

Augur: a Decentralized Oracle and Prediction Market Platform (v2.0)

Jack Peterson, Joseph Krug, Micah Zoltu, Austin K. Williams, and Stephanie Alexander

Forecast Foundation

(Dated: August 12, 2022)

Augur is a trustless, decentralized oracle and platform for prediction markets. The outcomes of Augur’s prediction markets are chosen by users that hold Augur’s native Reputation token, who stake their tokens on the actual observed outcome and, in return, receive settlement fees from the markets. Augur’s incentive structure is designed to ensure that honest, accurate reporting of outcomes is always the most profitable option for Reputation token holders. Token holders can post progressively-larger Reputation bonds to dispute proposed market outcomes. If the size of these bonds reaches a certain threshold, Reputation splits into multiple versions, one for each possible outcome of the disputed market; token holders must then exchange their Reputation tokens for one of these versions. Versions of Reputation which do not correspond to the real-world outcome will become worthless, as no one will participate in prediction markets unless they are confident that the markets will resolve correctly. Therefore, token holders will select the only version of Reputation which they know will continue to have value: the version that corresponds to reality.

Augur is a trustless, decentralized oracle and prediction market platform. In a prediction market, individuals can speculate on the outcomes of future events; those who forecast the outcome correctly win money, and those who forecast incorrectly lose money [1–3]. The price of a prediction market can serve as a precise and well-calibrated indicator of how likely an event is to occur [4–7].

Using Augur, people will have the ability to trade in prediction markets at very low cost. The only significant expenses participants assume is compensation to market creators and to users that report on the outcomes of markets once the event has taken place. The result is a prediction market where trust requirements, friction, and fees will be as low as competitive market forces can drive them.

Historically, prediction markets have been centralized. The simplest way to aggregate trades in a prediction market is for a trustworthy entity to maintain a ledger; similarly, the simplest way to determine the outcome of an event and distribute payouts to traders is for an impartial, trusted judge to determine the outcomes of the markets. However, centralized prediction markets have many risks and limitations: they do not allow global participation, they limit what types of markets can be created or traded, and they require traders to trust the market operator to not steal funds and to resolve markets correctly.

Augur aims to resolve markets in a fully decentralized way. Decentralized, trustless networks, such as Bitcoin[8] and Ethereum[9], eliminate the risk that self-interest will turn into corruption or theft. The only role of the Augur developers is to publish smart contracts to the Ethereum network. The Augur contracts are totally automated: the developers do not have the ability to spend funds that are held in escrow on-contract, do not control how markets resolve, do not approve or reject trades or other transactions on the network, cannot undo trades, cannot modify or cancel orders, etc. The Augur *oracle* allows information to be migrated from the real world to a blockchain without relying on a trusted intermediary. Augur will be the world’s first decentralized oracle.

I. HOW AUGUR WORKS

Augur markets follow a four-stage progression: *creation*, *trading*, *reporting*, and *settlement*. Anyone can create a market based on any real-world event. Trading begins immediately after market creation, and all users are free to trade on any market. After the event on which the market is based has occurred, the outcome of the event is determined by Augur’s oracle. Once the outcome is determined, traders can close out their positions and collect their payouts.

Augur has a native token, Reputation (REP)¹. REP is needed by market creators and by reporters when they report on the outcome of markets created on the Augur platform. Reporters report on a market by *staking* their REP on one of the market’s possible outcomes. By doing this, the reporter declares that the outcome on which the stake was placed matches the real-world outcome of the market’s underlying event. The consensus of a market’s reporters is considered the “truth” for the purpose of determining the market’s outcome. If a reporter’s report of a market’s outcome does not match the consensus reached by the other reporters, Augur redistributes the REP staked on the non-consensus outcome by this reporter to the reporters that reported with the consensus.

By owning REP, and participating in the accurate reporting on the outcomes of events, token holders are entitled to a portion of the fees on the platform. Each staked REP token entitles its holder to an equal portion of Augur’s market fees. The more REP a reporter owns, and reports correctly with, the more fees they will earn for their work in keeping the platform secure.

Although REP plays a central role in Augur’s operations, it is not used to trade in Augur’s markets. Traders are not required to participate in the reporting process, so they will never need to own or use REP.

¹REP is an ERC777 token on the Ethereum network.



Figure 1. Simplified outline of the lifetime of a prediction market.

A. Market Creation

Augur allows anyone to create a market about any upcoming event. The *market creator* sets the *event end time* and chooses a *designated reporter* to report the outcome of the event. The designated reporter does not unilaterally decide the outcome of the market; the community always has an opportunity to dispute and correct the designated reporter’s report.

Next, the market creator chooses a *resolution source* that reporters should use to determine the outcome. The resolution source may simply be “common knowledge”, or it may be a specific source, such as “The United States Department of Energy”, *bbc.com*, or the address of a particular API endpoint.² They also set a *creator fee*, which is the fee paid to the market creator by traders who settle with the market contract (see Section ID for details on fees). Finally, the market creator posts two bonds: the *validity bond*, and the *creation bond*.

The validity bond is paid in DAI and is returned to the market creator if the market resolves to any outcome other than *invalid*.³ The validity bond incentivizes market creators to create markets based on well-defined events with objective, unambiguous outcomes. The size of the validity bond is set dynamically, based on the proportion of invalid outcomes in recent markets.⁴

The creation bond, paid in REP, is returned to the market creator if and only if the market’s designated reporter actually reports during the first 24 hours after the market’s *event end time* and if the market ends up resolving to the same outcome that the designated reporter reported. If the designated reporter does not submit their report during the allotted 24 hours window, then the market creator forfeits the creation bond and it is given to the *first public reporter* who reports on the market (see Section IC6). This incentivizes the market creator to choose a reliable designated reporter who will report truthfully – which should help markets resolve quickly.

²For example, if a market on “The high temperature (in degrees Fahrenheit) on April 10, 2018 at the San Francisco International Airport, as reported by Weather Underground” specifies a resolution source of <https://www.wunderground.com/history/airport/KSFO/2018/4/10/DailyHistory.html>, reporters would simply go to that URL and enter the high temperature displayed there as their report.

³An *invalid market* is a market determined to be invalid by reporters because none of the outcomes listed by the market creator is correct, or because the market wording is ambiguous or subjective; see Section III G for discussion.

⁴See Appendix B 1 for details.

In the event that the designated reporter fails to report, the creation bond is given to the first public reporter in the form of stake on their reported outcome, so that the first public reporter receives the creation bond if and only if they report correctly. As with the validity bond, the creation bond is adjusted dynamically based on the proportion of designated reporters who failed to report on time during the previous dispute window and the proportion of markets that resolve to some outcome other than the one reported by the designated reporter.⁵

The market creator creates the market and posts all required bonds via a single Ethereum transaction. Once the transaction is confirmed, the market is live and trading begins.

B. Trading

Market participants forecast the outcomes of events by trading *shares* of those market outcomes. A *complete set of shares* is a collection of shares that consists of one share of each possible valid outcome of the event [10]. Complete sets are created by Augur’s on-contract matching engine as needed to complete trades⁶.

For example, consider a market that has two possible outcomes, A and B. Alice is willing to pay 0.7 DAI for a share of A and Bob is willing to pay 0.3 DAI for a share of B. First, Augur matches these orders and collects a total of 1 DAI from Alice and Bob.⁷ Then Augur creates a complete set of shares, giving Alice the share of A and Bob the share of B. This is how shares of outcomes come into existence. Once the shares are created, they can be traded freely.

The Augur trading contracts maintain an order book for every market created on the platform. Anybody can create a new order or fill an existing order at any time. Orders are filled by an automated matching engine that exists within Augur’s smart contracts. Requests to buy or sell shares are fulfilled immediately if there is a matching order already on the order book. It may be filled by buying shares from or selling shares to other participants, which, may involve issuing new complete sets or

⁵See Appendix B 2 for details.

⁶In practice, each share in a complete set is an ERC777 token on the Ethereum network.

⁷The 1 DAI figure is used here for ease of discussion. The actual cost of a complete set of shares is much smaller than this; see docs.augur.net/#number-of-ticks for details.

closing out existing complete sets. Augur’s matching engine always sequesters the minimum amount of shares and/or cash needed to cover the value at risk. If there is no matching order, or the request can be only partially filled, the remainder is placed on the order book as a new order.

Orders are never executed at a worse price than the limit price set by the trader, but may be executed at a better price. Unfilled and partially-filled orders can be removed from the order book by the order’s creator at any time. Fees are paid by traders only when complete sets of shares are sold; settlement fees are discussed in more detail in Section ID.

While most trading of shares is expected to happen before market settlement, shares can be traded any time after market creation. All Augur assets – including shares in market outcomes, participation tokens, shares in dispute bonds, and even ownership of the markets themselves – are transferable at all times. In practice, all of these assets take the form of ERC777 tokens.

C. Reporting

Once a market’s underlying event occurs, the outcome must be determined in order for the market to finalize and begin settlement. Outcomes are determined by Augur’s oracle, which consists of profit-motivated reporters, who simply report the actual, real-world outcome of the event. Anyone who owns REP may participate in the reporting and disputing of outcomes. Reporters whose reports are consistent with objective reality are financially rewarded, while those whose reports are not consistent with objective reality are financially penalized (see Section ID 2).

1. Dispute Windows

Augur’s reporting system runs on a cycle of consecutive 7-day long *dispute windows*. All fees collected by Augur during a given dispute window are added to the *reporting fee pool* for that dispute window. At the end of the dispute window, the reporting fee pool is paid out to REP holders who participated in the reporting process. Reporters receive rewards in proportion to the amount of REP they staked during that dispute window. Participation includes: staking during an initial report, disputing a tentative outcome, or purchasing *participation tokens*.

2. Participation Tokens

During any dispute window, REP holders may purchase any number of participation tokens⁸ for one at-

torep⁹ each. At the end of the dispute window, they may redeem their participation tokens for one attorep each, in addition to a proportional share of the dispute window’s *reporting fee pool*. If there were no actions (*e.g.*, submitting a report or disputing a report submitted by another user) needed of a reporter, the reporter may purchase participation tokens to indicate that they showed up for the dispute window. Just like staked REP, participation tokens may be redeemed by their owners for a *pro rata* portion of fees in this dispute window.

Participation tokens are primarily the means by which Augur pays fees to REP holders, but they may also serve as an additional incentive for REP holders to monitor the platform at least once per week. Even REP holders who do not want to participate in the reporting process may be incentivized to check-in with Augur once per 7-day dispute window in order to buy participation tokens and collect fees. This regular, active checking-in will ensure that they are familiar with how to use Augur, will be aware of forks when they occur, and thus should be more ready to participate in forks when they happen.

3. Market State Progression

Augur markets can be in seven different states after creation. The potential states, or “phases”, of an Augur market are as follows:

- Pre-reporting
- Designated Reporting
- Open Reporting
- Dispute Round
- Waiting for Window
- Fork
- Finalized

The relationship between these states can be seen in Fig. 2.

4. Pre-reporting

The *pre-reporting* or *trading* phase (Fig. 1) is the time period that begins after trading has begun in the market, but before the market’s event has come to pass. Generally, this is the most active trading period for any given Augur market. Once the event end date has passed, the market enters the *designated reporting* phase (Fig. 2a).

⁸Participation tokens are ERC777 tokens on the Ethereum network.

⁹One *attorep* is 10^{-18} REP.




Figure 2. Reporting flowchart.

5. Designated Reporting

When creating a market, market creators are required to choose a designated reporter and post a creation bond. During the designated reporting phase (Fig. 2a) the market’s designated reporter has up to 24 hours to report on the outcome of the event. If the designated reporter fails to report within the allotted 24 hours, the market creator forfeits the creation bond, and the market automatically enters the *open reporting* phase (Fig. 2b).

If the designated reporter submits a report on time then the creation bond is placed as stake on the reported outcome, which will be forfeited if the market finalizes to any outcome other than the one they reported.¹⁰ As soon as the designated reporter submits its report, the market

enters the *dispute round* phase (Fig. 2c), and the reported outcome becomes the market’s *tentative outcome*.

6. Open Reporting

If the designated reporter fails to report within the allotted 24 hours, the market creator forfeits the creation bond, and the market immediately enters the *open reporting* phase (Fig. 2b). As soon as the market enters the open reporting phase, anyone can report the outcome of the market. When the designated reporter fails to report, the first reporter who reports on the outcome of a market is called the market’s *first public reporter*.

The market’s first public reporter receives the forfeited creation bond in the form of stake on their chosen outcome, so they may claim the creation bond only if their reported outcome agrees with the market’s final outcome.

The first public reporter does *not* need to stake any of their own REP when reporting the outcome of the mar-

¹⁰Forfeited stake is used to reward honest reporters and disputers; see Section ID 2 for details.

ket. In this way, any market whose designated reporter fails to report is expected to have its outcome reported by *someone* very soon after entering the open reporting phase.

Once an *initial report* has been received by the initial reporter (whether it was the designated reporter or first public reporter), the reported outcome becomes the market’s tentative outcome, and the market enters the dispute round phase (Fig. 2c).

7. Dispute Round

The dispute round (Fig. 2c) is a phase during which any REP holder has the opportunity to dispute the market’s *tentative outcome*. A dispute round may last up to 7 days (with the exception of the very first dispute round, which may last up to 24 hours). At the beginning of a dispute round, a market’s tentative outcome is the outcome that will become the market’s final outcome if it is not successfully disputed by REP holders. A dispute consists of *staking* REP (referred to as *dispute stake* in this context) on an outcome *other than* the market’s current tentative outcome. A dispute is *successful* if the total amount of dispute stake on some outcome meets the *dispute bond size* required for the current round. The dispute bond size is computed as follows.

Let A_n denote the total stake over all of this market’s outcomes at the beginning of dispute round n . Let ω be any market outcome *other than* the market’s tentative outcome at the beginning of this dispute round. Let $S(\omega, n)$ denote the total amount of stake on outcome ω at the beginning of dispute round n . Then the size of the *dispute bond* needed to successfully dispute the current tentative outcome in favor of the new outcome ω during round n is denoted $B(\omega, n)$ and is given by:

$$B(\omega, n) = 2A_n - 3S(\omega, n) \quad (1)$$

The bond sizes are chosen this way to ensure a fixed ROI for reporters who successfully dispute false outcomes (see Section IID).

The dispute bonds need not be paid in their entirety by a single user. The Augur platform allows participants to crowdsource dispute bonds. Any user who sees an incorrect tentative outcome can dispute that outcome by staking REP on an outcome other than the tentative outcome. If any outcome (other than the tentative outcome) accumulates enough dispute stake to fill its dispute bond, the current tentative outcome will be successfully disputed.

In the case of a successful dispute, One of three things will happen: the market will either undergo another dispute round immediately, the market will wait until the next dispute window begins before undergoing another dispute round, or the market will enter the *fork* state (Fig. 2e). If the size of the filled dispute bond is greater

than or equal to 2.5% of all theoretical REP¹¹, then the market will enter the fork state. If the size of the filled dispute bond is less than 2.5% of all theoretical REP but greater than or equal to 0.02% of all theoretical REP, then the newly chosen outcome becomes the market’s new tentative outcome and the market enters the *waiting for window* phase (Fig. 2d). If the size of the filled dispute bond is less than 0.02% of all theoretical REP, then the newly chosen outcome becomes the market’s new tentative outcome and the market immediately enters another dispute round.

All dispute stake is held in escrow during the dispute round. If a dispute bond is unsuccessful, then the dispute stake is returned to its owners at the end of the dispute round. If no dispute is successful during the dispute round, then the market enters the *finalized* state (Fig. 2f), and its tentative outcome is accepted as its *final outcome*. A market’s final outcome is the tentative outcome that passes through a dispute round without being successfully disputed, or is determined via a fork. Augur’s contracts treat final outcomes as *truth* and pay out accordingly.

All unsuccessful dispute stake is returned to the original owners at the end of every dispute round. All successful dispute stake is applied to the outcome it championed, and remains there until the market is finalized (or until a fork occurs in some other Augur market). All dispute stake (whether successful or unsuccessful) will receive a portion of the *reporting fee pool*¹² from the current dispute window.

8. Waiting for Window

If a market’s tentative outcome is disputed with a bond greater than or equal to 0.02% of all theoretical REP but less than 2.5% of all theoretical REP, then the market enters the waiting for window phase (Fig. 2d), before undergoing another dispute round. The purpose of this is simply to slow down the dispute process as the bonds get larger – giving honest participants more time to crowd-fund the larger dispute bond. This reduces the risk of a critical failure mode: one where the oracle resolves incorrectly because honest participants didn’t have time to raise the funds needed to dispute a false tentative outcome.

During this phase, reporting for the market is on hold until end of the current dispute window. Once the next

¹¹ All theoretical REP means the total theoretical supply of REP in the universe. In other words, Sum of the total amount of REP which exist in the universe and the total amount of REP which exist in the other universe and can be migrated to the universe.

¹² Any reporting fees and validity bonds collected during a dispute window get added to that dispute window’s reporting fee pool. At the end of the dispute window, the reporting fee pool is paid out to users in proportion to the amount of REP they staked during that dispute window.

dispute window begins, the market enters the *dispute round* phase.

9. Fork

The fork state (Fig. 2e) is a special state that lasts up to 60 days. Forking is the market resolution method of last resort; it is a very disruptive process and is intended to be a rare occurrence. A fork is caused when there is a market with an outcome with a successfully-filled dispute bond of at least 2.5% of all theoretical REP. This market is referred to as the *forking market*.

When a fork is initiated, a 60-day¹³ *forking period* begins. Disputing for all other non-finalized markets is put on hold until the end of this forking period. The forking period is much longer than the usual dispute window because the platform needs to provide ample time for REP holders and service providers (such as wallets and exchanges) to prepare. A fork's final outcome cannot be disputed.

Every Augur market and all REP tokens exist in some *universe*. REP tokens can be used to report on outcomes (and thus earn fees) *only* for markets that exist in the same universe as the REP tokens. When Augur first launches, all markets and all REP will exist together in the *genesis universe*.

When a market forks, new universes are created. Forking creates a new *child universe* for each possible outcome of the forking market (including *Invalid*). For example, a "Yes/No" market has 3 possible outcomes: *Yes*, *No*, and *Invalid*. Thus, a "Yes/No" forking market will create three new child universes: universe *Yes*, universe *No*, and universe *Invalid*. Initially, these newly created universes are empty: they contain no markets or REP tokens.

When a fork is initiated, the *parent universe* becomes permanently *locked*. In a locked universe, no new markets may be created and no REP can be staked on any market. Therefore markets in a locked universe cannot be finalized. However users may continue trading shares in markets in locked universes. In order for markets or REP tokens in the locked universe to be useful, they must first be migrated to a child universe.

During the forking period, holders of REP tokens in the parent universe may migrate their tokens to a child universe of their choice. This choice should be considered carefully, because migration is one-way; it cannot be reversed. Tokens cannot be sent from one sibling universe to another. *Migration is a permanent commitment of REP tokens to a particular outcome of the forking market*. REP tokens that migrate to different child universes

ought to be considered entirely separate tokens, and service providers like wallets and exchanges ought to list them as such.

Any REP tokens that have not been migrated out of the parent universe 60 days after the fork started will be permanently locked in the parent universe. Such tokens are expected to lose all value, so it is of paramount importance that REP holders migrate their tokens anytime a fork happens.

When a fork is initiated, all REP staked on all non-forking markets is *unstaked* so that it is free to be migrated to a child universe during the forking period.¹⁴

Whichever child universe receives the most migrated REP by the end of the forking period becomes the *winning universe*, and its corresponding outcome becomes the final outcome of the forking market. Un-finalized markets in the parent universe may be migrated only to the winning universe and, if they have received an initial report, are reset back to the waiting for window phase.

Reporters that have staked REP on one of the forking market's outcomes cannot change their position during a fork. REP that was staked on a forking market's outcome in the parent universe can be migrated only to the child universe that corresponds to that outcome. For example, if a reporter helped fulfill a successful dispute bond in favor of outcome A during some dispute round, then the REP they have staked on outcome A can only be migrated to universe A during a fork.

Sibling universes are entirely disjoint. REP tokens that exist in one universe cannot be used to report on events or earn rewards from markets in another universe. Since users presumably will not want to create or trade on markets in a universe whose oracle is untrustworthy, REP that exists in a universe that does not correspond to objective reality is unlikely to earn its owner any fees, and therefore should not hold any significant market value. Therefore, REP tokens migrated to a universe which does not correspond to objective reality should hold no market value, regardless of whether or not the objectively false universe ends up being the winning universe after a fork. This has important security consequences, which we discuss in Section II.

10. Finalized

A market enters the finalized state (Fig. 2f) if it passes through a 7-day dispute round without having its tentative outcome successfully disputed, or after completion of a fork. The outcome of a fork cannot be disputed and is always considered final at the end of the forking period.

¹³Forking periods can be less than 60 days: a forking period ends when either 60 days have passed, or more than 50% of all theoretical REP is migrated to some child universe. However even after the end of a forking period, REP in the parent universe can be migrated to the child universe if it is within 60 day of fork initiation.

¹⁴The only exception is the REP staked by the initial reporter when they made the initial report. That REP remains staked on the initial reported outcome and is automatically migrated to the child universe that wins the fork. This happens for technical reasons unrelated to the mechanism design.

Once a market is finalized, traders can settle their positions directly with the market. When a market enters the finalized state, we refer to its chosen outcome as the *final outcome*.

D. Market Settlement

A trader can close their position in one of two ways: by selling the shares they hold to another trader in exchange for currency, or by settling their shares with the market. Recall that every share comes into existence as part of a complete set when a total of 1 DAI has been escrowed with Augur.⁷ To get that 1 DAI out of escrow, traders must give Augur either a complete set or, if the market has finalized, a share of the winning outcome. When this exchange happens we say traders are *settling with the market contract*.

For example, consider a non-finalized market with possible outcomes A and B. Suppose Alice has a share of outcome A that she wants to sell for 0.7 DAI and Bob has a share of outcome B that he wants to sell for 0.3 DAI. First, Augur matches these orders and collects the A and B shares from the participants. Then Augur gives 0.7 DAI (minus fees) to Alice and 0.3 DAI (minus fees) to Bob.

As a second example, consider a finalized market whose winning outcome is A. Alice has a share of A and wants to cash it in. She sends her share of A to Augur and in return receives 1 DAI (minus fees).

1. Settlement Fees

The only time Augur levies fees is when market participants are settling with the market contract. Augur levies two fees during settlement: the creator fee, and the reporting fee. Both of these fees are proportional to the amount being paid out. So, in the pre-finalized settlement example above, where Alice receives 0.7 DAI and Bob receives 0.3 DAI, Alice would pay 70% of the fees while Bob would pay 30%.

The creator fee is set by the market creator during market creation, and is paid to the market creator upon settlement. The reporting fee is set dynamically (see Section II C) and is paid to reporters who participate in the reporting process.

2. Reputation Redistribution

If a market finalizes without initiating a fork, all REP staked on any outcome other than the market's final outcome is forfeited. Twenty percent of the forfeited stake is burned, and the remainder is distributed to the users who staked on the market's final outcome in proportion to the amount of REP they staked. The dispute bond

sizes and the amount burned are chosen such that anyone who successfully disputes an outcome in favor of the market's final outcome is rewarded with a 40% ROI on their dispute stake.¹⁵ This is a strong incentive for reporters to dispute false tentative outcomes.

II. INCENTIVES AND SECURITY

There is a strong relationship between the market cap of REP and the trustworthiness of Augur's forking protocol. If the market cap of REP is large enough¹⁶, and attackers are economically rational, then the outcome that wins the fork should correspond to objective reality. In fact, it would be possible for Augur to function properly without using designated reporters and dispute rounds. Using *only* the forking process, the oracle would report truthfully.

However, forks are disruptive and time consuming. A fork takes up to 60 days to resolve a single market, and can resolve only one market at a time. During the 60 days in which the forking market is being resolved, all other non-finalized markets are put on hold.¹⁷ Service providers have to update, and REP holders have to migrate their REP to one of the new child universes. Therefore, forks should be used only when they are absolutely necessary. Forking is the nuclear option.

Fortunately, once it has been established that forks can be trusted to determine truth, incentives can be used to encourage participants to behave honestly without having to actually initiate a fork. *It is the credible threat of a fork, and the belief that the fork will resolve correctly, that are the cornerstones of Augur's incentive system.*

Next, we discuss the conditions under which the forking system can be trusted to determine truth. We then discuss the incentive system and how it encourages quick and correct resolution of all markets.

A. Integrity of the Forking Protocol

Here we discuss the reliability of the forking process and the conditions under which it can be trusted. For ease of discussion, when referring to forks, we will refer to the child universe that corresponds to objective reality as the *True* universe, and any other child universe as a *False* universe. We will refer to the child universe which receives the most REP migration during the forking period as the winning universe and all other child universes as losing universes.

Naturally, we always want the *True* universe to be the winning universe, and the *False* universes to be the losing

¹⁵See Theorem 2 in Appendix A.

¹⁶See Section II A for details.

¹⁷Traders can continue trading on those markets, but those markets cannot finalize until after the forking period.

universes. We say that the forking protocol has been successfully attacked whenever a **False** universe ends up being the winning universe of a fork – thus resulting in the forking market (and, potentially, all non-finalized markets) being paid out incorrectly.

Our approach to securing the oracle is to arrange matters such that the maximum benefit to a successful attacker is less than the minimum cost of performing the attack. We formalize this below.

1. Maximum Benefit to an Attacker

An attacker who successfully attacks the oracle would cause all non-finalized Augur markets to migrate to a **False** universe. If the attacker controls the majority of REP in the **False** universe, the attacker can then force all non-finalized markets to resolve however she wants. In the most extreme case, she would also be able to capture all funds escrowed in all of those markets.¹⁸

Definition 1. We define, and denote by I_a , Augur’s *native open interest* as the value of the sum of all funds escrowed in unfinalized Augur markets.¹⁹

Definition 2. We define a *parasitic market* as any market that does not pay reporting fees to Augur, but does resolve in accordance with the resolution of a native Augur market.

Definition 3. We define, and denote by I_p , the *parasitic open interest* as the value of the sum of all funds escrowed in all parasitic markets that resolve in accordance to non-finalized, native Augur markets.

In the most extreme case, an attacker would also be able to capture all funds in all parasitic markets which resolve in accordance to non-finalized, native Augur markets.

Observation 1. The maximum (gross) benefit to an attacker who successfully attacks the oracle is $I_a + I_p$.

2. Parasitic Open Interest is Unknowable

Augur can accurately and efficiently measure I_a . However, I_p cannot be known in general, as there may exist arbitrarily many offline parasitic markets, each with arbitrarily large open interest. Since the maximum possible benefit to an attacker includes the unknowable quantity I_p , one can never be objectively certain that the oracle is secure against economically rational attackers.²⁰

However, if we are willing to assert that I_p is reasonably bounded in practice, then we can define conditions under which we may assert that the oracle is secure.

3. Minimum Cost of a Successful Attack

Next, consider the cost of attacking the oracle. Let P denote the price of REP. Let ϵ denote one attorep²¹. Let M denote the total amount of REP in existence (the “money supply” of REP). Let S denote the proportion of M that will be migrated to the **True** universe during the forking period of a fork.

Thus the product SM represents the absolute amount of REP migrated to the **True** universe during the forking period of a fork, and the product PM is the market cap of REP.

Let P_f denote the price of REP migrated to a **False** universe of the attacker’s choosing. Note that if $P \leq P_f$ then the oracle would not be secure against economically rational attackers, because it would be at least as profitable to migrate REP to the **False** universe as it would be to not migrate at all.

4. Integrity

Assumption 1. *Reporters that are not attackers will never migrate REP to a False universe during a fork.*²²

By design, a successful attack on the oracle requires more REP to be migrated to some **False** universe than to the **True** universe during the forking period of a fork. By assumption, only the attacker will migrate REP to a **False** universe. The amount of REP migrated to the **True** universe during the forking period is denoted by SM . Thus, for an attacker to be successful, they must migrate at least $SM + \epsilon$ REP. For simplicity, we will ignore the negligible ϵ , and say that a successful attack requires migrating at least SM REP, which has a value of SMP before the migration, to some **False** universe.

If an attacker migrates SM REP during the forking period of a fork, they will receive SM REP on the child universe to which they migrate. If the attacker migrates to a **False** universe then the value of those coins becomes SMP_f . Thus the minimum cost to the attacker is $(P - P_f)SM$.

Observation 2. The minimum amount of REP a successful attacker must migrate to a **False** universe during a fork is SM , which costs the attacker $(P - P_f)SM$.

Note that if $S > \frac{1}{2}$ then an attack is *impossible* because there does not exist enough REP outside of the **True** universe for any **False** universe to become the winning universe.

¹⁸This would require the attacker to capture *all* shares of some given outcome, and then force the market to finalize to that outcome.

¹⁹This includes external markets that pay reporting fees to Augur.

²⁰This is true of all public oracles, even centralized ones.

²¹One attorep is 10^{-18} REP.

²²There may be cases where some non-malicious reporters do migrate REP to a **False** universe accidentally or carelessly. However, such behavior is, in practice, indistinguishable from collaborating with an attacker.

Pitted against economically rational attackers, the oracle will resolve to outcomes that correspond to objective reality if the maximum benefit to an attacker is less than the minimum cost of attack. By observations 1 & 2 we can see that this occurs whenever $S > \frac{1}{2}$ or $I_a + I_p < (P - P_f)SM$. This gives us our formal definition of integrity.

Definition 4. (Integrity Property) The forking protocol has *integrity* whenever $S > \frac{1}{2}$ or whenever $I_a + I_p < (P - P_f)SM$.

The above inequality can be solved for PM to see the relationship between forking protocol integrity and the market cap of REP.

Theorem 1. (Market Cap Security Theorem) *The forking protocol has integrity if and only if:*

1. $S > \frac{1}{2}$, or
2. $P_f < P$ and the market cap of REP is greater than $\frac{(I_a + I_p)P}{(P - P_f)S}$.

Proof. Suppose the forking protocol has integrity. Then, by definition, $S > \frac{1}{2}$ or $I_a + I_p < (P - P_f)SM$. Suppose $I_a + I_p < (P - P_f)SM$. Since $I_a + I_p \geq 0$ and $SM > 0$, we know that $P_f < P$. Then, solving $I_a + I_p < (P - P_f)SM$ for PM , we see that $\frac{(I_a + I_p)P}{(P - P_f)S} < PM$. Thus the first direction is proved.

Now suppose that $S > \frac{1}{2}$, or that $P_f < P$ and $\frac{(I_a + I_p)P}{(P - P_f)S} < PM$. If $S > \frac{1}{2}$, then the forking protocol has integrity by definition. If $P_f < P$ and $\frac{(I_a + I_p)P}{(P - P_f)S} < PM$, then, solving the inequality for $I_a + I_p$, we see that $I_a + I_p < (P - P_f)SM$, and the forking protocol has integrity. \square

B. Our Assumptions and Their Consequences

We believe traders will not want to trade on Augur in a universe where reporters have lied. We also believe that market creators will not pay to create Augur markets in a universe where there are no traders. In a universe without markets or trading, REP does not provide any dividends to those holding it. Therefore, we believe REP sent to a False universe will hold no non-negligible market value and we model this by letting $P_f = 0$.

An attacker should never migrate more than 50% of all theoretical REP (plus on attorep) to a false universe²³. We believe that approximately all theoretical REP will be migrated out of the parent universe during a forking period, because failure to do so is expected to result in loss of token value for no benefit. Therefore, we think

it is reasonable to expect at least 50% of REP (minus one attorep) to be migrated to the True outcome during the forking period of a fork, and we model this by letting $S \geq \frac{1}{2}$. We are also willing to accommodate parasitic open interest as large as 50% of the native open interest, and so we let $I_a \geq 2I_p$.

Under these assumptions, Theorem 1 tells us that the forking protocol has integrity whenever the market cap of REP is at least 3 times the native open interest.

C. Market Cap Nudges

Augur gets information about the price of REP via a collection of trusted third-party price oracles²⁴. This gives Augur the ability to compute the current market cap of REP. Augur can also measure the current native open interest, and can thus determine what market cap ought to be targeted in order to meet Augur's integrity requirements.

Every universe begins with a default reporting fee of 1%. If the current market cap is below the target, then reporting fees are automatically increased (but will never be higher than 33.3%), putting upward pressure on the price of REP and/or downward pressure on new native open interest. If the current market cap is above the target, then reporting fees are automatically decreased (but will never be lower than 0.01%) so that traders are not paying more than needed to keep the system secure.

The reporting fees are determined as follows. Let r be the reporting fee from the previous window, let t be the target market cap, and let c be the current market cap. Then the reporting fee for the current dispute window is given by $\max \left\{ \min \left\{ \frac{t}{c}r, \frac{333}{1000} \right\}, \frac{1}{10,000} \right\}$.

D. Leveraging the Threat of a Fork

As mentioned above, forks are a disruptive and slow way for markets to reach finalization. Rather than using the forking process to resolve every market, Augur leverages the *threat* of a fork to resolve markets efficiently [11].

Recall that any stake successfully disputing an outcome in favor of the market's final outcome will receive a 40% ROI on their dispute stake.²⁵ In the event of a fork, any REP staked on any of the market's false outcomes should lose all economic value, while any REP staked on the market's true outcome is rewarded with 40% more REP in the child universe that corresponds to the market's true outcome (regardless of the outcome of the fork). Therefore, if pushed to a fork, REP holders

²⁴In the future, Augur may be able to learn the price of REP without having to trust third parties. This is an active area of research.

²⁵Measured in REP that exists in a universe that corresponds to the market's final outcome; see Theorem 2 in Appendix A.

²³Doing so unnecessarily increases the cost of attack.

who dispute false outcomes in favor of true outcomes will always come out ahead, while REP holders who staked on false outcomes will see their REP lose all economic value.

We believe this situation is sufficient to guarantee that all false tentative outcomes will be successfully disputed.

III. POTENTIAL ISSUES & RISKS

A. Parasitic Markets

Recall that a parasitic market is any market that does not pay reporting fees to Augur, but does resolve in accordance with the resolution of a native Augur market. Because parasitic markets do not have any reporters to pay, they can offer the same service as Augur with lower fees. This can have serious consequences for the integrity of Augur’s forking protocol.

In particular, if parasitic markets attract trading interest away from Augur, then Augur’s reporters will receive less in reporting fees. This would put downward pressure on the market cap of REP. If the market cap of REP falls too low, the integrity of the forking protocol is put in jeopardy (Theorem 1). As a result, parasitic markets have the potential to threaten the long term viability of Augur, and should be vehemently opposed.

The oracle parasite problem has not been solved – and may be provably unsolvable – even for centralized systems. That said, at the time of this writing, we have observed no significant parasitic interest leveraging Augur’s oracle.

B. Volatility of Open Interest

Large, sudden, unexpected, and short-lived increases in open interest – like those that may be seen during a popular sporting event – result in rapid increases in the market cap requirement for forking protocol integrity (Theorem 1). When the market cap requirement exceeds the market cap, there is a risk of economically rational attackers causing a fork to resolve incorrectly. While Augur does attempt to nudge the market cap and/or the open interest back into a safe ratio during such situations (see Section II C), these nudges are reactionary and are adjusted only once per 7-day dispute window.

It is worth noting, however, that speculators who witness the sudden increase in open interest may buy REP in anticipation of the reactionary market cap nudge, thus driving the market cap of REP up, perhaps to a point where the integrity of the forking protocol is no longer threatened. So the length of time during which the oracle is vulnerable may not be long enough for an attacker to successfully exploit the vulnerability.

Additionally, we have increased the security multiplier

from the theoretical minimum of 3^{26} , to a more comfortable value of 5. This means that the fee adjustment algorithm targets a market cap of 5 times the native open interest, rather than 3 times. This should help us absorb large (up to a 66.6% increase), sudden, unexpected, and short-lived increases in open interest without threatening oracle integrity.

C. Inconsistent or Malicious Resolution Sources

During market creation, market creators chose a resolution source that reporters should use to determine the outcome of the event in question. If the market creator chooses an inconsistent or malicious resolution source, honest reporters may lose money.

For example, suppose the market in question has outcomes A and B, and the market creator, Serena, has chosen her own website, `attacker.com`, as the resolution source. After the market’s event end time, Serena – who is also the designated reporter for the market – reports outcome A, and updates `attacker.com` to indicate that outcome B is the correct outcome. Honest reporters who check `attacker.com` will see that the initial report is incorrect and, during the first dispute round, should successfully dispute the tentative outcome in favor of outcome B. Serena would update `attacker.com` to indicate that outcome A is the correct outcome, and the market would then enter its second dispute round. Again, reporters who check `attacker.com` will see that the tentative outcome (outcome B) is incorrect, and may successfully dispute it. Serena can repeat this behavior until the market resolves. No matter how the market resolves, some honest reporters will lose money.

Several variations of this attack exist. Simply ignoring markets with dubious resolution sources is not sufficient, for in the event that such a market causes a fork, all REP holders will have to choose a child universe to which to migrate their REP. Reporters should remain vigilant against markets with dubious resolution sources. Such markets should be publicly identified so reporters can coordinate to make sure such markets finalize as invalid.

D. Self-Referential Oracle Queries

Markets that trade on the future behavior of Augur’s oracle may have undesirable effects on the behavior of the oracle itself [12]. For example, consider a market that trades on the question, “Will any designated reporter fail to submit a report during their three-day forking period

²⁶The theoretical minimum for the security multiplier is 2 when not accounting for any parasitic interest. When account for parasitic interest of up to 50% of the native open interest, the theoretical minimum for the security multiplier is 3.

before December 31, 2018?” Bets placed on the No outcome of this market may act as a perverse incentive for designated reporters to intentionally fail to report. If a designated reporter can buy up enough Yes shares at a low enough price to compensate for the loss of the creation bond, they may intentionally fail to report.

If the market cap of REP is large enough (Theorem 1) then these self-referential oracle queries will not threaten the integrity of the forking protocol. However, they may negatively affect the performance of Augur by causing delays in market finalizations. While markets would still finalize correctly, this sort of behavior is disruptive and undesirable.

E. Uncertain Fork Participation

We cannot know in advance how much REP will be migrated to the True universe during the forking period of a fork, thus we cannot know in advance whether the market cap is large enough for the oracle to have integrity (Theorem 1). Our belief in the integrity of the forking protocol can be no stronger than our belief in our assumption that users will behave rationally and in their own economic self-interest. We assume that at least 50% (minus one attoprep) of all theoretical REP will migrate to the True child universe during the forking period of a fork because that is consistent with self-interested, rational decision making on the behalf of the participants. However, we cannot guarantee this will happen.

F. Responsibility During a Fork

Augur forks differ from blockchain forks in one important respect: after a blockchain fork, a user who owned a coin on the parent chain will now own a coin on both forks. Ignoring replay attacks, blockchain forks pose little risk to users. During an Augur fork, however, a user who owns a REP token in the parent universe can migrate that coin to only one of the child universes. If the user migrates their token to any universe other than the consensus universe, their token may lose all value. Thus migrating REP during the forking period of a fork, before it is clear which child universe has achieved consensus, exposes the user to risk.

This risk is an inherent part of REP ownership. Abstaining from migration during a fork will result in near-certain loss of value for the REP holder. This is a necessary design decision, as the integrity of the oracle relies upon REP holders reporting truthfully during the fork. REP holders cannot be absolved of this responsibility without greatly increasing the cost of system security.

G. Ambiguous or Subjective Markets

Only events that have objectively knowable outcomes are suitable for use in Augur markets. If reporters believe that a market is not suitable for resolution by the platform – for example, because it is ambiguous, subjective, or the outcome is not known by the event end date – they should report the market as *Invalid*. If a market resolves as *Invalid*, traders are paid out at equal values for all possible outcomes; for scalar markets, traders are paid out halfway between the market’s minimum price and maximum price.

It is possible to imagine markets where some reporters are certain that the outcome is A and others are certain that the outcome is B. For example, in 2006, TradeSports allowed its users to speculate on whether North Korea would fire a ballistic missile that would land outside of its airspace before the end of July 2006. On July 5, 2006, North Korea successfully fired a ballistic missile that landed outside of its airspace, and the event was widely reported by the world media and confirmed by many U.S. government sources. However, the U.S. Department of Defense had not confirmed the event, as was required by TradeSports’ contract. TradeSports concluded that the contract’s conditions had not been met, and paid out accordingly.²⁷

This is a case where the spirit of the market – to predict the missile launch – was clearly satisfied, but the letter of the market – to predict whether the U.S. Department of Defense would confirm the launch – was not. TradeSports, being a centralized website, was able to unilaterally declare the outcome of the market. If such a situation arises in an Augur market, REP holders may have differing opinions about how the market should resolve, and stake their REP accordingly. In the worst case, this could result in a fork where REP in more than one child universe maintains a non-zero market value.

ACKNOWLEDGMENTS

We thank Abraham Othman, Alex Chapman, Serena Randolph, Tom Haile, George Hotz, Scott Bigelow, and Peronet Despeignes for their helpful feedback and suggestions.

²⁷See <https://en.wikipedia.org/wiki/Intrade#Disputes> for details.

-
- [1] J. Wolfers and E. Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18(2):107–126, 2004.
 - [2] James Surowiecki. *The Wisdom of Crowds*. Anchor, 2005.
 - [3] R. Hanson, R. Oprea, and D. Porter. Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior & Organization*, 60(4):449–459, 2006.
 - [4] D.M. Pennock, S. Lawrence, C.L. Giles, and F.A. Nielsen. The real power of artificial markets. *Science*, 291:987–988, 2001.
 - [5] C. Manski. Interpreting the predictions of prediction markets. *NBER Working Paper No. 10359*, 2004.
 - [6] J. Wolfers and E. Zitzewitz. Interpreting prediction market prices as probabilities. *NBER Working Paper No. 10359*, 2005.
 - [7] S. Goel, D.M. Reeves, D.J. Watts, and D.M. Pennock. Prediction without markets. In *Proceedings of the 11th ACM Conference on Electronic Commerce*, EC ’10, pages 357–366. ACM, 2010.
 - [8] S. Nakamoto. Bitcoin: a peer-to-peer electronic cash system. <https://bitcoin.org/bitcoin.pdf>, 2008.
 - [9] V. Buterin. A next generation smart contract and decentralized application platform. <https://github.com/ethereum/wiki/wiki/White-Paper>, 2013.
 - [10] J. Clark, J. Bonneau, E.W. Felten, J.A. Kroll, A. Miller, and A. Narayanan. On decentralizing prediction markets and order books. In *WEIS ’14: Proceedings of the 10th Workshop on the Economics of Information Security*, June 2014.
 - [11] A.K. Williams and J. Peterson. Decentralized common knowledge oracles. *Ledger*, 4:157–190, 2019.
 - [12] A. Othman and T. Sandholm. Decision rules and decision markets. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1 - Volume 1*, AAMAS ’10, pages 625–632. International Foundation for Autonomous Agents and Multiagent Systems, 2010.

Appendix A: Finalization Time & Redistribution

We begin with some notation, definitions, and observations.

Definition 5. For a given market M , let Ω_M be the outcome space (or set of outcomes) of M .

Definition 6. For $n \geq 1$ and $\omega \in \Omega_M$, let $S(\omega, n)$ denote the total amount of stake on outcome ω at the beginning of dispute round n . This includes all stake from all successful dispute bonds in favor of ω over all previous dispute rounds.

Definition 7. For $n \geq 1$ and $\omega \in \Omega_M$, let $S(\bar{\omega}, n)$ denote the amount of stake on all outcomes in Ω_M *except for* ω at the beginning of dispute round n :

$$S(\bar{\omega}, n) = \sum_{\substack{\gamma \in \Omega_M \\ \gamma \neq \omega}} S(\gamma, n)$$

Definition 8. For $n \geq 1$, let A_n denote the total stake over all outcomes M at the beginning of dispute round n :

$$A_n = \sum_{\omega \in \Omega_M} S(\omega, n)$$

Observation 3. It follows that $A_n - S(\omega, n) = S(\bar{\omega}, n)$.

Definition 9. For $n \geq 1$, let $\hat{\omega}_n$ denote the tentative outcome at the beginning of dispute round n . For example, $\hat{\omega}_1$ is the outcome reported by the initial reporter.

Definition 10. For $n \geq 1$ and $\omega \neq \hat{\omega}_n$, let $B(\omega, n)$ denote the amount of stake required to successfully fill a dispute bond in favor of outcome ω during dispute round n .

Recall that the amount of stake required to successfully fill a dispute bond in favor of outcome ω during dispute round n , where $\omega \neq \hat{\omega}_n$ is given by Eq. 1, $B(\omega, n) = 2A_n - 3S(\omega, n)$.

Observation 4. If a dispute bond is successfully filled in favor of outcome ω during dispute round n , then $S(\omega, n+1) = B(\omega, n) + S(\omega, n)$. That is, the successful dispute stake is the only new stake applied to outcome ω at the end of dispute round n .

Observation 5. For all $\omega \neq \hat{\omega}_n$, $S(\omega, n-1) = S(\omega, n)$. That is, if a dispute bond is not entirely filled in favor of outcome ω , then no additional stake is added to outcome ω at the beginning of the next dispute round. This is due to the fact that all unsuccessful dispute stake is returned to the users at the end of the dispute round.

Observation 6. For all $n \geq 2$, $A_n = A_{n-1} + B(\hat{\omega}_n, n-1)$. That is, the total stake over all outcomes at the beginning of a dispute round is simply the total stake from the beginning of the previous dispute round plus the successful dispute stake from the previous dispute round. All other stake is returned to users at the end of the previous dispute round.

Lemma 1. $S(\hat{\omega}_n, n) = 2S(\bar{\hat{\omega}_n}, n)$, for $n \geq 2$.

Proof. Suppose a market enters dispute round n , where $n \geq 2$. During dispute round $n-1$, the outcome $\hat{\omega}_{n-1}$ must have been successfully disputed in favor of outcome $\hat{\omega}_n$. According to Eq. 1, the size of that dispute bond is $B(\hat{\omega}_n, n-1) = 2A_{n-1} - 3S(\hat{\omega}_n, n-1)$. Using observation 3, this can be rewritten as

$$B(\hat{\omega}_n, n-1) + S(\hat{\omega}_n, n-1) = 2S(\bar{\hat{\omega}_n}, n-1) \quad (\text{A1})$$

We know the dispute bond was successfully filled during round $n-1$. Using observation 4, we see that $B(\hat{\omega}_n, n-1) + S(\hat{\omega}_n, n-1) = S(\hat{\omega}_n, n)$. Observation 5 tells us that the total amount staked on $\bar{\hat{\omega}_n}$ is unchanged from round $n-1$ to n , $2S(\bar{\hat{\omega}_n}, n-1) = 2S(\bar{\hat{\omega}_n}, n)$. Thus, Eq. A1 reduces to $S(\hat{\omega}_n, n) = 2S(\bar{\hat{\omega}_n}, n)$. \square

Theorem 2. *Any REP holders successfully disputing an outcome in favor of a market's final outcome will receive a 40% ROI on their dispute stake (measured in REP that exists in a universe that corresponds to the market's final outcome), unless the market is interrupted by some other market causing a fork.*

Proof. During a fork, all users who successfully filled dispute bonds in favor of the forking market's final outcome are given (via coins minted during the fork) a 40% return on their dispute stake when they migrate their dispute stake to the corresponding child universe. Thus, in the case where the market in question has caused a fork, the theorem is immediately true.

Now consider the case where the market in question resolves without causing a fork, and reporting is not interrupted by some other market causing a fork.

Denote the market's final outcome by ω_{Final} and suppose the market resolves at the end of dispute round n , where $n \geq 2$. That means the tentative outcome for round n is ω_{Final} , and that outcome is not successfully disputed during round n . In other words: $\hat{\omega}_n = \omega_{\text{Final}}$. Then by Lemma 1 we know that $S(\omega_{\text{Final}}, n) = 2S(\bar{\omega_{\text{Final}}}, n)$.

Since the market resolves at the end of round n with no further stake added to any outcome, the above equation shows the final amount of stake on the market's final outcome, ω_{Final} , and the sum of all stake on all of the market's other outcomes, $\bar{\omega_{\text{Final}}}$. Note that there is exactly twice as much stake on the market's final outcome as there is on all other outcomes combined.

Augur burns 20% of the all stake on the non-final outcomes and redistributes the rest to users who staked on ω_{Final} , in proportion to the amount of REP they staked. Therefore the users who successfully filled a dispute bond in favor of ω_{Final} get a 40% ROI on their staked REP. \square

Next, consider the maximum number of dispute rounds required to resolve a market. Eq. 1 is minimized when ω is chosen to be the non-tentative outcome that begins the dispute round with the greatest amount of stake. Lemma

1 implies that the non-tentative outcome with the greatest amount of stake is the previous dispute round's tentative outcome. Therefore, the smallest possible dispute bond size that can be successfully filled during dispute round n , where $n \geq 2$, is $B(\hat{\omega}_{n-1}, n)$.

In other words, the dispute bond size grows *slowest* when the same two outcomes are repeatedly disputed in favor of one another. It follows that the number of dispute rounds required for a market to initiate a fork is *maximized* when the same two outcomes are repeatedly disputed in favor of one another. Therefore we can determine the maximum number of dispute rounds that any market may undergo before initiating a fork by finding the maximum number of dispute rounds that can occur in the particular case where the same two market outcomes are repeatedly disputed in favor of one another. We examine that case now.

Suppose that every successful dispute bond is filled in favor of the previous dispute round's tentative outcome. Then the two tentative outcomes that are iteratively disputed in favor of one another are $\hat{\omega}_1$ and $\hat{\omega}_2$.

Observation 7. In the case where the same two tentative outcomes are repeatedly disputed in favor of one another, $\hat{\omega}_n = \hat{\omega}_{n-2}$ for all $n \geq 3$.

Definition 11. Let d denote the amount of stake placed on $\hat{\omega}_1$ during the initial report. Because the tentative outcome for each round is known in this situation, we can simplify our notation for the dispute bond sizes. Define a shorthand B_n to denote the bond size required for round n , so that $B_1 = 2d$ and $B_n = B(\hat{\omega}_{n-1}, n)$ for all $n \geq 2$. This will make for easier reading and comprehension.

Observation 8. In the case where the same two tentative outcomes are repeatedly disputed in favor of one another, $S(\hat{\omega}_{n-1}, n) = S(\hat{\omega}_{n-1}, n-2) + B_{n-2}$ for $n \geq 3$. (That is, every other successful dispute bond is added to the same outcome.)

Lemma 2. If the same two tentative outcomes are repeatedly disputed in favor of one another, then for all n where $n \geq 3$:

1. $S(\hat{\omega}_{n-1}, n) = \frac{2}{3}B_{n-1}$
2. $A_n = 2B_{n-1}$ and
3. $B_n = 3d2^{n-2}$

Proof. (By induction on n)

Suppose the same two tentative outcomes are repeatedly disputed in favor of one another.

(Base Case) By definition and Eq. 1 we make the following observations.

- $S(\hat{\omega}_1, 1) = d$, $S(\hat{\omega}_2, 1) = 0$, $A_1 = d$, and $B_1 = 2d$
- $S(\hat{\omega}_1, 2) = d$, $S(\hat{\omega}_2, 2) = 2d$, $A_2 = 3d$, and $B_2 = 3d$
- $S(\hat{\omega}_1, 3) = 4d$, $S(\hat{\omega}_2, 3) = 2d$, $A_3 = 6d$, and $B_3 = 6d$

$S(\hat{\omega}_{3-1}, 3) = S(\hat{\omega}_2, 3) = 2d = \frac{2}{3}(3d) = \frac{2}{3}B_2 = \frac{2}{3}B_{3-1}$, so part 1 of the lemma holds for $n = 3$.

$A_3 = 6d = 2(3d) = 2B_2 = 2B_{3-1}$, so part 2 of the lemma holds for $n = 3$.

$B_3 = 6d = 3d2^{3-2}$, so part 3 of the lemma holds for $n = 3$.

Therefore the lemma, in its entirety, holds true for the base case of $n = 3$.

(Induction) Suppose the lemma is true for all n such that $3 \leq n \leq k$. We want to show that the lemma holds for $n = k + 1$. That is, we want to show that:

- (a) $S(\hat{\omega}_k, k + 1) = \frac{2}{3}B_k$
- (b) $A_{k+1} = 2B_k$ and
- (c) $B_{k+1} = 3d2^{k-1}$

First, we prove part (a). By observation 8:

$$S(\hat{\omega}_k, k + 1) = S(\hat{\omega}_k, k - 1) + B_{k-1}$$

By observation 7 we can rewrite the above as:

$$S(\hat{\omega}_{k-2}, k + 1) = S(\hat{\omega}_{k-2}, k - 1) + B_{k-1}$$

By the induction hypothesis, we can rewrite $S(\hat{\omega}_{k-2}, k - 1)$ as $\frac{2}{3}B_{k-2}$ on the right-hand side to get:

$$S(\hat{\omega}_{k-2}, k + 1) = \frac{2}{3}B_{k-2} + B_{k-1}$$

By the induction hypothesis, we can write B_{k-2} as $3d2^{k-4}$ and B_{k-1} as $3d2^{k-3}$:

$$S(\hat{\omega}_{k-2}, k + 1) = d2^{k-1}$$

Applying observation 7 to the left-hand side we get:

$$S(\hat{\omega}_k, k + 1) = d2^{k-1}$$

Finally, note that by the above equation and the induction hypothesis, $S(\hat{\omega}_k, k + 1) = d2^{k-1} = \frac{2}{3}(3d2^{k-2}) = \frac{2}{3}B_k$. This proves part (a).

Next, we prove part (b). By observation 6:

$$A_{k+1} = A_k + B_k$$

By the induction hypothesis, $A_k = 2B_{k-1}$:

$$A_{k+1} = 2B_{k-1} + B_k$$

By the induction hypothesis, $B_{k-1} = 3d2^{k-3}$, so the right-hand side can be simplified to

$$A_{k+1} = 3d2^{k-2} + B_k$$

By the induction hypothesis, $B_k = 3d2^{k-2}$ to rewrite the right-hand side as

$$A_{k+1} = 2B_k,$$

and part (b) is proved.

Finally, we prove part (c). By Eq. 1:

$$B_{k+1} = 2A_{k+1} - 3S(\hat{\omega}_k, k+1)$$

By observation 8, we can write $S(\hat{\omega}_k, k+1)$ as $S(\hat{\omega}_k, k-1) + B_{k-1}$:

$$B_{k+1} = 2A_{k+1} - 3(S(\hat{\omega}_k, k-1) + B_{k-1})$$

By observation 7, $\hat{\omega}_k = \hat{\omega}_{k-2}$:

$$B_{k+1} = 2A_{k+1} - 3(S(\hat{\omega}_{k-2}, k-1) + B_{k-1})$$

By observation 6, $A_{k+1} = A_k + B_k$:

$$B_{k+1} = 2(A_k + B_k) - 3(S(\hat{\omega}_{k-2}, k-1) + B_{k-1})$$

By the induction hypothesis, $A_k = 2B_{k-1}$ and $S(\hat{\omega}_{k-2}, k-1) = \frac{2}{3}B_{k-2}$:

$$B_{k+1} = 2(2B_{k-1} + B_k) - 3(\frac{2}{3}B_{k-2} + B_{k-1})$$

By the induction hypothesis, $B_k = 3d2^{k-2}$, $B_{k-1} = 3d2^{k-3}$ and $B_{k-2} = 3d2^{k-4}$. Making these substitutions and simplifying yields:

$$B_{k+1} = 3d2^{k-1}$$

This proves part (c), and concludes the proof of the lemma. \square

Theorem 3. *If not interrupted by some other market causing a fork, a given market may undergo at most 20 dispute rounds before finalizing or causing a fork.*

Proof. Suppose that a given market is not interrupted by some other market causing a fork. Then, as shown above, we know that the number of dispute rounds required for a market to initiate a fork is maximized when the same two outcomes are repeatedly disputed in favor of one another. Part 3 of Lemma 2 tells us that, in this situation, the dispute bond size required for successfully disputing the tentative outcome during round n is given by $3d2^{n-2}$, where d is the amount of the stake placed during the initial report.

We know that forks are initiated after the successful fulfillment of a dispute bond with size at least 2.5% of all existing REP, and we know that there are 11 million REP in existence. Thus a fork is initiated when a dispute bond of size 275,000 REP is filled. We also know that $d \geq 0.35$ REP, because the minimum amount of stake on the initial report is 0.35 REP²⁸.

Solving $3(0.35)2^{n-2} > 275,000$ for $n \in \mathbb{Z}$ yields $n \geq 20$. Thus, we can guarantee that a market will resolve or cause a fork after at most 20 dispute rounds. \square

Appendix B: Bond Size Adjustments

The validity bond and the creation bond are dynamically adjusted based on the behavior of participants during the previous dispute window. The creation bond is the maximum of the *no-show bond* and the *designated reporter bond* – two values that the system tracks in order to compute the creation bond but does not expose to users.

Here we describe how we adjust these bonds.

We define the function $f : [0, 1] \rightarrow [\frac{1}{2}, 2]$ by:²⁹

$$f(x) = \begin{cases} \frac{100}{99}x + \frac{98}{99} & \text{for } x > \frac{1}{100} \\ 50x + \frac{1}{2} & \text{for } x \leq \frac{1}{100} \end{cases} \quad (\text{B1})$$

The function f is used to determine the multiple used in these adjustments, as described in the subsections below. In brief, if the undesirable behavior occurred exactly 1% of the time during the previous dispute window, then the bond size remains the same. If it was less frequent, then the bond size will be reduced by as much as half. If it was more frequent, then the bond size will be increased by as much as a factor of 2.

1. Validity Bond

During the very first dispute window after launch, the validity bond will be set at 0.01 ETH. Then, if more than 1% of the finalized markets in the previous dispute window were invalid, the validity bond will be increased. If less than 1% of the finalized markets in the previous dispute window were invalid, then the validity bond will be decreased (but will never be lower than 0.01 ETH).

In particular, we let ν be the proportion of finalized markets in the previous dispute window that were invalid, and b_v be the amount of the validity bond from the previous dispute window. Then the validity bond for the current window is $\max\{\frac{1}{100}, b_v f(\nu)\}$.

2. No-Show Bond

During the very first dispute window after launch, the no-show bond will be set at 0.35 REP. As with the validity bond, the no-show bond is adjusted up or down, targeting a 1% no-show rate with a floor of 0.35 REP.

Specifically, we let ρ be the proportion of markets in the previous dispute window whose designated reporters failed to report on time, and we let b_r be amount of the no-show bond from the previous dispute window. The the amount of the no-show bond for the current dispute window is $\max\{0.35, b_r f(\rho)\}$.

²⁸See appendixes B2 and B3

²⁹This formula may change once empirical data from live markets is obtained.

3. Designated Reporter Bond

During the very first dispute window after launch, the amount of the designated reporter bond will be set at 0.35 REP. The amount of the designated reporter bond is dynamically adjusted according to how many designated reports were incorrect (failed to concur with the final

market outcome) during the previous dispute window.

In particular, we let δ be the proportion of designated reports that were incorrect during the previous dispute window, and we let b_d be the amount of the designated reporter stake during the previous dispute window, then the amount of the designated reporter bond for the current window is $\max\{0.35, b_d f(\delta)\}$.