followed by the submission of the order after the batch period ends, that gives market participants the option of failing to submit the order if it turns out to be unfavorable to them. This is why we have opted for full encryption in this case.

9.6 GnosisDB

Ethereum promises censorship resistant, decentralized applications. However, as of today some parts of a DApp stack are still served from some central services, which are controlled by a third party and can be censored. One of those services is document search. Searching in a larger amount of data is hardware intensive task, which benefits from centralized execution and thus is most efficiently done by a remote service indexing data. In the case of Gnosis thousands of event descriptions for prediction markets have to be served to DApps. We present an approach using smart contracts to verify search results by using a smart contract as a judge to decide on censorship of results and punishing bad behavior.

Ethereum allows to run decentralized applications, saving their state on the blockchain. However, changing the state of applications and storing data is very costly. According to the yellow paper, the fee is 20k gas to store a 256 bit word. Storing 1KB of data costs 0.0128 ETH (0.18 USD; 1ETH = 14.5USD): $20k * (1024/32) * 20$ GWei = 0.0128 ether.

Besides limited storage capabilities, searching for data on the blockchain is very difficult as well. As of today, the only way to query the blockchain via JSON RPC is executing a call to a predefined contract function. Searching across different contracts is very limited as the same access controls apply as for regular transactions.

IPFS and Swarm on the other hand are protocols designed to create a permanent and decentralized method of storing and sharing files. Storing is very cheap but both protocols come without search capabilities.

Our approach allows everyone to add records to a public database, making all records searchable from any DApp. Search results can be verified and censorship can be detected, resulting in a punishment of the search provider.

9.6.1 Indexing Data

To allow a decentralized search, all indexed data should be available to everyone, so anyone is able to start an indexing service based on the same data source. As the indexed data is too large to be saved on the blockchain, alternative storage solutions like IPFS have to be used. Everyone can easily add files to IPFS but the hash used to retrieve the data from IPFS is only known to the user adding the file. To allow everyone discovering files included in the index, the IPFS hash of every file added to the index should be saved on the blockchain as a reference.

This can be done by adding the IPFS hash to a merkle tree maintained in a smart contract. By doing so, additional information available at transaction execution time, like the sender and timestamp can be included in the root hash:

bytes32 rootHash = sha256(rootHash, documentHash, now, msg.sender);

By adding the IPFS hash to the hash chain, we can show anytime later, that an IPFS document has been part of index at the time it was added. A merkle tree allows to prove membership but not non-membership. An indexing service could provide search results containing documents, which are not part of the merkle tree and it would be impossible to show that they are not part of the tree. A simple approach to prove non-membership is by adding every IPFS hash to a Solidity mapping:

mapping (bytes32 => uint) documents;

The mapping maps IPFS hashes to the block number at the time the document was added. Proof of non-membership is now very simple: Any IPFS object, which is not referenced in this mapping, was never added to the index.

To inform indexing services about new documents, an Ethereum event should be triggered when a new record is added. Every indexing service can listen to the blockchain for events and retrieve the referenced files from IPFS to add them to their index. A newly started indexing service can iterate over all events to load all IPFS objects and restore the entire index.

9.6.2 Query Data

One option to guarantee that a searchresult was not manipulated is indexing all data locally and perform the query against the local index. This won't be a feasible option for most DApps as search indexes can become very large. Another approach to ensure that results are not manipulated is to make manipulation very costly. Everyone, who offers an indexing service, should put up a security deposit. The indexing service will lose the deposit if manipulation can be detected. To detect manipulation search results should be shared between clients receiving search results and other indexing services, which can validate results.

Discover Indexing Service

Before any query can be answered, a connection has to be established between the DApp and the indexing service.

The DApp has to perform a handshake with the indexing service:

1. A DApp sends out a whisper message asking for indexing services.
2. The indexing service responds with signed connection information (protocol, host, port).
3. The DApp waits for the first response and validates, that the indexing service address derived from the signature has put up a security deposit. Two indicators can be used to determine the selected indexing service:
   (a) How big is the security deposit. The larger the deposit, the more expensive is manipulation.

   (b) The faster a whisper response arrives, the faster future queries will be responded, as the indexer is most likely closer to the DApp user. This is important, if the communication is solely done via whisper.
4. A direct IP connection can be established with the given information. In case the indexer or the DApp user don't want to reveal their IP or have restricted access, queries can be send via a 1:1 whisper connection.

Sending Query

To make sure a result cannot be censored a query object has to define all conditions a document has to meet. This can include keywords, which should appear in a document as well as an indication, which documents should be included(e.g. all documents added up to block 300,000).

The query object sent by the user includes:

1. Query string Defines required document properties and sorting. The query could be defined in a SQL like query language: Contains(title, "Gnosis") LIMIT 10; Block number
2. All documents added until this block number should be included in the search process.

The indexing service processes the query and returns a result object:

1. Result
   A list of IPFS hashes, sorted by their ranking.
2. Indexing service signature
   Indexing service sings its result together with the user defined query string and defined block number.

The DApp receives the result object and validates the indexing service signature. To validate the result the DApp can forward the result to a validation service. Validation can be done by any other indexing service. The validator runs the query and validates that its result is equal to the result received by the DApp. If the result should differ, the validation service can challenge the result via the judge smart contract. The validation service has to prove that there is one document, which matches the query and has a higher ranking score than a document included in the result with a higher ranking. If this can be proven, the indexing service delivering the wrong result loses its security deposit, which could be credited to the validator.

There are three ways an indexing service can manipulate results:

1. Indexing service includes documents, which are not part of the index.
2. Indexing service excludes documents, which are part of the index.
3. Ranking score of objects is manipulated and they are returned in false order.

To prove manipulation the judge contract has to receive the following information for validation:

1. Query string
2. Block number
3. Result documents
4. Challenged document
5. Challenge document
6. Indexing service signature

The Judge contract validates:

1. *Is indexing service signature is valid?*
   If the address derived from the signature doesn't match the address of the security deposit holder, the validation process is stopped.

2. *Is **challenged document** part of **result documents**?*
   The challenged document has to be part of the result documents returned by the indexing service. The validation process stops otherwise.
3. *Is **challenged document** not part of index up to the defined block number?*
   Using the document mapping, the validator can easily prove that the challenged document was not part of the index at the defined block number. The indexing service would lose its security deposit in this case.
4. *Is **challenge document** part of index up to the defined block number.*
   The challenge document provided by the validator has to be part of the index at the time of the block defined in the query. If this is not the case, the validation process stops.
5. *Has **challenge document** higher ranking score than **challenged document** but has lower ranking position in result or is not included in result at all.*
   To prove manipulated ranking positions, the ranking algorithm itself has to be validated in the smart contract. The ranking scores for the challenge and challenged document are calculated and compared. If the challenged document has a higher ranking but a lower ranking score than the challenge document, the result was manipulated and the indexing service loses its deposit.

Exclude Data

IPFS has no guarantee for data being available. Unintentional loss of data or intentional not publishing of IPFS data can occur. A possible attack would be to add an IPFS document hash to the smart contract maintaining the index merkle tree but never publish the document. The attacker could reveal the missing document in the validation process later on and show that this document should have been included in the search result. Because it is impossible to detect this in a smart contract, the validator would challenge the result successfully and receive the indexing service's security deposit. For this reason, an indexing service should be able to exclude missing documents from its results explicitly. A hash of all missing docs should be included in the signed result. If an indexing service is excluding available documents, it won't be punished but a validator can respond to the DApp that it owns one of the documents excluded by the indexing service. If the DApp detects multiple cases of falsely declared missing docs, the DApp should consider switching the indexing service. In addition, if a DApp receives results containing IPFS hashes of objects, which are missing and cannot be loaded, the DApp should consider switching the indexing service as well.