Vega Protocol

A liquidity incentivising decentralised trading protocol for smart financial products

G. Danezis
george@vega.xyz

D. Hrycyszyn
dave@vega.xyz

B. Mannerings
barney@vega.xyz

T. Rudolph
tamlyn@vega.xyz

D. Šiška
david@vega.xyz

3rd June 2019

First published: 4th September 2018

Abstract

The majority of existing financial markets exhibit several key problems: the need to trust third parties, high costs driven by rent-seeking intermediaries, and the presence of organisations and structures that act as gatekeepers and censors, control the availability of products, and stifle innovation.

Vega is a protocol for the decentralised trading and execution of financial products. It is designed for fully automated, end-to-end margin trading on open public networks, secured with proof of stake. We outline a novel incentivisation scheme which leverages a dynamic liquidity marketplace to solve the problem of attracting and allocating market making resources. Permissionless innovation is enabled by *Smart Products*, which allow anyone to create products and propose new markets, and works in tandem with decentralised risk management to enable the safe trading of arbitrarily complex instruments.

Products on a Vega network can reference practically any underlying price or other data feed, allowing participants to define and trade a wide range of instruments across the full spectrum of global markets. Cross-chain settlement means that the protocol is blockchain agnostic and allows trades to settle in any crypto-asset residing on a supported chain, paving the way for physically settled and cash settled products, as commodity and asset tokenisation become widespread.

We believe that this technology can be transformational for the financial system, changing the dynamics of power and forming part of a wave of change that could radically alter the operation of markets and their relationship with society.

THIS WHITEPAPER PROVIDES AN INITIAL SUMMARY OF CERTAIN TECHNICAL AND BUSINESS ESSENTIALS UNDERLYING THE VEGA PROTOCOL. THIS DOCUMENT IS EXPECTED TO EVOLVE OVER TIME, AS THE PROJECT PROCEEDS. THE VEGA TEAM MAY POST MODIFICATIONS, REVISIONS AND/OR UPDATED DRAFTS FROM TIME TO TIME, INCLUDING BEFORE, DURING, AND AFTER THE CREATION OF ANY TOKENS, AND WHILST NETWORK(S) BASED ON THE VEGA PROTOCOL ('VEGA NETWORKS') ARE IN OPERATION.

THIS DOCUMENT SETS FORTH A DESCRIPTION OF THE VEGA PROTOCOL, REFERENCE SOFTWARE IMPLEMENTATION, AND POTENTIAL VEGA NETWORKS. THIS INCLUDES DESCRIPTIONS OF THE PROTOCOL ITSELF, 'SMART PRODUCTS', AND THE USE OF TOKENS SUCH AS THE PROPOSED VEGA TOKEN. THIS DOCUMENT IS PROVIDED FOR INFORMATION PURPOSES ONLY AND IS NOT A BINDING LEGAL AGREEMENT. ANY SALE OR OTHER OFFERING OF VEGA TOKENS WOULD BE GOVERNED BY SEPARATE TERMS & CONDITIONS. IN THE EVENT OF CONFLICT BETWEEN APPLICABLE TERMS & CONDITIONS AND THIS DOCUMENT, THE TERMS & CONDITIONS GOVERN.

THIS WHITEPAPER IS NOT AN OFFERING DOCUMENT OR PROSPECTUS, AND IS NOT INTENDED TO PROVIDE THE BASIS OF ANY INVESTMENT DECISION OR CONTRACT.

Legal disclaimer

As of the date of publication, the Vega team have no plans to launch any public Vega Networks, and Vega Tokens are a proposed token with no known potential uses outside of Vega Networks, and no such use is intended. This document does not constitute advice nor a recommendation by the Vega team, its officers, directors, managers, employees, agents, advisers or consultants, or any other person to any recipient of this document on the merits of purchasing, otherwise acquiring, or holding Vega Tokens or any other cryptocurrency or token. The purchase and holding of cryptocurrencies and tokens carries substantial risks and may involve special risks that could lead to a loss of all or a substantial portion of any money invested. Do not purchase tokens unless you are prepared to lose the entire amount allocated to the purchase.

Vega Tokens, if and when they are created and made available, should not be acquired for speculative or investment purposes with the expectation of making a profit or immediate re-sale. They should be acquired only if you fully understand the intended functionality of the Vega Tokens, and you intend to use the Vega Tokens for those purposes only, and it is legal for you to do so. No promises of future utility or performance or value are or will be made with respect to Vega Tokens, including no promise any Vega Networks will be launched, no promise of inherent value, no promise of any payments, and no guarantee that Vega Tokens will hold any particular value.

Vega Tokens are not designed and will not be structured or sold as securities. Vega Tokens will hold no rights and confer no interests in the equity of the Vega business or any future Vega Networks. Vega Tokens are designed and intended for future use on public Vega Networks that may be created using the Vega protocol, for the purposes of trading and governance transactions, or for the operation of a node. Proceeds of any sale of Vega Tokens may be spent freely by Vega for any purpose, including but not limited to the development of its business and underlying technological infrastructure, absent any conditions set out in this document.

This whitepaper is not a prospectus or disclosure document and is not an offer to sell, nor the solicitation of any offer to buy any investment or financial instrument or other product in any jurisdiction and should not be treated or relied upon as one. Any distribution of this whitepaper must be of the complete document including the cover page and this disclaimer and the accompanying boilerplate in their entirety.

All information in this document that is forward looking is speculative in nature and may change in response to numerous outside forces, including technological innovations, regulatory factors, and/or currency fluctuations, including but not limited to the market value of cryptocurrencies.

This whitepaper is for information purposes only and will be subject to change. The Vega team cannot guarantee the accuracy of the statements made or conclusions reached in this whitepaper. The Vega team does not make and expressly disclaims all representations and warranties (whether express or implied by statute or otherwise) whatsoever, including but not limited to: any representations or warranties relating to merchantability, fitness for a particular purpose, suitability, wage, title or non-infringement; that the contents of this document are accurate and free from any errors; and that such contents do not infringe any third party rights.

The Vega business, Vega team, and operators of any Vega Networks shall have no liability for damages of any kind arising out of the use, reference to or reliance on the contents of this whitepaper, even if advised of the possibility of such damages arising.

This whitepaper includes references to third party data and industry publications. The Vega team believes that the information reproduced in this whitepaper is accurate and that the estimates and assumptions contained herein are reasonable. However, there are no assurances as to the accuracy or completeness of this data. The information from third party sources contained herein has been obtained from sources believed to be reliable; however, there are no assurances as to the accuracy or completeness of any included information. Although the data is believed to be reliable, the Vega team has not independently verified any of the information or data from third party sources referred to in this whitepaper or ascertained the underlying assumptions relied upon by such sources.

Please note that Vega is in the process of undertaking a legal and regulatory analysis of the functionality of the protocol, proposed Vega Tokens, and the operation of its business. Following the conclusion of this analysis, the Vega team may decide to amend the intended functionality of Vega Tokens in order to ensure compliance with any legal or regulatory requirements to which it is subject, which may affect the utility, fungibility, or any other properties of Vega Tokens.

Any Vega Tokens could be impacted by regulatory action, including potential restrictions on the ownership, use, or possession of such tokens. Regulators or other competent authorities may demand that the mechanics of the Vega Tokens be altered, entirely or in part. Vega may revise the Vega protocol or Vega Token mechanics to comply with regulatory requirements or other governmental or business obligations. Nevertheless, Vega believes it has taken all commercially reasonable steps to ensure that the design of Vega Tokens is proper and in compliance with currently considered regulations as far as reasonably possible.

No regulatory authority has examined or approved any of the information set out in this whitepaper. The publication, distribution or dissemination of this whitepaper does not imply compliance with applicable laws or regulatory requirements.

Contents

1 Introduction

Vega is a technology protocol and associated crypto-asset [28] for an open, blockchain-backed public network for fully automated end-to-end trading and execution of financial products. The network is secured with *proof of stake* and implements pseudonymous margin trading using a novel liquidity incentivisation scheme based on market forces to solve the problem of attracting and allocating market making resources in a decentralised system.

Permissionless innovation is enabled by *smart products* which allow anyone to create products and propose new markets. This works in tandem with a decentralised margin system using a suite of *risk models* based on *coherent risk measures* [5] to enable the safe trading of arbitrarily complex *instruments* in an environment with zero expected recovery in the event of default.

Products can reference practically any *underlying* price or other data feed, allowing participants to define and trade a wide range of instruments across the full spectrum of global markets. Cross-chain settlement means that the protocol is blockchain agnostic and allows trades to settle in any crypto-asset residing on a supported chain, paving the way for physically settled[1] in addition to cash settled products, as commodity and asset tokenisation become widespread.

This paper describes a protocol that defines how traders, *market makers*, and node operators interact to collectively run high performance, fully decentralised markets in a deterministic way without the need for human intervention. This includes a robust market framework which manages the network, markets, and participants, and provides a strong foundation for the functionality of Vega. The protocol covers end-to-end trading including price determination, margining, and settlement, as well as collateral management and on-chain market governance.

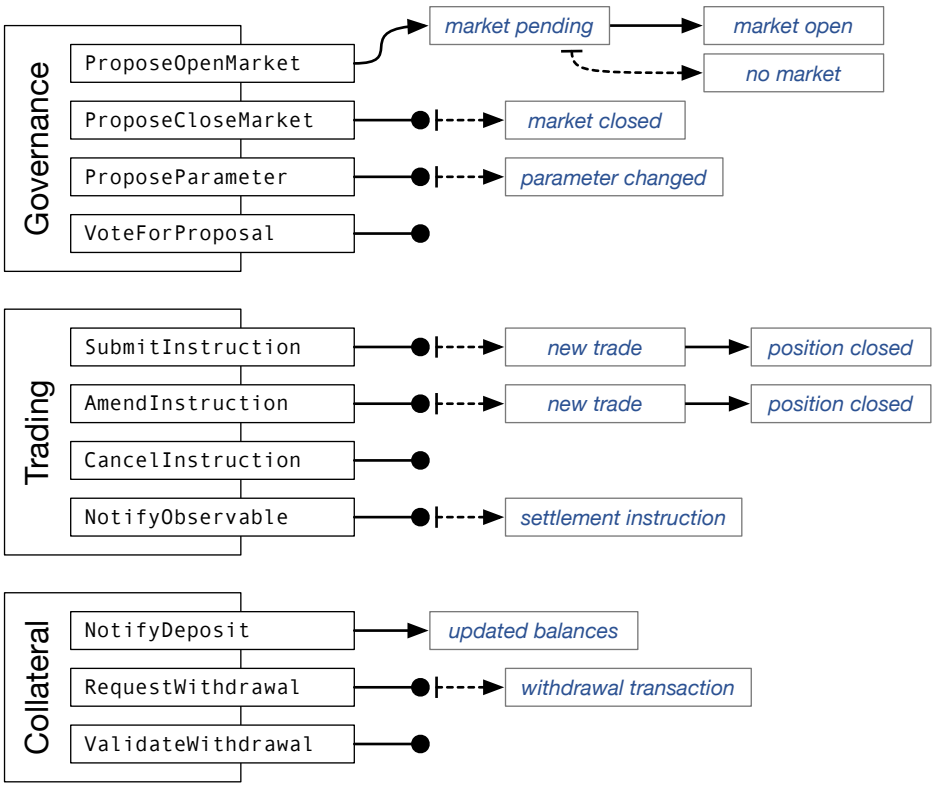

Figure 1: Map of core protocol transactions

[1] Via settlement of an appropriate crypto-asset that represents a claim on the real asset.

1.1 Overview of core Vega concepts

Vega opens and decentralises markets by fully automating the processes and incentives for trading and settling financial products between pseudonymous participants. This requires carefully designed mechanisms of economic rewards and penalties, and a protocol that balances the desire for permissionless innovation with the need to protect markets and participants. Such a system is described in brief below, and more fully in the rest of this paper.

Instruments are defined by a combination of a product, *risk model*, and their required parameters and are traded either in *open markets* or *over the counter* (OTC), depending on the type of *instrument* and level of market making support. *Open markets* are created by *market makers*, and match trades between any willing and sufficiently collateralised participants, whereas *OTC* trading occurs on a more ad-hoc basis. *Trading modes* include *continuous trading* on a limit order book, *discrete trading* via frequent batch auctions [19] and *request for quote* (RFQ).

Smart products are a special type of smart contract designed to allow the creation of a wide range of financial products from a toolkit of standard features and economic primitives. They are written at a higher level of abstraction than most smart contracts, making them easier to develop and test, and better suited to static analysis and automated risk modelling. The *smart product language* (the details of which will be the subject of a separate paper) used to create new products is designed such that settlement instructions, risk calculations, and other useful outputs for trading[2] can be derived from a single *smart product* definition.

Collateral is managed by Vega networks via links to other blockchains[3], with funds deposited by paying in to a smart contract on the 'host' chain. Collateral can be used as margin for orders and positions, meaning the required funds will be *allocated* to a market until they are no longer needed and are *released*. Allocated funds can't be withdrawn or used for trading in other markets. Withdrawals are requested with a Vega transaction which, if successful, results in a transaction for the host chain that has been signed by a quorum of Vega nodes, and will cause the funds to be released to the requested destination address.

Trading and settlement are designed to be fair and predictable for all participants. Trading on *open markets* will use a defined *trading mode* unless market conditions dictate otherwise, for example when an auction period is used to identify the fair price after large moves, or when protective measures are required due to low liquidity. Positions are settled continuously as they are closed, when products generate interim cashflows, and finally at expiry of the instrument, when collateral held in margin is also released. We define a *position resolution* algorithm to fairly handle situations where there is a shortfall in the available collateral and no remaining funds in the market's *insurance pool*.

Risk management is of particular importance for pseudonymous trading, as there is no practical recourse in the event that a participant owes more than they have, or is allowed to withdraw more than is rightfully theirs. To mitigate this, trading is margined, with *risk models* that have been selected and calibrated for a zero recovery rate environment. Margin requirements take into account the slippage incurred when closing a position, and positions that present an unacceptably high risk of loss to the network are closed automatically. The rules are designed so that on average, closeouts will occur with a net positive margin remaining allocated to the position. This is added to an *insurance pool* that is used to cover the difference when a closeout leaves a negative balance. This mechanic ensures that most markets, and the network as a whole — *insurance pool* balances are redistributed to other markets at expiry — will become safer over time.

Liquidity provision is incentivised through the protocol rather than as an offline activity. Liquidity rewards are distributed to the price makers of a trade, to the market makers of a market and to the proof of stake token holders who are supporting the infrastructure. Price takers incur a fee at the point of trade. This fee represents a cost for accessing liquidity and the level is dynamically calculated according to how valuable that liquidity is to the market. Since liquidity providers may decide where to provide liquidity on Vega, this model effectively operates as a marketplace for liquidity, with the ultimate goal being to minimise cost per trade by efficiently

[2] For example, legal contracts, pricing models and human readable descriptions.
[3] The protocol does not create a cryptocurrency. Trades are settled in existing coins and tokens.

allocating market making resources. Key to this approach is the role of *market makers* as owner-operators of markets, which involves committing to provide order book depth in return for a share of the liquidity value realised by participants, secured by a financial bond deposited with the network.

Market governance is necessary to ensure that the network can operate and grow unencumbered and without manual intervention whilst minimising the risk posed by bad actors. Vega's market governance features are designed around the concept of stake weighted voting, with various actions such as the creation and closure of markets, and the setting of parameters that influence their behaviour being the main focus of on-chain governance. Off-chain governance around the protocol's ongoing development and the reference implementation are out of scope for this paper and will be covered in more detail in later work.

1.2 Architecture of a public network

The Vega protocol is designed to be implemented in a distributed and decentralised manner on a network of nodes that may be the same or distinct from trading parties participating in markets. Nodes will maintain a mirror of the state of their Vega network, and process transactions to operate markets and their governance. Nodes are included in the infrastructure through a proof-of-stake mechanism: a certain stake is locked by a node, and as a surety they will operate correctly. These infrastructure nodes jointly run a byzantine consensus protocol [6], that ensures all honest parties sequence operations consistently, and thus feed the protocol implementation with actions in the same order across the network. Our reference implementation currently uses[4] the Tendermint distributed smart contracts platform for the consensus and the proof-of-stake protocol.

Clients may connect to any infrastructure node and send orders for any available market, perform market actions, and participate in the governance of the network or markets. The current reference node implementation includes a REST API for light clients to be able to access the platform, a GraphQL API for web applications, and a native (GRPC) API to interact with the infrastructure, as well as a reference HTML based decentralised trading application that can connect to a local or remote node. Clients, as anyone, may participate in the network as a full infrastructure node — and we expect institutional actors, including market makers to use this option.

Once all infrastructure nodes have sequenced actions consistently, those actions are processed by software implementing the Vega protocol. Since all nodes execute the actions in the same order, and the protocol is deterministic, the network as a whole arrives at the same mirrored state. Thus any orders, trades, collateral calculations, or governance decisions will be the same across all nodes.

We also define a bridge between other platforms carrying cryptocurrencies and tokens of value, and the Vega protocol, in order to trade assets native to those platforms. The first integration is with the Ethereum blockchain: an Ethereum smart contract allows users to send assets, either Ether or ERC20 tokens, to the Vega platform – effectively locking them into a holding account while they are available as collateral. Vega nodes read the Ethereum ledger and transform any locked assets into collateral that can be used for trading on relevant markets. Once trades have settled, clients may choose to withdraw their assets, by inserting into the Vega ledger a withdrawal request transaction. The Ethereum smart contract interprets those transactions, then unlocks and transfers the assets to the receiving Ethereum wallet.

Simpler read-only bridges are also established between Vega and *oracles* providing data feeds to power certain markets. Those *oracles* are authenticated, and the feeds of data signed to ensure their integrity and authenticity. A market may combine data from multiple similar *oracles* to mitigate the impact on the market if any given *oracle* fails or is malicious.

To secure the Vega platform we use a mixture of modern cryptography and distributed systems security mechanisms. All user assets are associated with the verification key of a digital signature algorithm (public key). Any actions relating to those assets, such as transferring them,

[4]Note that we do not run a public network on 'production' blockchains. However, the Vega protocol and reference node software implementation are designed to be capable of supporting such a network.

placing orders, amending or cancelling orders, or governance decisions will carry a valid digital signature from the originating user to be considered valid; unless it is determined automatically by the protocol rules, e.g. when closing overexposed positions. The properties of modern Byzantine consensus algorithms guarantee safety —- namely that all infrastructure nodes see the same sequence of client actions — and liveness —- namely that the system eventually will perform all client actions. The Tendermint library, which we currently use, guarantees this subject to the set of nodes with two-thirds stake being honest [18].

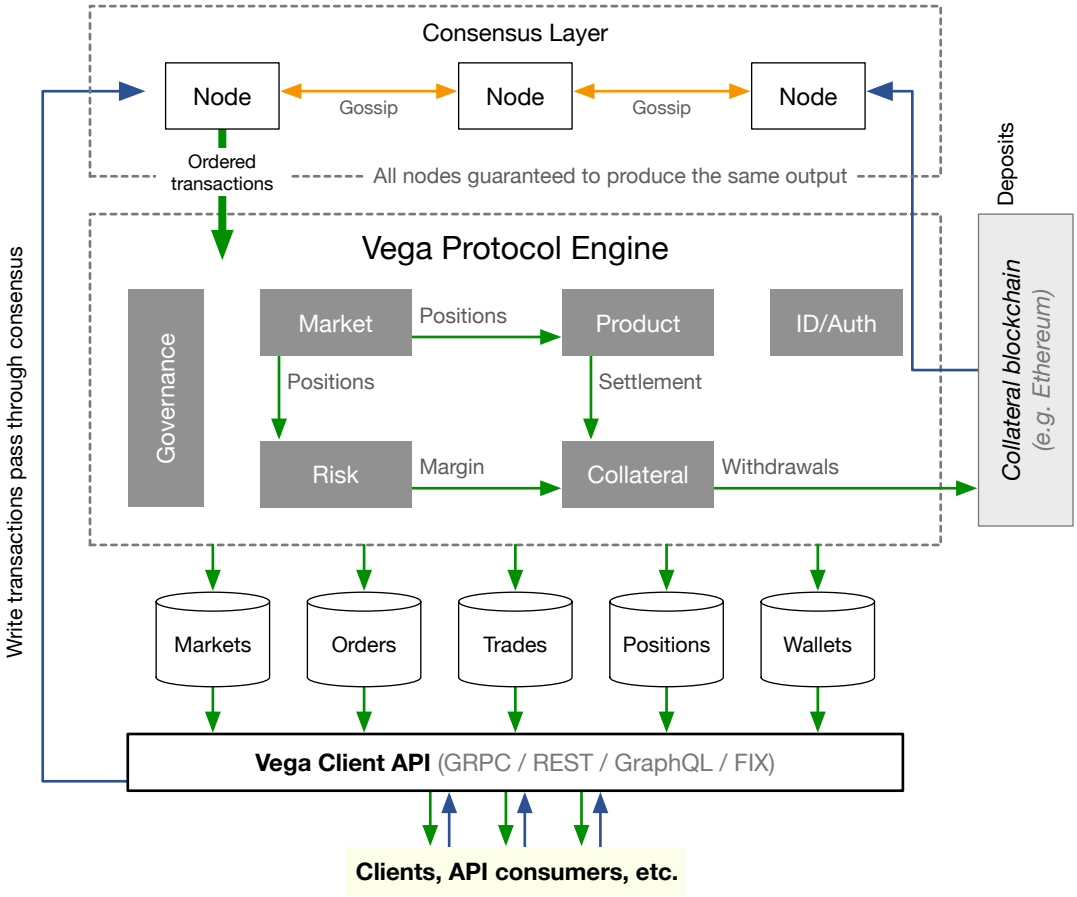

Figure 2: Logical architecture of a public Vega network

2 Background & motivation

Traditionally, financial products are created and traded in markets consisting of various organisations (and sometimes individuals) connected by technology systems and contractual obligations that simultaneously facilitate trading and create barriers to entry. These vary in complexity, cost, and sophistication from extremely slow, manual, and error prone paper-based operations to high speed, high liquidity electronic markets.

Regardless of the precise methods and systems in use, these markets all exhibit several key problems. They include the need to trust third parties, rent-seeking intermediaries that drive up costs, and the presence of organisations and structures that act as gatekeepers and censors for access to existing markets, and equally importantly, controlling the availability of products and creation of markets themselves.

A decentralised platform for financial products — in which no individual node or party being compromised[5] would pose a risk to the ongoing operation of the network or availability of markets — could provide a genuine alternative, allowing everyone to trade on a more equal basis. Access would be available to all, and the creation of markets would no longer be dependent on central parties or organisations.

2.1 Desirable properties of a decentralised trading protocol

In considering our design approach for Vega, we find ourselves examining existing centralised and decentralised solutions for the trading and execution of financial products in markets of varying size and maturity. From these, we have identified nine key properties, enumerated below, that describe a good trading platform, a good decentralised system, or both.

i) **Fair and trusted:** To be credible, a market must operate in a way that is predictable, comprehensible, and fair. Participants will not place their money on the line for a trading platform that they don't trust to work as expected. Most markets hold the concepts of trust and fairness in high regard, however, for a decentralised solution this is even more important, as we are proposing a system that would widen access greatly. We must therefore demonstrate that our proposal treats all potential participants fairly and predictably, even in a variety of potentially hostile conditions and under the constraints of decentralisation.

ii) **Liquid:** It is important for any trading venue that sufficient liquidity is available to sustain and grow each market. This is best achieved through different mechanisms depending on the maturity and size of each market:

- new markets will initially have no natural traders, and so require a commitment[6] from *market makers* to provide liquidity on first opening to kickstart their growth;
- relatively illiquid markets will not have a constant flow of traders and thus need strong incentives to attract liquidity providers and catalyse growth; and
- mature, liquid markets will be of interest to *market makers* based on volume and therefore benefit most from the trading activity attracted by low fees.

iii) **Low latency:** In general, the faster a market can react, the more accurately the price reflects all available information. Whilst there are limits to the social utility of ever decreasing latency, it is important that any solution operates with low enough latency that traders are not disadvantaged compared to other trading venues.

iv) **High throughput:** To provide a fair trading venue with wide appeal, markets must support the trading needs of as many participants as possible, as a system that is too easily congested will either need to randomly exclude participants or create an undesirably high cost barrier to

[5] Whether due to malicious actors, bankruptcy, or as a consequence of some other unanticipated situation.
[6] Traditionally provided or sourced by the market operator in centralised markets.

trading[7]. Furthermore, it is important to minimise the chance of external network conditions disrupting trading activities, as exemplified by the network slowdowns on the Ethereum blockchain caused by popular token sales and trending 'dApps'.

v) **Scalable:** Related to throughput is scalability. It is important that the system supports the needs of an individual market, and that there is also a clear route to supporting many hundreds or thousands of markets. Of particular practical interest is ensuring that the underlying protocol design (regardless of any constraints that may exist for a given implementation) doesn't rely on all markets sharing global state or resources, for instance by requiring that all transactions write to a single blockchain.

vi) **Flexible and widely applicable:** This work is driven by the desire to create a common, end-to-end infrastructure for financial products. The usefulness of such a system is heavily dependent on its ability to effectively support the current and future needs of the many industries, organisations, and societies that interact with the global financial system. As such, we must design a platform that allows for the permissionless creation of markets in bespoke and arbitrarily complex instruments, and for their ongoing evolution.

vii) **Trust-minimising:** In order to effectively deliver on the benefits of decentralisation, it's important to minimise the need for participants to identify — other than pseudonymously — or trust each other, which creates requirements to:

- protect the traders and markets from unrecoverable defaults;
- prevent market creators and liquidity providers from using their position to gain an unfair advantage against other participants; and
- ensure participants have no overall net incentive for malicious behaviour.

viii) **Self governing:** As previously discussed, we posit a pseudonymous and permissionless global trading network, which poses a number of risks of the kind usually mitigated with strong centralised governance. These include but are not limited to the potential for fraudulent instruments, fragmentation of liquidity, and reputational damage from unethical markets. To mitigate this, a strong system of decentralised governance is required that addresses these risks without jeopardising the permissionless nature of the overall system.

ix) **Product and market independent:** Given the global nature of a decentralised network and the existence of fundamental regional incompatibilities in terms of financial markets legislation, it is important that any solution is able to act as a neutral infrastructure layer that does not itself offer products nor operate markets, with responsibility for legal compliance falling on the participants in each jurisdiction who create products, make markets, and trade.

2.2 Real world uses and adoption

By definition, a permissionless platform places the development of new products and the launching of markets in the hands of participants, so it is not possible to predict which instruments would be made available or when, however we have identified below a small subset of use cases with clear potential benefits to give an idea of the potential of this technology.

Traders in the blockchain and cryptocurrency sector form an obvious starting point, given their technical familiarity with the concepts underpinning Vega. The existence of an active and informed trading community with significant (and growing) trade volume, and increasing sophistication and institutional interest means there is already demonstrable demand for a decentralised platform for trading in various crypto-asset derivatives[8].

Where the cost of creating products and operating markets has previously been prohibitive, and in situations where barriers to access exist, there is clear demand for innovation that is not met by the current system. For example:

[7] This could be implicit, through the need to acquire computational resources or network connectivity to achieve access, or explicit by including only the transactions offering the highest fee.

[8] Examples include options, futures and other derivatives on various crypto-assets, particularly where these will allow traders to take a market view that has previously been difficult to trade.

- relevant, locally targeted insurance products, for instance providing unbanked and underbanked agriculture with compensation in the event of catastrophic weather;
- wider access and reduced cost for common hedging strategies, such as offsetting FX risk for small business owners with overseas cost and/or revenue; and
- sovereign credit default swaps (CDS) settled in cryptocurrency, mitigating the risk of fiat devaluation in the event of a credit event on a fiat settled sovereign CDS.

As the wider crypto-economic ecosystem develops[9], we see the opportunity for such a decentralised and highly automated system to replace costly, slow, and error prone processes governing many types of business transactions and agreements such as:

- asset backed loans, with the tokenised asset title automatically transferred in the event of default by the borrower;
- 'delivery vs payment' for industrial and bulk goods or services, with payment tied to delivery proof and the ability to apply penalties for late delivery; and
- replication of many existing core banking and insurance products, providing lower execution costs, better price discovery, and increased transparency.

Finally, a decentralised financial products platform could eventually attract trading away from markets that are already cheap, liquid, and generally accessible such as FX, equities, and others. However, we see the migration of this trading activity as a much longer term outcome predicated on successful adoption in other markets along with wider industry trends including automation, cost reduction, and increasing adoption of blockchain technology in general.

2.3 Summary

This work is re-imagining the functionality of key economic mechanisms and institutions in light of recent innovations in decentralisation. We have identified a clear and well defined opportunity to augment the decentralised financial system, taking it beyond cryptocurrencies, tokenised assets, and relatively slow general purpose smart contracts by adding a high performance financial products trading and execution layer.

The platform we imagine would open the door to a fairer and more accessible financial system that doesn't favour some groups, nor exclude others. We see the potential of permissionless global markets to not only improve on things that exist today, but perhaps to fundamentally rethink how we harness the tools of economics to drive better outcomes for everyone [27].

To conclude: while the proposed platform is only a small piece of the puzzle, we believe that as a response to centralisation and inequality in financial markets, and an enabler for experimentation around our future relationship with money, it may be an important one, and as such are excited to introduce the Vega protocol.

[9] For instance, with increasing availability and maturity of services like physical asset tokenisation, blockchain based identity, algorithmic or asset-backed 'stablecoins', and oracle services.

3 Market framework

To satisfy the need for a flexible, self governing system (see Section 2.1 vi, viii), Vega must provide a standardised framework for creating and interacting with markets that is both rigid enough to provide certainty to users, and flexibly designed, so as to allow for the future expansion of market and product types, and trading modes. Additionally, it's important to keep this core of the protocol as simple and predictable as possible, to ensure that implementations are testable and the behaviour of the various transactions can be reliably verified by any observer. This is achieved with a hierarchical, parametric framework for describing markets (see Figure 3), combined with the rules (described throughout this document) that specify how the components interact and transactions are processed.

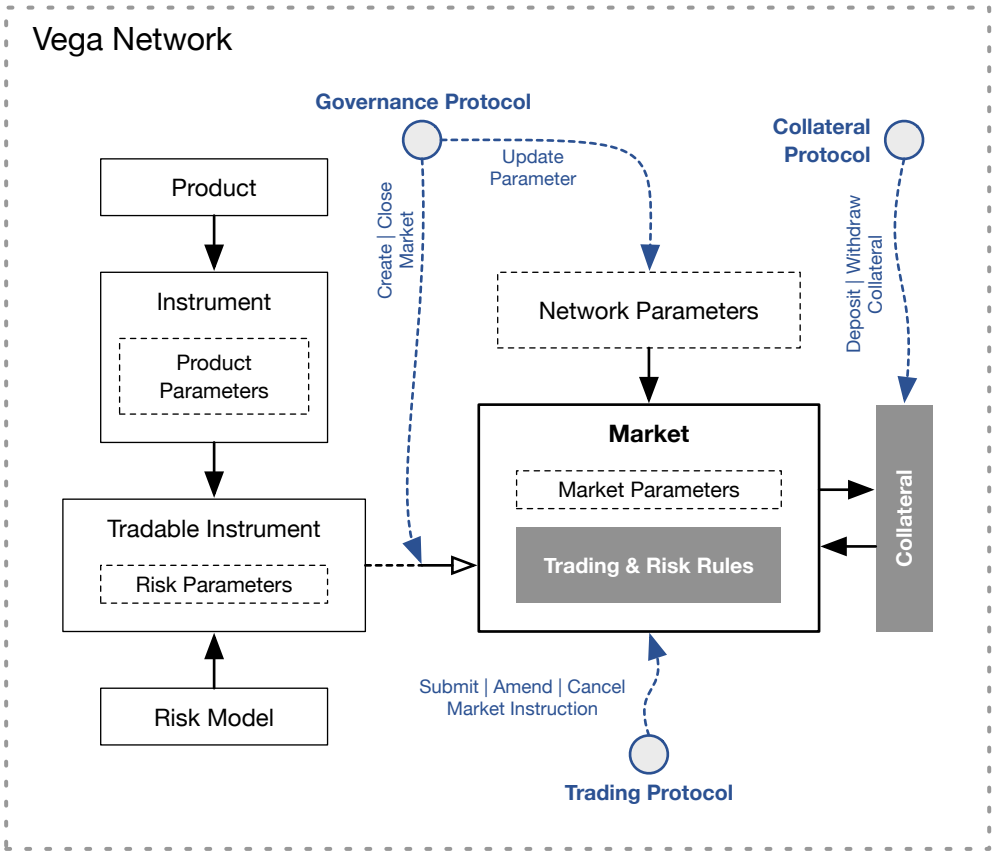

Figure 3: Vega's hierarchical, parametric market framework[10]

3.1 Network

Conceptually the top level of the Vega framework is the network, which encompasses all of the markets, products, participants, and governance actions contained by a physical Vega network (see Section 1.2), which may potentially be partitioned into multiple *shards* to achieve the required scalability (see Section 2.1 iv, v). The network is the level at which most governance transactions (see Section 8) act, and also where a number of *network parameters* are maintained. Collateral balances are maintained at the network level.

[10]Note, this diagram is focused on the interactions between incoming transactions and the market framework, and the transactions shown are not exhaustive.

3.2 Products

A product defines how trades will behave. It specifies how and when to calculate settlement cashflows (see Section 5.2), and contains metadata to tell the protocol about any *product parameters* — such as the source for an *underlying* data feed, a strike price, maturity date, or other more exotic inputs — that are required to create an instrument (see Section 3.3) using the product. Products can also make use of external data — such as a price from another market, or whether or not some event has occurred — either in the form of data from Vega markets, or special signed transactions from pre-approved outside parties known as *oracles*.

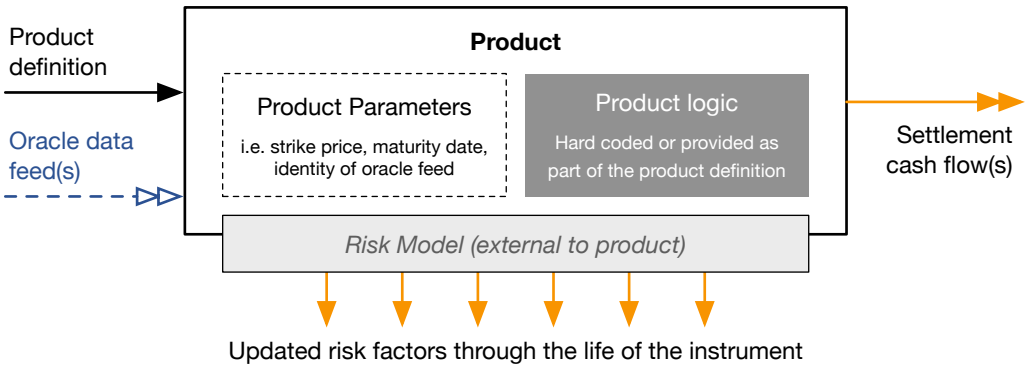

Figure 4: The flow of data and actions in and out of Vega products

Products are intentionally treated as somewhat of a black box, meaning that the protocol defines the interaction points between products and the other components in a network but does not interpret or control the functionality of products themselves. In this respect they can be considered as domain specific smart contract. Products interact (see Figure 4) directly with the protocol when a position is settled, and indirectly through the associated *risk model*[11], which determines the amount of collateral (see Section 4) to be allocated to, or released from, a risk universe (see Section 3.5) for a given position.

The primary advantage of this approach is to decouple product definition from the general protocol, allowing for significant evolution in products within the market framework. This is exemplified by the fact that although prototypes of the Vega reference implementation use a fixed number of hard coded products, the protocol has been designed from the outset as an execution platform for *smart products*, a type of smart contract designed for financial products that will allow Vega, once complete, to host almost any product or market imaginable. *Smart products* will be constructed using a specially designed *smart product language* and will be discussed further in an upcoming paper to be published at vega.xyz.

3.3 Instruments

An **instrument** represents something that can be settled, defined by the combination of a product (see Section 3.2) and all of the required *product parameters* to fully satisfy that product; e.g. for an option, an instrument may consist of the product plus a reference to an *oracle* for the *underlying* price, a strike price, and a maturity date [8, p. 151–155]. This combination of a product and all of its parameters uniquely identifies an instrument, meaning that an instrument cannot be duplicated on a Vega network[12]. Instruments are the building blocks of markets and the level at which markets are de-duplicated, but are not themselves tradable.

[11] Note that products do not contain *risk models*, however they are linked, as a *risk model* must have some understanding of the product for which it is calculating risk factors (see Section 6.1).

[12] Although it is possible for an instrument to be considered redundant (which may cause market creation to be vetoed) if its product and parameters too closely match another, see Section 8.2.

A **tradable instrument** is a complete *instrument* (product and parameters) as described above, combined with a validated *risk model* (see Section 6.1) and its required *risk parameters*. A valid *tradable instrument* contains all of the data needed for Vega to execute trades, calculate margin, and carry out settlement. A complete and validated *tradable instrument* is required to initiate a trade managed by the Vega protocol or create a market of any sort.

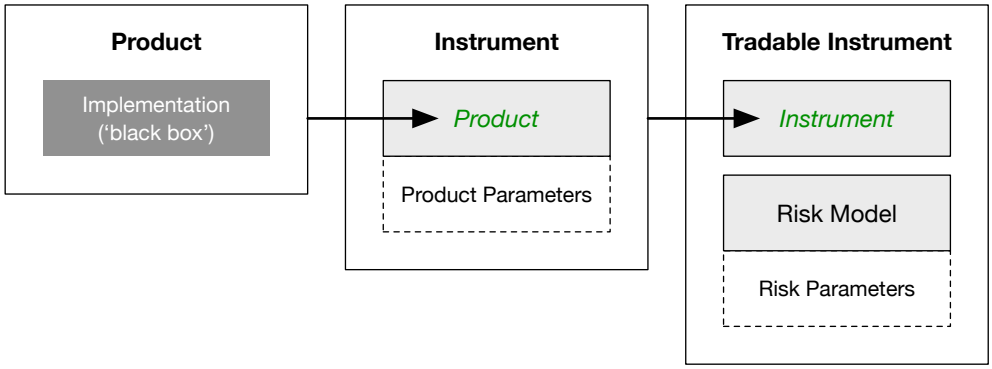

Figure 5: The relationship between *products*, *instruments*, and *tradable instruments*

3.4 Markets

Markets represent a *tradable instrument* that has been configured for a specific method of trading. A market may be *proposed*, *active*, *suspended*, or *closed* depending on where the instrument is in its lifecycle and whether any risk management (see Section 6) or governance (see Section 8) actions are currently impacting trading. Markets are specified via their *market parameters*, which include the *trading mode* (see Section 5.1) and any applicable parameters for that trading mode.

Markets fall into two categories:

- **Open markets** which, as the name suggests, are open for any sufficiently collateralised participant to buy or sell. These are analogous to the public markets that exist in many countries in stocks, foreign exchange, and other asset classes. They generally trade using a limit order book (a.k.a *continuous trading*), although other modes such as *discrete trading* via frequent batch auctions are also considered (see Section 5.1). Open markets require the support of *market makers* and must be created and approved via the network's governance process (see Section 8.2) to ensure that the *risk model* and *risk parameters* are set correctly, and that new markets meet the community's standards.
- **Ad-hoc or *OTC* markets.** These are created as needed by a participant who wishes to initiate a trade, either at a price they have agreed offline, or by using the Vega protocol to facilitate a price discovery process, for instance with an *RFQ* process. These *trading modes* are covered in more detail in Section 5.1 (items v, vi).

3.5 Risk universes

A *risk universe* is a set of markets sharing the same *risk model* and *risk parameters*, to permit perfect netting of margin between markets beyond the partial netting achieved by *coherent risk measures* (see Section 6.1). This allows more efficient capital usage for margins in highly related product sets, such as futures of different maturities in the same *underlying* asset.

Very often a given market will not be a member of an explicitly defined *risk universe*, so, by default, and unless specified otherwise during market creation, a market is said to reside in its own implied *risk universe* with a population of one. Addition of a new market to an existing risk universe must be carefully scrutinised as part of the creation process (see Section 8.2) as adding the new market will affect margin calculations.

4 Collateral management

Trade settlement occurs in crypto-assets hosted on other blockchains. To ensure safe trading in a trust-minimised environment, all trades on a public Vega network will need to be sufficiently collateralised in order to avoid being closed out (see Section 6.1). Collateral in the base currency and any other settlement asset(s) will therefore need to be held and managed in a decentralised fashion by the network.

The protocol is designed to keep track of balances on a per-participant basis (e.g. per public key, or wallet address) in an arbitrarily large number of crypto-assets. Collateral management is furthermore designed to use a simple deposit and withdrawal protocol to interface between a Vega network and a crypto-asset's host chain, allowing for more assets to be added relatively easily based on demand. The performance of deposits and withdrawals for a given crypto-asset is dependent on the performance and finality properties of its host chain, but once deposited, any crypto-asset can be transferred, settled, allocated, and released at the speed of the Vega infrastructure, regardless of the properties of the asset's host chain.

4.1 Acceptable types of collateral

The defining factor for whether a given asset can be used with Vega is the network's ability to manage it with certainty. Collateral may be held by a participant in any fungible digital asset[13] for which the following conditions hold true:

i) Vega nodes can digitally verify the participant's access to the asset;

ii) it can be placed indisputably under the sole control of the network;

iii) it can be irrevocably transferred to any other participant by the network; and

iv) asset balances can be released by the network when required.

In practice, this requires a system that supports sufficiently powerful scripting or smart contracts to operate a deposit address with the ability to pay out funds using a threshold (k-out-of-n) authorisation scheme [1].

Currently, the only assets that meet all four requirements are blockchain based cryptocurrencies and digital tokens. We envision that the first public implementation would integrate with the Ethereum blockchain, due to its smart contract capabilities and the existence of more than 1500 ERC20 tokens[14], followed soon after by Bitcoin because of its dominant size.

4.2 Depositing and withdrawing collateral

Collateral is deposited by placing it under a Vega network's control. This generally means paying it to a wallet or contract address from which it can be released with consensus agreement from the Vega nodes — and by no other method — to any address on the collateral chain as required. Nodes monitoring the collateral chain will post updates that have not been included by the network as new 'notify deposit' transactions (see Figure 1). These may initially be rejected by a majority of nodes until a quorum is reached that has observed the new deposit.

Participants may request to withdraw unallocated collateral from Vega at any time. Withdrawal transactions are processed by Vega nodes signing either an approval or a rejection message. In the case of approval, the message will also contain a payload that is a signed transaction for the host blockchain on which the collateral resides. By combining approval payloads from a supermajority of nodes and submitting them to the collateral blockchain, the requester will be able to complete the withdrawal.

[13] Subject to the existence of an appropriate inter-blockchain bridge, which in many cases may be non-trivial.

[14] Which allows for earlier experimentation and innovation, in areas such as settlement of tokenised assets and multi-asset settlement, versus starting with an effectively single asset chain such as Bitcoin.

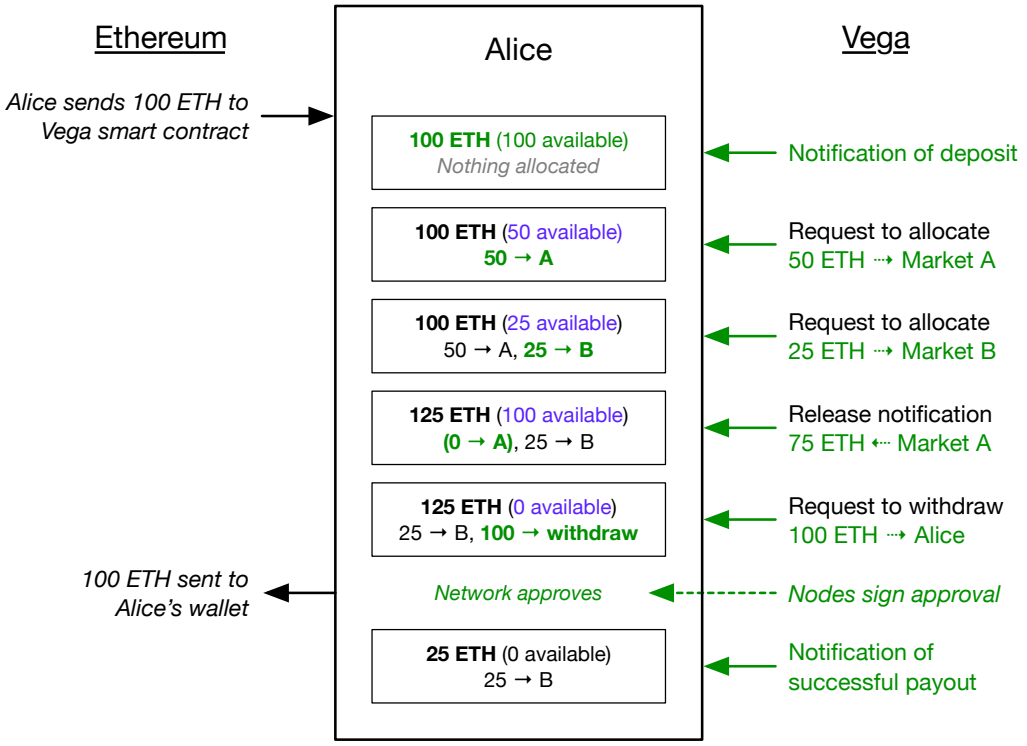

Figure 6: Mechanics of collateral deposit, allocation, and withdrawal

4.3 Allocation and release of collateral

When collateral is required for an order or trade, it is allocated to the relevant *risk universe* and becomes unavailable for use in other markets or withdrawal. A *risk universe* in which a trader has active orders or positions will from time to time need to allocate or release collateral (see also Figure 6, Figure 7), when:

- a new order is submitted and insufficient collateral is reserved to cover the participant's updated worst case requirement;
- a position moves against the trader breaching the configured collateral call level and unallocated collateral is available;
- a position moves in the trader's favour, causing them to have excessively high allocated margin, in which case some collateral will be released; and
- a position is fully or partially closed and the collateral is no longer required.

Depending on the architecture of the Vega network in question[15], the allocation and release of collateral may be able to be done atomically within a single transaction if collateral management is handled locally, or may require an asynchronous request to another blockchain or *shard* where it is done elsewhere.

The cost of asynchronous collateral requests, along with the potential limitations they place on scalability — where collateral requests may become a bottleneck, even where trading itself is distributed among many chains — present the biggest performance risks in the design of the Vega protocol. To reduce the latency impact on trading in this scenario, orders may be accepted optimistically and the network may implement penalties to deter attempted double spending of collateral[16]. No ideal solution for the throughput limit on collateral requests is known at the time

[15] Particularly, if collateral is managed on the same blockchain as the *risk universe* requiring the allocation/release, i.e. whether the Vega network in question implements *sharding*.

[16] Note that this is possible for a short period as the P&L of a newly opened position is ≈ 0.

of writing; this is an area of active research, however it is not considered critical as (a) it will not present a problem early in the life of a Vega network, and (b) an acceptable if not ideal work-around exists in the form of maintaining multiple balances per asset class for each participant such that the collateral transactions themselves are partitioned.

Prior to processing any transaction requiring additional collateral allocation, the participant's unallocated collateral must be checked against the minimum requirement to process the transaction (for more on margin calculations, see Section 6.1). Assuming the initial check has passed, the collateral may be reserved (allocated to the *risk universe*) and the transaction processed; if not, the transaction is rejected.

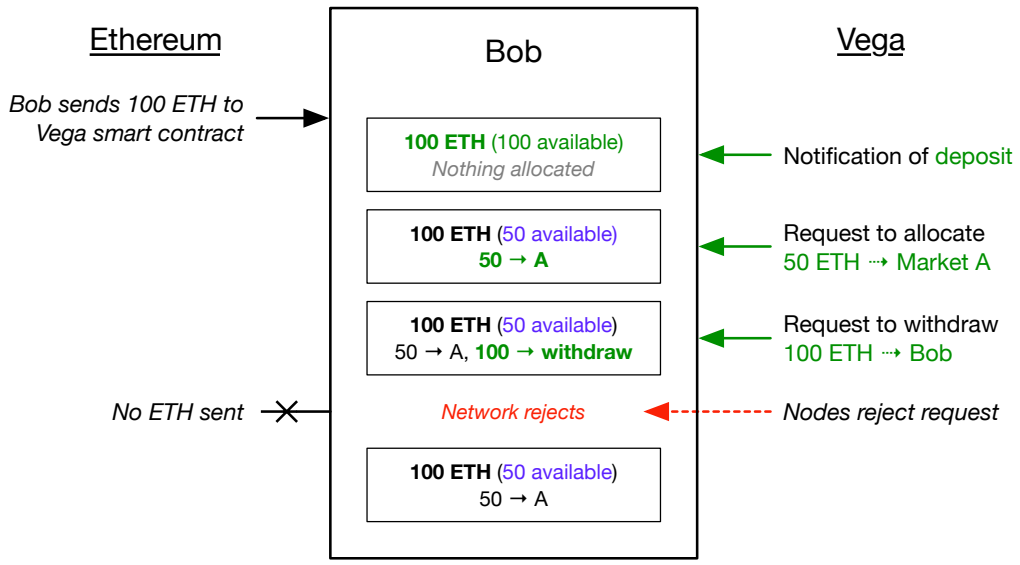

Figure 7: Example of a rejected collateral withdrawal request

4.4 Collateral maintenance levels and margin calls

The minimum collateral requirement to maintain an open position (*maintenance margin*) and all open orders (*initial margin*) is calculated periodically. Table 9 in Section 6.1 describes the four collateral zones that the network tracks; the zone exceeding initial margin; the maintenance zone; the collateral search zone and the closeout zone. The *initial margin* is the amount required to enter into a trade. Entering the collateral search zone prompts the collateral manager to attempt to allocate more collateral to a particular *risk universe*, attempting to forestall a forced closeout. Entering the closeout zone prompts forced closeout of positions, with any remaining allocated collateral transferred to the risk universe's *insurance pool* (see Section 6.4)[17].

A trader may also inform the network that if a position reaches the collateral search level whilst they have unallocated collateral available, the network should automatically attempt to allocate additional collateral, in which case the position will be recapitalised automatically by the network if possible. If recapitalisation is not configured or not possible, the position will enter the collateral search zone, however the network will not take further action until another collateral zone is reached.

Unlike centralised markets, margin calls will be performed by client software, algorithms or third party services, and will not be issued by the network. The collateral levels for margin calls are the responsibility of the trader and may be set according to their trading strategy and risk appetite.

[17]Open orders will be cancelled with no collateral penalty to the trader.

5 Trading & settlement

Decentralisation requires that trading instructions, like all other transactions, are executed deterministically on all nodes. This is necessary to ensure that precisely the same results for matching, risk management, and settlement of all orders and positions are seen by every participant. Furthermore, it is necessary that every node also evaluate the state of each market, to determine the appropriate *trading mode* according to the applicable *market parameters*, liquidity, and recent price history. A consequence of this requirement for absolute determinism is that every action[18] in the markets — including processing of normal orders, margin related closeouts, and responses to risky price or liquidity conditions — must be completely automated from end-to-end, with no manual intervention or exceptions process. This is unlike almost every centralised market and presents unique challenges.

This fully automated trading and settlement logic works within the *market framework* (see Section 3) and interprets the data structures and transactions to determine the state of a market at any given moment in time. The workflow during normal trading is illustrated in Figure 8. The protocol's internal state includes the full order, trade (including any closeouts), and price history for the market, in addition to the current state of the order book, open positions, and the margin requirements for each participant — all of which are transparent and publicly visible.

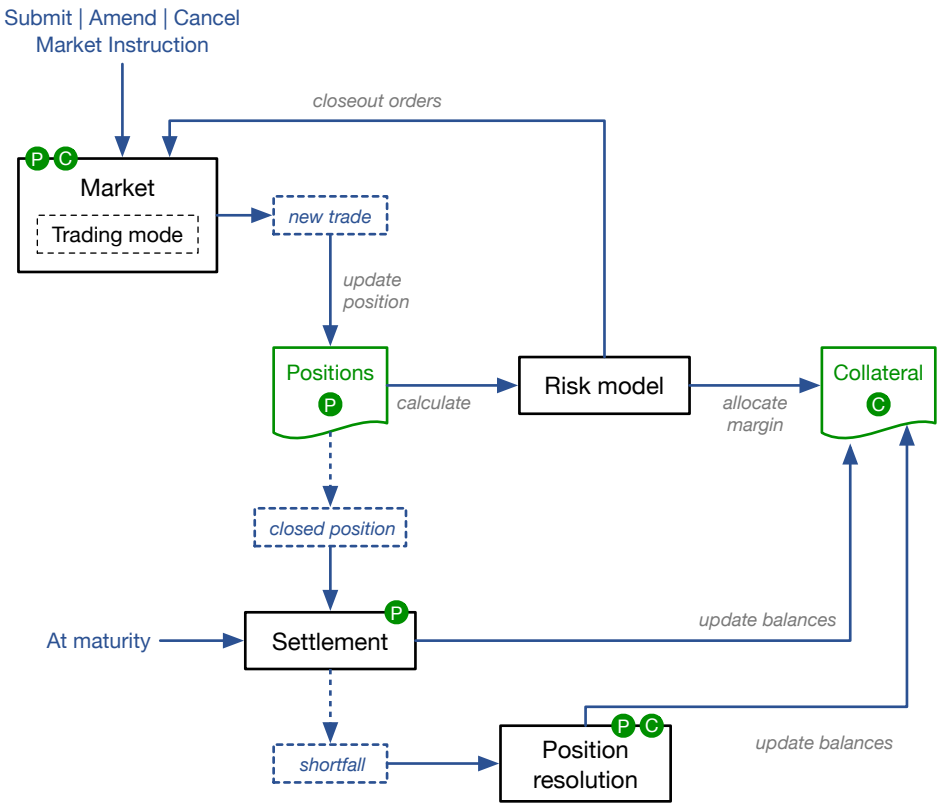

Figure 8: Fully automated trading and settlement workflow

[18] Aside from a small number of governance actions (see Section 8), which involve participant voting and occur over extended time periods in comparison to market events.

5.1 Trading modes

Trading modes are the methods by which new trades are generated, some of which are designed for liquid *open market* trading whilst others facilitate *OTC* markets in bespoke or illiquid instruments. All trading modes may support multiple price determination methods, although a precise discussion of trading and pricing algorithms is out of scope for this document, as the protocol is agnostic to the specific implementation details. Trading modes are set for *open markets* during market creation (see Section 8.2); for *OTC* markets, which are created in a more ad-hoc fashion, the trading mode is defined as needed.

The following trading modes meet the requirements for flexibility and fairness set out in Section 2.1 and cover the needs of both *open market* and *ad-hoc* trading (see Section 3.4). The list also includes the protective modes required by the risk and governance protocols (see Section 6 and Section 8 respectively).

i) **Continuous trading** operates a decentralised version of the limit order books used by most public (and many private) markets worldwide. The most common configuration is for orders to be matched in price/time priority and trades to take the passive price from the order book, however alternative prioritisation schemes and pricing algorithms may be implemented, configurable via *market parameters*.

ii) **Discrete trading** provides *open market* trading via frequent batch auctions [19]. Originally designed as a response to the high frequency trading 'arms race', frequent batch auctions are well suited to a blockchain environment in which transactions are batched into blocks. This mode at its core implements a series of very short auctions and, as such, price determination can use any pricing algorithm available in an auction period.

   This method also shows interesting potential for operating semi-private decentralised markets via a two-stage process known as commit and reveal in which participants first submit orders encrypted with a secret-sharing [1] [26] algorithm, and in round two submit portions of the decryption key set to nodes on the network. This creates a simultaneous reveal of orders and subsequent trade execution round, ensuring that participants cannot front run or adjust prices based on short term micro-structural phenomena.

iii) **Auctions** can be used in standalone form to create a special type of one-time market, or more commonly in an *open market* that normally uses another trading method. In this case, a market is temporarily moved to an auction period — with the expectation that it reverts to its normal trading mode after the auction concludes — under certain situations, such as:

   - during market creation (see Section 8.2) before the market first opens;
   - when trading resumes after suspension due to illiquidity (see Section 6.2); or
   - to find the fair market price during large price moves (see Section 6.3).

   During an auction, participants are able to submit, amend, and cancel buy and sell orders for a pre-determined period of time, called the *auction call period*. This will end either at a given logical time, or after the network reaches consensus that the call period is over. The price and volume that would be achieved if the auction were to end immediately at any given time are known as the *indicative uncrossing price* and *indicative uncrossing volume*, and can be used by participants in much the same way as the evolving trade price in a continuously traded market to inform their trading decisions. Once the call period is over, an *auction uncrossing* occurs with trades being generated at a price in the range that maximises the total trading volume. There are multiple potential price determination algorithms, the details of which are outside the scope of this paper.

iv) A **suspended market** occurs when an *open market* is temporarily stopped from trading to protect the market or the network from various types of risk. Suspension is a last resort used when the system has determined it is either not safe or not reasonable to operate the market at the current time, for example due to extremely low liquidity (see Section 6.2).

   Suspension operates like an auction call period with no defined end: orders will be accepted to the book but no trades will be executed. In some cases, the criteria for lifting the suspension

may be linked to the depth of orders. If the suspension is to be lifted, this will be achieved by changing the trading mode to an auction with a defined end time, after which the market will return to its normal trading mode.

v) **Request for quote (RFQ)** can be used for trading more bespoke or illiquid instruments on *OTC* markets. This mode allows traders to advertise their interest to trade any instrument. Other traders and *market makers* will be free to make prices at their discretion and quotes, once accepted, will result in trades margined and settled like any other. Quoted markets differ from *open markets* in that the risk models used for margin calculations will always require external data feeds to mitigate the lack of liquid prices.

vi) **Matched trades** are a method for any two participants to directly submit a trade to be managed on the Vega network. As long as both parties submit the same or compatible trade details within a given time window, the trade will be accepted and managed like an *RFQ* driven trade.

5.2 Settlement

Settlement is the mechanism for ensuring the accounts of all traders are credited and debited with the correct amount of collateral at specific points in the trade lifecycle:

i) **At expiry of an instrument**, the product's valuation function (see Section 3.2) will be passed any required settlement data, causing it to generate settlement instructions to determine the net asset flows to and from each trader with an open position. After settlement at expiry, all positions are closed and collateral is released.

   To prevent gaming of the settlement due to the probabilistic nature of time in a decentralised environment, trading will be closed sufficiently prior to expiry such that there is no chance of any new transactions being accepted once the settlemen data may have become available.

ii) **When interim cashflows are generated** by a *product*, collateral positions are updated to reflect these cashflows, and margin requirements are recalculated based on the updated collateral position and market data.

iii) **When closing all or part of a position** by entering into a counter trade, including if that trade is created as part of a forced risk management closeout.

   The closeout cashflow is calculated as the difference between the opening and closing price for the position, where the closing price is the volume weighted average price (VWAP) of the closing trades, and the opening price is the VWAP of the oldest v in volume that has not yet been closed, where v is the total size being closed.

In the event of a shortfall in which a participant does not have enough collateral to cover the required cashflows, the *position resolution* algorithm (see Section 5.3) is used to resolve the situation fairly and deterministically.

5.3 Position resolution

When a trader's liabilities in a market — comprised of any outstanding settlement cashflows plus the required margin for their active orders and open positions — are found to be greater than their available collateral, their position will be called *distressed* and the *position resolution* algorithm will be invoked to unwind their obligations.

When a market has one or more *distressed positions*, it is in a potentially unstable situation in which essentially randomly selected counterparties[19] could experience large losses without warning. We solve this problem by designing Vega such that in reality, all trades are back-to-back trades with the decentralised network itself acting as the intermediary, and the pool of funds available

[19]Given that *open markets* are pseudonymous and it is not possible for traders to be selective about the identity or creditworthiness of their counterparties.

to a market being limited to the total collateral allocated to the *risk universe* (see Section 3.5) and the market's *insurance pool* to protect the network from contagion between markets.

The *position resolution* algorithm defined below attempts to ensure that *distressed positions* are closed fairly and deterministically, to minimise the risk of excessive impacts on any individual participant from a defaulting trader, and to protect the market as far as possible in volatile and illiquid situations. The algorithm is based on three key principles:

- credit events should impact the whole market, not individual counterparties;
- *distressed positions* never threaten other markets or an entire Vega network; and
- in the event of uncovered shortfall after exhausting available collateral and the *insurance pool* (see Section 6.4), contributions by profitable traders will be based linearly on their profitability.

Position resolution algorithm

When *position resolution* is required, we must first evaluate all outstanding settlement instructions to ensure that all *distressed positions* have been identified, after which the following process is executed in order:

1. Open orders are cancelled, oldest first. Participants with remaining *distressed positions* at this point, after closing all open orders, will lose all of their collateral that is allocated to the *risk universe*, regardless of any remaining balance after closeout trades.
2. The net closeout volume NCV is calculated as:

$$NCV = \textit{Long Closeout Volume} - \textit{Short Closeout Volume},$$

   where *Long Closeout Volume* and *Short Closeout Volume* are the sum of the volume of distressed long positions and distressed short positions respectively.
3. A market order is placed for the absolute net closeout volume unless $NCV = 0$; the order is a **sell** if $NCV > 0$, otherwise $NCV < 0$ and a **buy** order is placed.
4. Any shortfall between the settlement cashflows for the closeout trades and the traders' available collateral is paid out of the *insurance pool*, and any remaining collateral after settling the closeout trade is transferred to the *insurance pool*.
5. In the event that the *insurance pool's* (see Section 6.4)collateral is exhausted with unpaid settlement cashflows still outstanding, the final remaining shortfall will be made up through contributions from market participants with profitable open positions. The relative contribution of each trader holding an open position takes into account their relative (unrealised) profitability and size of their open position.

Note that during this process, market protection measures such as those that protect against excessive price moves and low liquidity remain in effect and could result in the market being suspended or entering a price discovery auction before the closeout process completes.

6 Network risk

Risk management is a complex topic, certainly one too large to do full justice in this paper, and also one that will evolve with the protocol and the ecosystem around it. In this section we have considered the most critical risks, particularly those that must be definitively addressed from day one by rules in the trading protocol itself, along with various forms of gaming, manipulation, and unfair advantage that might arise in a public network running the Vega protocol.

By far the most important risk under consideration in a pseudonymous environment is that of credit risk, given that there can be no expectation of recourse in the event of a counterparty walking away from their losses, should they be given the chance. We also define some rules to detect dangerous market conditions and apply protective measures, and discuss some of the theoretical strategies that a malicious actor may use to attempt to extract unfair value from a Vega network or deny service to participants.

We recognise the huge scope and serious implications of this topic, and the need for further work, and have dedicated significant resources to all areas of risk, about which more details, additional research papers, and the latest updates will be published at vega.xyz.

6.1 Credit risk

The primary financial risk facing a Vega network is credit risk. On a platform where counterparties may be identified by no more than a public key, there is no recourse in the event that a trader owes more in settlement than their posted collateral. It is therefore essential that the protocol be designed to constantly maintain effective collateralisation for all positions.

This property of cryptocurrencies and decentralised systems in general leads us to design Vega as a collateralised, margined, and leveraged[20] platform with margin calculations taking into account the probability of the liquidation value of a position falling short of the available capital. Although the inability to assume any creditworthiness whatsoever for participants serves to push up margin requirements somewhat, this is mitigated by the ability of the network to re-evaluate the risk frequently and pre-emptively close positions.

Margin and risk are assessed across a *risk universe* (see Section 3.5) which is comprised of a single instrument or a collection of instruments that may provide collateral offsets to each other. If the risk arising from a trader's net liabilities exceeds the minimum required margin amount, then the *position resolution* algorithm (see Section 5.3) is used to settle the *distressed position*.

Approach to risk modelling

Margin calculations are performed on the net riskiest composition of a trader's open positions and live orders (assuming the scenario that the orders were successfully executed). This resultant worst case hypothetical position is referred to as a liable position for risk calculation purposes.

For most products their value is only known at the time when they were traded (the value equal to the price at which the trade executed) and at maturity. At any intervening time one has to use a stochastic *risk model* to provide a reasonable estimate of the value. These are provided using risk-neutral pricing (see [22, 13] or [25]). The most well known example of such a model is the Black–Scholes model[21]. However there is now a whole universe of models used by various institutions for different products. In general these models do not provide simple formulae for calculating derivative prices (and hence their risk). Instead Monte-Carlo simulation methods are used in pricing. For a general introduction see e.g. [9]. For the problem of nested simulation arising specifically in risk calculations see [20].

[20] When appropriate. Vega will also support full collateralisation and profit/loss ceilings where no suitable risk model exists to calculate a safe margin level for given instrument.

[21] This model is known not to capture essential market features and thus has been superseded by more advanced ones. However it is easy to calibrate and provides closed-form formulae for prices.

Vega requires the development of a variety of appropriate *risk models* to allow for margined trading of the widest possible range of *tradable instruments* that satisfy the following requirements:

i) calculations are as transparent as possible to the trader and maintain a manageable pace of change in margin requirements;

ii) models are robust and theoretically sound to support leveraged trading; and

iii) netting is supported across products and asset classes so that unrealised gains in one market may be used to offset margin requirements in another.

Note that for certain types of products there will be no reasonable risk model that will capture the liabilities well. In this case the margin will be set to the full liable position, which precludes leveraged trading but protects the network and its participants.

Margin calculation

With this in mind, Vega will use the following methodology for margin calculations. The maintenance margin level $m^{\text{maintenance}}$ for one unit of the traded product will be determined as

$$m^{\text{maintenance}} := \text{Net Closeout P\&L} + \text{Market Observable} \times \text{Risk Factor}\,. \tag{1}$$

The *Net Closeout P&L* is the volume weighted (as seen on the order book) price of the trade that would be required to close the trader's position minus the trade price (the price at which the relevant trade was entered).[22] The *Market Observable* is any value that can be directly seen in the market, for example for futures it is simply the current price at which the futures are trading and for options it is the current price of the underlying. The *Risk Factor* is a number that will be calculated by an appropriate stochastic risk model, and will be different for long and short positions. If there is no appropriate stochastic risk model then the *Risk Factor* will reflect the entire liability while the *Market Observable* will be set to 1.

The *Risk Factor* should remain relatively stable during the lifetime of the product so that most changes in the minimum margin requirement can be explained by a trader considering their *Net Closeout P&L* together with moves in the *Market Observable*. This ensures that the margin calculations meet the first key requirement (i) above. This can be further enhanced by applying additional 'padding' to the effective *Risk Factor* used by traders in order to reduce the frequency with which margin requirements change independent of market moves and/or to allow for advance warning of changes.

To meet the remaining two requirements, we turn to the well established theory of *coherent risk measures*. If a trader has liable positions that will have value X (calculated from a *risk model*) at a future time $\tau > 0$, then the corresponding *liability* is $-X$. The margin has to reflect the risk inherent in such a liability and Vega will use an appropriate *coherent risk measure* to calculate the required margin.

Mathematically, a *risk measure* is a function, say ρ, which assigns a real number (the risk) to a random variable (the portfolio). For a *risk measure* to be called *coherent*, it will satisfy the following properties [12]:

i) *Monotonicity*: if Y is another portfolio payout such that $X \leq Y$ then $\rho(X) \geq \rho(Y)$. So the portfolio that always pays less is deemed more risky.

ii) *Cash invariance:* if $m \in \mathbb{R}$ then $\rho(X + m) = \rho(X) - m$. So adding a non-random cash amount to a portfolio reduces the risk exactly by that amount.

iii) *Positive homogeneity:* if $c > 0$ then $\rho(cX) = c\rho(X)$. For example, if we double our portfolio then the risk also doubles.

iv) *Subadditivity:* if Y is another portfolio payout then $\rho(X + Y) \leq \rho(X) + \rho(Y)$. Thus a diversified portfolio $X + Y$ will have risk that does not exceed that of X and Y taken separately.

[22] For products that become illiquid this can be determined from the risk model price of the product.

Margin evaluation, maintenance margin and closeout

The Vega protocol will use a *coherent risk measure* which can be scaled by a parameter $\lambda \in (0,1)$. This parameter can be adjusted to ensure that the amount of collateral in the *insurance pool* doesn't get depleted and also doesn't grow beyond size appropriate for the market.

If a trader has open positions that will have value X at a future time $\tau > 0$ then we choose the *Risk Factor* for equation (1) such that

$$m^{\text{maintenance}} = \rho_\lambda(-X)\,.$$

However, if the risk inherent in the liabilities given by $m^{\text{maintenance}}$, exceeds m^{amount}, the amount of collateral a given participant has in a *risk universe*, then there is a probability that the network risk is high enough to justify initiating a close-out trade to reduce the risk. Therefore a trade will be closed out when

$$m^{\text{amount}} \leq \rho_\lambda(-X) = m^{\text{maintenance}}\,. \tag{2}$$

The amount $m^{\text{maintenance}}$ is the *maintenance margin*. Any margin in the trader's account that is left after the closeout trade is executed will be transferred to the network *insurance pool*.

To protect participants from unwanted closeouts the network will offer a user-configurable $\alpha^{\text{search}} > 0$ and if

$$m^{\text{maintenance}} \leq m^{\text{amount}} \leq (1+\alpha^{\text{search}})m^{\text{maintenance}}\,, \tag{3}$$

then the network will to attempt to allocate more collateral to this risk universe, see Section 4.4. Even if no more collateral has been allocated, the position will only be closed if (2) is satisfied.

Initial margin

A market parameter will specify $\alpha^{\text{initial}} > \alpha^{\text{search}}$ and the minimum collateral amount required for a new trade to be entered into is the *initial margin*.

$$m^{\text{initial}} := (1+\alpha^{\text{initial}})m^{\text{maintenance}}\,.$$

Having the initial margin level m^{initial} higher than the margin search level $(1+\alpha^{\text{search}})m^{\text{maintenance}}$ ensures that a small negative price move won't lead to a situation where the network has to attempt to allocate more collateral to this risk universe immediately after a trade has been entered into.

Coherent risk measures versus Value at Risk

Using *coherent risk measures* as the basis for margining means that the risk calculation is robust and theoretically sound. Moreover the *subadditivity* property of *coherent risk measures* means that we can net the margin requirements even across different markets. We will use *expected shortfall* (ES) as the risk measure for margin calculations. The average *Value at Risk* (VaR) / *ES* for a r.v. X representing payoff of portfolio and $\lambda \in (0,1)$ is

$$\text{ES}_\lambda(X) := \frac{1}{\lambda}\int_0^\lambda \text{VaR}_\alpha(X)\,d\alpha\,, \tag{4}$$

where *VaR* is defined as the minimum amount x such that $\mathbb{P}(X+x<0)$ is smaller than α. This is

$$\text{VaR}_\alpha(X) := \inf\{x \in \mathbb{R} : \mathbb{P}(X+x \leq 0) \leq \alpha\}\,.$$

We note that there are many other approaches to margin calculations. Of particular note is the ISDA SIMM approach [17, 24]. We believe that our approach is more robust when dealing with correlations (or the lack of knowledge about correlation) by using a *coherent risk measure*. In contrast, the ISDA SIMM model uses *Value at Risk* which is not *subadditive*. To overcome this, a correlation matrix has to be used and calibrated and moreover an assumption that risk factors

Margin Level	Network risk categorisation
$+\infty$ $\vdots$ $(1+\alpha^{\text{initial}})m^{\text{maintenance}}$	Sufficient to enter into a position, no financial risk to network.
$(1+\alpha^{\text{initial}})m^{\text{maintenance}}$ $\vdots$ $(1+\alpha^{\text{search}})m^{\text{maintenance}}$	No financial risk to network.
$(1+\alpha^{\text{search}})m^{\text{maintenance}}$ $\vdots$ $m^{\text{maintenance}}$	Collateral search zone, see (3).
$m^{\text{maintenance}}$ $\vdots$ 0	Closeout zone, see (2), acceptable probability of loss.

Figure 9: Collateral categorisation for at-risk trades

are jointly normal is made. Part of the reason the ISDA SIMM model is designed to be simple is that it is intended for non cleared derivatives. This means that each counterparty has to have an implementation and data sets to run the model. For Vega networks this issue does not arise as the network nodes will run the risk models on behalf of all participants.

Vega will use the expected shortfall (4) as the basis for margin calculations. While value at risk is a building block for expected shortfall Vega will never use it as a risk measure in its own right. Finally, to make the models robust it is important to use appropriate stochastic models. In particular the volatility parameters have to be set conservatively and / or models that imply fat tailed time-marginal distributions should be used, see [16]. For more details of the models and calculations used by the Vega network see [29].

6.2 Liquidity risk

The ability of markets to remain functional and fairly priced is dependent, among other things[23], on the presence of sufficient liquidity relative to the trading volume and open interest. Liquid markets allow participants the best chance of trading at a fair price, absorb new information faster, and are less susceptible to extreme price moves and 'flash crashes'.

In Vega's decentralised environment, illiquidity is an even bigger problem, as it impacts risk models [4], which could lead to erroneous margin requirements, and jeopardises the ability of the *position resolution* algorithm (see Section 5.3) to execute closeout trades for participants that are close to default (see Section 6.1), increasing the risk of loss to the network. Furthermore, when closeout trades do occur, they are more likely to move the price substantially in a low liquidity environment which may trigger extreme price move protection (see Section 6.3).

In the event that the buy or sell side of the order book volume is insufficient to close the largest single counter position held by a trader, the market will be suspended (see Section 5.1). The market will be reopened when the provision of additional liquidity by market makers or traders meets this requirement.

The aim of this mechanism is to ensure that rather than allowing a market's behaviour to

[23] Including the presence of informed traders and arbitrageurs, and timely dissemination of relevant information.

descend into farce through a self-reinforcing cascade of closeouts and ever larger price moves, the situation is arrested before it becomes untenable and no trading or settlement occurs[24] until the market is viable once again. This protects traders with open positions from unnecessary closeouts and *position resolution*, and makes it more likely that new buyers and sellers enter the market and allow trading to resume than if nothing were done. For example, a market may be awaiting an announcement that will cause a step shift in price, causing order book depth to reduce as traders await the outcome. This may trigger a market suspension until the news is received and the order book volume recovers sufficiently to reopen the market.

6.3 Extreme price moves

In some circumstances, an instrument's price may diverge unacceptably from the fair market price. This is more likely to happen in illiquid markets, and can often be exacerbated by interactions between automated trading systems that may react to large price moves in a way that reinforces them. This creates a need to monitor and react to extreme price moves, particularly given the fully automated nature of trading.

On Vega, however, the requirement to close positions at unacceptably high risk of default (see Section 6.1) creates the additional potentially amplifying force of closeout trades created during *position resolution* (see Section 5.3), which is most likely to happen during times of market disruption. The experience of other trading venues tells us that the failure to properly react to such events can create negative outcomes for both participants and the trading venue.

The protocol includes a 'circuit breaker' which is designed to ensure market price moves are reflective of the true supply and demand. The circuit breaker is triggered in the event that the change in price on an *open market* would be deemed excessive[25]. By excessive we mean a price move that is more likely to be an artefact of market microstructure rather than reflecting a genuine market move. What is excessive will be determined by the risk model view of likelihood of a given price move together with a *network parameter* specifying the exact threshold. For markets where no risk model is needed for margins, the protocol can use the assumption of price moves being normally distributed and determine the standard deviation from historical price data using a rolling window.

Upon triggering the circuit breaker, a market is transitioned to an auction (see Section 5.1) of a pre-defined length. The price resulting from the *auction uncrossing* is deemed to be a *fair price* and the market returns to its normal *trading mode*.

The goal of this monitoring and intervention is not to prevent large valid price moves, but to ensure that the fairest possible price is achieved by giving market participants time to respond to extreme price action. For example, a large sell order may initiate a significant downward price swing if there are not a lot of orders near the best bid; this would place the market (including the large sell order) into an auction, during which other buyers may enter the market in reaction to the low *indicative uncrossing price*.

6.4 Insurance pool

Each *market* (see Section 3.4) has an *insurance pool* that is utilised as a layer of collateral protection in case of a shortfall when closing out distressed trades. The *insurance pool* for a market will contain no funds when the market is first created but will gain funds over time from trading activity and as other markets mature and release their *insurance pool* funds.

In Section 4 and Section 6.1 we explained that a trader's position will be closed out using the *position resolution* algorithm (see Section 5.3) if the amount of collateral they have assigned to a *risk universe* (see Section 3.5) falls below the level required by the *risk model*. By virtue of the fact that all remaining collateral allocated to the *risk universe* for a position is held by the network once

[24] Unless the instrument reaches maturity while suspended, when settlement will be performed per the usual rules.

[25] Note the use of the words 'would be', as the circuit breaker must be triggered **before** and **instead of** executing any trades that would breach the *market limits* to prevent an anomalous price being reached.

position resolution is invoked, a closeout trade may result in either a profit or a loss to the network in question, depending on the slippage experienced by the closeout trade. If after the closeout trade, there is remaining collateral allocated to the *risk universe*, this is added to the *insurance pool* while a loss-making closeout trade will be funded from the *insurance pool*.

This mechanism to confiscate funds is designed both to ensure the *insurance pool* grows over time, and as a deterrent against collateral mis-management. The *insurance pool* is further funded in case a trader attempts to double spend collateral (see Section 6.1), and from penalties applied to a market maker who doesn't fulfil their obligations (see Section 7.2). If there are no funds left in the *insurance pool* to make good on a loss-making closeout trade, then step 5 of the *position resolution* algorithm is triggered to manage the socialisation of the loss (see Section 5.3).

At expiry of the order book and after final settlement, or if a market is closed by governance action (see Section 8.3) the *insurance pool* will be distributed between markets with the same *base currency*, volume weighted by the relative open interest. If a market is suspended, the *insurance pool* remains with the market until the market expires or is otherwise closed. Where there are no other markets for a *base currency*, the *insurance pool* will remain available for a period of time (a *network parameter* suggested at six months) and will be attributed to the next market successfully created with that *base currency*. Should no market be created, the *insurance pool* proceeds may be divided between the network's node operators, relative to their stake at the conclusion of the waiting period.

6.5 Risks from decentralisation and proof of stake

Markets using Vega are exposed to most of the traditional forms of market manipulation that traditional centralised exchanges suffer from. However, the decentralised nature of Vega networks also creates new ways of gaming the system via the transparency of transactions and the mechanics of the consensus protocol, which must be mitigated in the design of the protocol mechanics, and via the inclusion of protective logic where that is not possible.

Anyone who subverts the consensus mechanism can potentially obtain all the funds held in users' margin accounts, market maker stakes, and all the funds in the network *insurance pools*. The safety of the consensus algorithms is beyond the scope of this section, however proof of stake consensus algorithms are an area of active research both by Vega and also the wider blockchain community.

If someone accumulates a sufficient stake to control the outcome of the consensus algorithm then, again, they can potentially control all the funds in margin accounts and insurance pools. It thus follows that the theoretical minimum value of the crypto-assets used for *proof of stake* will be the total value of collateral that such an attacker could access. This may be reduced with the inclusion of mechanisms to restrict withdrawal of funds, or that ensure funds can only be depleted slowly. This is an ongoing area of research at Vega.

Depending on the exact consensus mechanism and algorithm implemented, at any point in time, some node known as the 'lead node' may know a definite list of proposed transactions in the next block before others. This may give rise to opportunities for front running. At Vega we are researching a number of mechanisms to prevent this. Broadly we focus on two approaches: one is to have the lead node chosen in a hard-to-predict fashion and changed frequently, thus making any chance of profit hard to realise.

The other, known as commit and reveal, is based on cryptographic methods and involves submitting transactions in a two step process: first participants submit binding transactions that are encrypted using a secret-sharing algorithm [1] and thus invisible to other participants. The lead node then proposes an ordered list of transactions but does not see other participants' intentions and hence has no scope for front running. The participants finally submit keys to decrypt the transactions from the second step, after which nodes can execute the transactions. This method impacts latency and would likely only be usable on markets using *discrete trading* (see Section 5.1).

6.6 Market manipulation and gaming

The aim of this work is to design a trading protocol that provides all participants with the same opportunities to profit from trading and leads to fair and efficient markets. However, in addition to exploiting the decentralised nature of Vega networks, there may be other ways to game the protocol and/or manipulate markets and prices. Market manipulation takes many forms but broadly speaking it is a deliberate attempt to interfere with the free and fair operation of the market and to benefit from such artificially created price moves.

In this category we broadly distinguish two types of gaming and manipulation: a nefarious participant might try to manipulate the settlement process (see Section 5.2) by feeding a node with incorrect information that benefits said participant directly, or they may attempt to benefit from the rules of the Vega protocol — for instance how it closes out trades (see Section 5.3), protects against illiquidity (see Section 6.2) and extreme price moves (see Section 6.3), and incentivises and penalises *market makers* (see Section 7).

Settlement may be based on data from another 'local' market on the same Vega network or one or more *oracles*. In the case of settlement based on a local market, there is an opportunity to attempt to shift the price by trading. In general such an attack will be avoided if products reference deep and liquid markets as source of settlement data. Getting an *oracle* to provide false data is theoretically possible but again can be mitigated by only using proven and reliable *oracles* that can provide cryptographically signed data feeds, and by using multiple *oracle* sources along with an algorithm that makes manipulation harder (e.g. remove outliers and average the rest), see also [2]. Finally, to mitigate attacks on settlement, a Vega network may enforce a 'cooling off' period before any settled funds are released, which can be augmented with a mechanism for re-running settlement if manipulation has been observed and corrected.

The question of how *position resolution* (see Section 5.3) may be attacked essentially boils down to: does the position resolution algorithm create arbitrage[26]? One clear possible source of arbitrage is if the margin level at which closeout is initiated (see Section 6.1) is set too low, thus allowing a participant to trade in a risky fashion, walking away from losing positions while pocketing gains. This can clearly be prevented by setting the risk parameters for the closeout level conservatively. At the moment the authors are not aware of any theoretical methods that would allow one to set the correct level and so this will have to be determined through simulation.

Another possible attack vector is via market manipulation that will put one participant into a position to benefit from another participant being closed out by the position resolution algorithm. This is a situation that needs further research and testing. There are models that capture behaviour of trading in limit order books [23], and these can serve as a starting point. Another approach is to simulate such situations on a test network. We note that market manipulation is prevented in part through the application of trading and price limits (see Section 6.3). These are calculated limits which take into account the instrument's volatility and open interest. These limits will restrict a single order trade size as well as a maximum allowable price move from a single order without triggering an auction that would likely prevent an egregious price move.[27]

Since *market maker* orders will automatically be posted and refreshed, it may be theoretically possible for a nefarious participant to take advantage of this. Careful design and testing of exactly how this happens should minimise the potential for abuse.

Most classical forms of market manipulation e.g. wash trading[28] are possible in a Vega market. These are mitigated by having deep and liquid markets, and the Vega protocol is designed from the ground up to provide such markets. We believe that with careful research, design and testing, and considered bootstrapping, it will be possible to prevent most market abuse.

[26] There are many ways to define arbitrage. One formulation is that arbitrage is a self-financing trading strategy that is guaranteed to never produce a loss and has strictly positive probability of turning a profit while not allowing unlimited exposure. The most general is that of 'no free lunch with vanishing risk' [11].

[27] Calculation methodologies for these limits will be tested, with details provided at a later date.

[28] Since the Vega protocol only identifies participants via their public key, it is possible for traders to hold multiple accounts and trade with themselves (wash trade). Wash trading requires two separate transactions to be submitted to the distributed network and there is no guarantee that a trader will trade with themselves, rather than another participant. Moreover, it will cost a trader in fees on at least one side of the transaction.

7 Liquidity

Unlike centrally managed markets — in which the operator's profits are derived largely from trading volume — an entirely peer-to-peer trading venue has no natural incentives for any individual party to source liquidity. This creates a bootstrapping problem in which low volume markets are unprofitable for *market makers* and the resulting lack of depth restricts the growth in trading that would attract market making activity. See also [3], [21] and [10].

The simplest solution is for the protocol to include a fee, payable to some organisation that is tasked with attracting liquidity. This has several disadvantages that, in our view, make it unworkable given the goals discussed in Section 2:

- it introduces unacceptable centralisation around the organisation in question, allowing it to favour some markets and effectively censor others by controlling the provision of liquidity;
- the organisation would be a bottleneck for the launch of new markets, which could severely compromise the ideal of permissionless market creation; and
- it would eventually give rise to a large and bureaucratic operation that would be liable to misallocate resources and lose out to more nimble ecosystem based alternatives.

To achieve our design goals therefore requires that incentivisation of market making be built into the protocol, and that this caters for markets with different trading volumes, and at different points in their lifecycle. This is achieved through facilitation at the protocol level of dynamically priced liquidity, which recognises that market making is capital intensive and thus aims for a market-driven solution that efficiently balances the need for order book depth on the one hand with a preference for low fees on the other.

The remainder of this section describes role of *market makers* and the mechanics of liquidity incentivisation in the Vega protocol in more detail.

7.1 Mechanics of the liquidity marketplace

Vega is in essence a peer-to-peer liquidity facilitation protocol with liquidity able to be priced [10] individually for each market (see Section 7.3). The liquidity fee is incurred — in a quantity determined by the volume and price of the potential trade, as well as the market's current liquidity price — by an aggressive, or price-taking order when it trades. In situations such as auctions (see Section 5.1) where there is no maker-taker relationship between the counterparties, the cost of liquidity is shared equally.

The liquidity cost is later credited during settlement to the participants responsible for the provision of liquidity in the market: the price maker, the market's infrastructure operators, and the *market makers*. The price maker and infrastructure operators will receive appropriate amounts, with the remainder divided between the instrument's *market makers*, with relative allocations based on individual contributions to the order book liquidity[29]. Market making volume is more valuable to the market when it is more competitively priced and consequently, the relative allocations between *market makers* of the liquidity reward takes their historical pricing into account.

Since *market makers* have a choice of where they deploy their capital, they will rationally select markets that offer the highest potential liquidity returns over their investment horizon. Their liquidity returns depend on: trading volume, their share of the liquidity rewards, and the liquidity price, which is always calculated in terms of the *base currency* of a market. Liquidity pricing is the protocol's mechanism to ensure liquidity is being attracted to markets that have the greatest need and that all markets operate at the most efficient costs for participants.

[29] Specific calculations for these values and their calibration will be covered in future work.

7.2 The role of market makers

Electing to become a *market maker* enables participants to be rewarded for the value they create by providing committed liquidity to markets by receiving a share of the proceeds from liquidity transactions across the entire market. This mechanic is designed, along with the requirement for a minimum market making commitment to create a new market, to incentivise market makers to act as owner-operators in their markets, and is a key part of the way the Vega protocol aims to create a thriving ecosystem.

To earn market making rewards, a participant must deposit a financial stake with their elected market (known as the market making stake). These funds are held by the network in the *base currency* of the market, with the size of the stake determining the minimum amount of volume a *market maker* will deploy into the market on each side of the order book. Rewards are calculated according to the total volume that is traded on the market, much like on a centralised exchange, with the percentage of the reward that an individual *market maker* receives depending on:

- their market making stake (relative to the total);
- their price making activity (active and competitive prices receive greater rewards)[30]; and
- the longevity of their market making commitment.

There is no limit on the amount of market making stake that a market will accept. Therefore, any Vega participant who wishes to be a *market maker* is able to do so. When an instrument's total market making stake is deemed to be above the instrument's comfortable minimum stake[31], market making stake may be withdrawn by any *market maker* wishing to 'resign' from market making or reduce their commitment.

Active market making requires the participant to actively manage a pricing strategy for their market making volume. All market making orders must be within a threshold specified by a market parameter of the last traded price, and must be actively priced by the *market maker* for a minimum percentage of the time (a *network parameter* initially proposed at 80%).

Market maker orders will be auto-refreshed by the network at the least competitive price within the threshold[32], such that the *market maker* always has their minimum obligated volume present on both the buy and sell side of the order book, however they will be penalised if they do not act to set a price for this volume and meet their active pricing requirements. The *market maker* may update the volume of their orders to be higher than their commitment so that they need to price their orders less frequently.

Passive market making, on the other hand, allows participants to support a market and provide liquidity without the need to actively make prices or risk-manage a portfolio of positions. Passive market making collateral is automatically deployed to supply liquidity algorithmically based on the prices posted by other *market makers*. Passive *market makers* share their market making reward with the active *market makers*, meaning the rewards are commensurately lower for the much lighter effort required. Note that a market's creator cannot be a passive *market maker*, and there is a threshold for the minimum active market making stake to ensure that there is always a significant active presence pricing orders.

The financial stake of the *market maker* is held as a bond against the fulfilment of their market obligations. *Market makers* must provide sufficient capital to support their minimum margin requirement, derived from the continuous provision of liquidity, in addition to open positions. Whilst the network will refresh *market maker* orders, the obligation remains with the *market maker* to ensure these orders will be accepted by the network (see Section 6.1).

If these orders are not accepted, the *market maker* will be penalised as they have failed to place an order. Active *market makers* will also be penalised if they fall below the threshold for the percentage of the time they are required to actively price orders.

[30] This will be calculated from historical price and volume activity, weighting recent activity more highly.
[31] Taking into account, for each instrument, its volatility, open interest and total market making stake.
[32] This prevents a *market maker* being penalised for having no orders because they have traded away.

Specifically,

i) a portion of the *market maker's* stake will be transferred to the *insurance pool* as a penalty, with the amount based on the scale of their failure to meet commitments; and

ii) a portion of the *market maker's* stake will be forcibly utilised as collateral if necessary to fund their order book obligations.[33] If the market contains sufficient market making stake to permit *market maker* resignations, their commitment will be reduced instead.

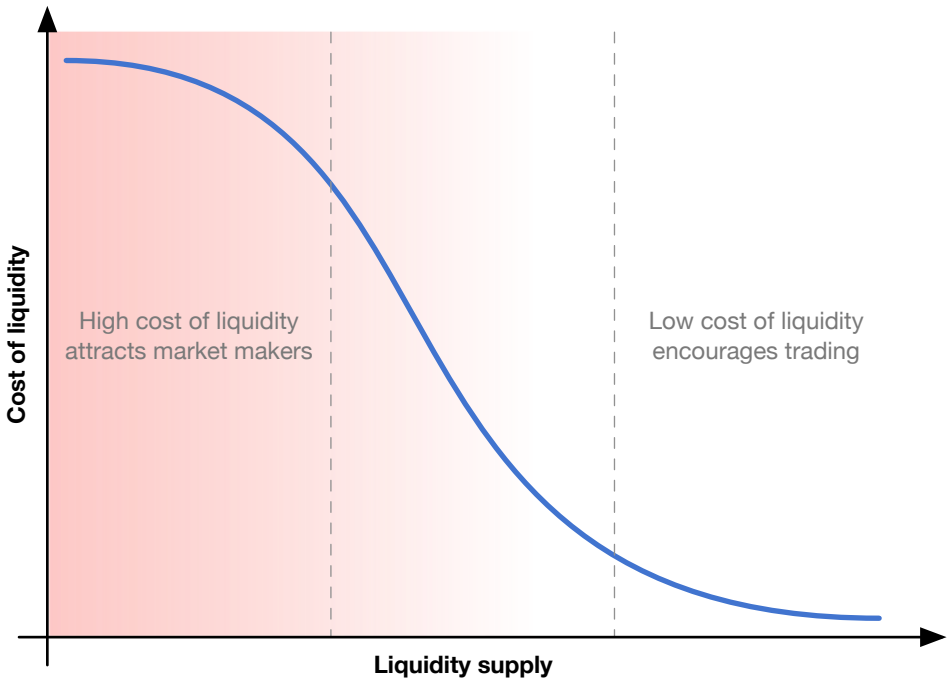

Figure 10: Illustrative example of a liquidity pricing curve

7.3 Dynamic liquidity pricing

In an efficient market, the liquidity price should capture the marginal benefit of attracting order book depth. We define this as the increase to trading volume as a result of incremental additional liquidity. Exogenous demand for the product at a price point is difficult to measure beyond what is displayed on the order book. We do, however, have endogenous demand within a Vega network due to the existence of closeout trades; this provides a useful lower bound to the value of additional liquidity, and thus the liquidity price at a point in time.

The liquidity price is a function of factors such as the open interest in the market, order book depth, observed trading volume, and current market making stake. The market's open interest captures the maximum endogenous demand at any point in time, the order book depth and trading volume calibrate the relative size of a unit of additional liquidity, and the current market making stake tells us the 'guaranteed' supply of liquidity. This is an area of ongoing research, results of which will be shared at vega.xyz. However, in the absence of this work, we can still usefully examine liquidity price dynamics as the balance of supply and demand in a closed system such as a Vega network.

Naturally, as the demand for liquidity increases, this should drive up the price. A market has an increased need for liquidity when:

[33] The amount to will be reallocated depends on the *market maker's* commitment versus their collateral shortfall.

i) **It has high credit risk**, which increases as the open interest in a market increases relative to the order book depth, because the protocol relies on the ability to immediately offload market risk from a *distressed position* to another participant. It follows that the more positions as measured by the open interest, the more potentially *distressed positions* and therefore, ceteris paribus, the higher the level of order book depth required for a given level of safety.

ii) **There is low trading volume** due to either a reduction in demand for the instrument (regardless of pricing), or the market being uncompetitive in price and/or accessible volume. In both cases, attracting *market makers* is beneficial as it will both increase the available volume, and likely improve prices, by provoking competition between *market makers*.

Conversely, a highly liquid market with high churn and trading volume has a reduced need for market makers and thus the cost of liquidity decreases accordingly. Market makers will benefit from a high volume of liquidity payments and traders will benefit from a lower liquidity price.

To capture these contributing factors, a liquidity pricing curve will be derived for each market. Initially we will develop a naive but workable pricing model, which will be improved over time with the application of our ongoing research and analysis of empirical data. Figure 10 shows an illustrative liquidity curve which captures the aforementioned price dynamics.

Note, the liquidity price is a single percentage value (e.g. 0.01%) from which the total liquidity reward per trade is calculated as a percentage of the trade's value (this could be notional volume or premium). Liquidity prices are adjusted periodically to balance the need to capture the changing market dynamics with the benefit of stable pricing. The time between updates will be a market or network parameter on the order of one day.

8 Decentralised governance

Vega is designed to run without centralised human intervention, and as such, operational governance is defined by rules embedded in the code, allowing for permissionless instrument creation and unlimited horizontal scalability. Network governance in Vega is not intended to replace all other forms of governance for the public Vega ecosystem[34], but to specifically provide on-chain governance for the key functions required to create and maintain high quality, well functioning markets in a decentralised environment.

8.1 Stake-weighted voting

The primary on-chain governance mechanic in Vega is voting by network participants based on their stake within the scope of the poll, for example:

i) for network governance issues, like the creation of a new market or determination of network parameters, stake would be measured in terms of a participant's holding of the network's native crypto-asset [28] (the *governance asset*); whereas

ii) for market governance decisions, stake may be measured by the notional value of a participant's net position, their market making stake or a combination of both.

Where a crypto-asset is allocated to a *risk universe*, either as fees or as trading collateral, it is considered to be held by the network rather than the participant and so will not be counted against the participant's votes. In some cases, the network may 'vote' for a default option with the weight of any such assets. However, assets held by a Vega network but not allocated to a *risk universe* will be included in a participant's stake for any votes.

A participant can only vote one way in any given poll, and will always be deemed to have voted with their full available stake[35]. The required majority for a decision and minimum participation will be defined for each type of poll, with a 2/3 majority and 0% minimum participation[36] being standard.

For votes requiring a certain (non-zero) participation level, the proposer will forfeit some of their *staking asset* if the minimum required participation level is not reached. This measure is in place to ensure that 'proposal spam' has a cost and that governance proposals do not create a vector for liveness attacks. In order to remove the incentive for various malicious behaviours such as voter bribery, there is the potential to introduce secret ballots for on-chain governance voting using principles from homomorphic encryption and secure electonic elections technologies. [15] [14].

8.2 Market creation

Open markets in Vega can be proposed for any *unlisted instrument* — i.e. an instrument with no active *open market* — by a participant holding a positive balance of the staking asset (see Section 8.1) and of the proposed market's *base currency*, some of which must be committed as market making stake, see Section 7.2.

A proposal must specify the *tradable instrument*, including product, *product parameters*, *risk model*, *risk parameters*, *trading mode*, and *market parameters* (see Section 3), and the size of the participant's market making commitment, which will become their market making stake (See Section 7.2). Proposals will be visible to all participants, and must successfully complete the process described below before a market becomes tradable.

[34] For example, the development of the reference Vega implementation, the ongoing development of the *smart product language*, and the suite of *risk models* for margining will occur offline.

[35] If, for some reason a participant wishes to vote with less than their full stake or support multiple options in the same vote, they will need to transfer some of their stake to a different wallet.

[36] This implies that, for example, markets are created even if only the proposer votes.

When a new instrument is proposed to a Vega network, the potential new market enters a proposal period — even if sufficient market making stake is immediately committed — to allow the network to ensure the market meets the community's standards, and during which the proposer must hold a positive amount of the *governance asset*. A market proposal is also deemed to be a vote in favour of the market's creation and a commitment to market making (see Section 7.2).

Other stakeholders may vote **for** or **against** the market, which will remain pending during the proposal period, a *network parameter* with a suggested length of five days. Participants can vote against the proposal for any reason, which may include:

- use of an incorrect, badly worded, or fraudulent market name or other parameter;
- poor choice of *risk model* or *risk parameters*;
- the proposed market being too similar to another market[37];
- proposal for an unethical market;
- concerns about the accuracy or trustworthiness of any *oracles*;
- the voter would prefer the market to use a different *trading mode*; or
- the voter prefers a different proposal for a similar market.

An instrument will remain pending indefinitely while the minimum market making stake is not met. During this period, other participants may join the proposer as a potential *market maker* for the instrument in the same way as for any other market. The proposer may also revoke their own proposal by voting against it, a revoked proposal will deemed to be cancelled and the market will not be created regardless of other market making commitment[38].

At the conclusion of the proposal period[39], the market will be created if sufficient collateral is allocated in market making stake and the **for** votes won. When a market exits the proposal period, the *tradable instrument* is considered to be a *listed instrument* and the market starts to accept orders. New markets will launch into an auction period before transitioning to their defined normal *trading mode*.

8.3 Market closure

In rare circumstances it may be preferable to close a market rather than allow it to continue trading. This may be the case if a market is later found to be fraudulent, if there is a failure or major disruption in an external price source, or for ethical reasons. A closed market will cease trading and settle all open positions immediately, opting either to settle as if the open trades had not been placed or at the price at the time of closure, with the choice of method proposed and voted on by market participants.

Due to the serious and irrevocable nature of a market closure, and the potential for reputational impact, both the majority required to close a market and the minimum participation needed will be significantly higher than for most other votes.

8.4 Parameter changes

Voting may also be used to update the configurable parameters of a Vega network. Proposals will be in the form { *parameter, proposed value, expiry time, effective time* } with the new parameter taking effect at *effective time* if the proposal is approved at *expiry time*.

[37] For example, exactly the same underlying but a different price source. Similar but different *underlying* assets, such as Brent Crude vs. WTI for oil should not be a reason for rejection.

[38] This is necessary to ensure that markets are not created due to information asymmetry if the proposer discovers an issue that should prevent the market going ahead that would be hard for others to spot, e.g. a bug in the *smart product* or an issue with an external data source.

[39] The maximum length of the proposal period is indefinite and depends on the time taken to attract sufficient market making stake and achieve the required majority of **for** votes.

9 Future work

Whilst this paper sets out a coherent and, we believe, compelling vision for a fully permissionless, trust-minimising *smart products* layer to augment the nascent decentralised financial system, it is not a complete technical specification and leaves open many implementation details and other specifics. These topics, listed below, are our focus at Vega as we work to bring the concept to life.

- **Smart product language.** The design of the smart product language and the associated risk models and framework is under active research and will be the subject of a detailed paper in due course.
- **Staking economics.** Proof of stake networks and their economics are a subject of interest and research not just within Vega but within the blockchain community in general. We will both observe and be actively involved in determining how to run a secure trading network on proof of stake.
- **Dynamic liquidity pricing.** We are actively working on a methodology to better quantify the value of liquidity in realistic market scenarios with the goal of improving on the liquidity pricing algorithm in Vega.
- **Risk and scenario testing.** In order to gain confidence in the quality and correctness of both the protocol design and our implementation of Vega, we will perform and publish the results of extensive testing, in addition to conducting further research on the risks inherent to the platform. We will also engage third parties to further audit and test our code.
- **Open source.** We plan to release the reference implementation of Vega as an open source project and transition development to a foundation over the medium to long term. We have no set timelines or further commitments at the time of writing.
- **Consensus protocol.** We are currently using Tendermint as the consensus layer and evaluating several alternatives. This process will be ongoing with the goal of improving the performance, security, and stability of Vega.

Glossary

auction call period the time during which an auction is open to receive bids and offers, after which auction uncrossing will occur to determine the trades that will be created. 15

auction uncrossing the algorithm which generates trades at the conclusion of an auction by processing, in price-time priority, the set of crossed orders that maximises traded volume. Trades occur within the price range allowed by the limit prices of the trading orders, calculated by the price determination algorithm in use (a *market parameter*). 15, 22

base currency the currency or crypto-asset in which orders are priced, and that is used for settlement and margining on any given market. 23, 25, 26, 29

coherent risk measure a function, say ρ, which assigns a real-number (the risk) to a random variable (the portfolio). For a *risk measure* to be called *coherent*, it will satisfy monotonicity, cash invariance, positive homogeneity and subadditivity. See Section 6 and [12]. 1, 10, 19, 20

continuous trading a trading mode that uses a limit order book to match trades in price-time priority. New or amended orders are immediately evaluated, and if the order crosses one or more orders on the other side (i.e. the best buy price is greater than or equal to the best sell price), trades are generated. 2, 10

discrete trading a trading mode by which prices are determined via frequent batch auctions [19]. While traders can submit orders at any time, they are held without trading as the market is constantly in an *auction call period*, with regular *auction uncrossings* occurring at the end of each discrete period. 2, 10, 23

distressed position an open position held by a trader for which the net combined value of their current margin requirements and outstanding settlement cashflows is greater than their available collateral for any given asset. See Section 5.3. 16–18, 28

indicative uncrossing price at any point during an *auction call period*, the price at which an auction would *uncross* if the auction ended at that point. 15, 22

indicative uncrossing volume at any point during an *auction call period*, the volume of trades that would be executed by *uncrossing* if the auction ended at that point. 15

initial margin the minimum amount of collateral available in a risk universe required for a new trade. See Section 6.1. 13, 20

instrument a product and all parameters required for settlement, in particular the *underlying* and *base currency*. Example: BTC/USD Dec 2019 Future. Contrast this with *tradable instrument*. See Section 3.3. 1, 2

insurance pool capital associated with a market that is used to complete settlement in the rare case of a shortfall when closing *distressed positions* in that market. Insurance pools are funded with allocated collateral that remains when closing *distressed positions*, which is the normal case, and when other markets with the same *base currency* are closed and their *insurance pools* are redistributed. See Section 6.4. 2, 13, 17, 20, 22, 23, 27

maintenance margin the minimum amount of collateral available in a risk universe required to keep trades from being forcibly closed-out. See Section 6.1. 13, 20

market maker a trader who commits to engage in market making by placing a stake on one or more markets, and receives a share of the market making rewards in that market for doing so. See Section 7.2. 1–3, 5, 10, 16, 24–28, 30

market parameter a parameter on a Vega network that applies to a given market, and may be different for each market. These are generally set at market creation and may be modified via a governance vote. See Section 3 and Section 8.4. 10, 14, 15, 29

network parameter a parameter on a Vega network that applies across the network and is the same in every market. These may be modified via a governance vote. See Section 3 and Section 8.4. 8, 22, 23, 26, 30

open market a publicly visible market on which any participant with sufficient collateral may trade or become a market maker. Open markets are created through the protocol's governance processes. See Section 3 and Section 8.2. 2, 15, 16, 22, 29

oracle a definite external source of price information or other relevant data used in the calculation of settlement cashflows and/or risk and margins. See Section 3.2. 3, 9, 24, 30

over the counter a trade that occurs between two parties without using an *open market*, for example via a request for quote or using matched trades. See Section 5.1. 2, 10, 15, 16

position resolution the methodology by which distressed positions are settled through by closing or deleveraging market positions. See Section 5.3. 2, 16–18, 21–24

product parameter a parameter required by a product in order for it to be fully specified. A product plus its required parameters is known as an *instrument*. See Section 3.2. 9, 29

proof of stake a method of securing decentralised, byzantine fault tolerant systems against sybil attacks by weighting their vote in the consensus protocol proportionally to their holding of a specific crypto-asset. 1, 23

request for quote an *over the counter* trading mode in which a participant signals to the market their interest in trading a specific *instrument* in a given size. Other participants are then free to provide a price quote for a certain period of time. At the end of the period, trades are created if the participant chooses to accept one or more of the quotes provided. 2, 10, 16

risk model a stochastic model that can calculate the required margin and any other relevant risk numbers for a given product. See Section 6.1. 1, 2, 9, 10, 18, 19, 22, 29, 30

risk parameter a parameter that is required by a risk model in order for it to be used to calculate margin and any other risk numbers for a product. An instrument with all required risk parameters is known as a *tradable instrument*. See Section 3.2. 10, 29, 30

risk universe a collection of one or more order books for related products, for which risk may be netted. For example, for futures, a risk universe may consist of order books for different maturity dates, and for European options, a risk universe could contain all order books for a given underlying and exercise date across a range of strikes. 10, 12, 13, 17, 18, 20, 22, 23, 29

sharding a method for scaling software by splitting work between multiple servers or networks ('shards'), in a way such that most work done on the network only touches one of the shards. The Vega protocol is designed to be shardable by risk universe, as each risk universe operates independently of the others, but within a risk universe the margin requirements depend on a participant's position across all the constituent markets. 8, 12

smart product a type of smart contract used to specify the behaviour of products traded on a Vega network. Smart products are written in a domain specific language (DSL) designed for financial products and specify, among other things, a product's inputs (*product parameters*), and how and when to calculate settlement cashflows. More information on smart products will be provided in a future paper to be published at vega.xyz. 1, 2, 9, 30, 31

smart product language a language for creating *smart products*. See also [7]. 2, 9, 29

tradable instrument an *instrument* that has a *risk model* and all the required parameters for risk and margin calculations specified. 10, 19, 29, 30

trading mode the set of rules in use for a given *open* or *over the counter* market, specified by a *market parameter*. See Section 5.1. 2, 10, 14, 22, 29, 30

underlying an asset, index, or other data point that gives a derivative its value. 1, 9, 10, 30

References

[1] S. Adi. How to share a secret. *Communications of the ACM*. 22(11): 612–613, 1979.

[2] D. Dolev, N. A. Lynch, S. S. Pinter, E. W. Stark, W. E. Weihl. Reaching approximate agreement in the presence of faults. *JACM*, 33(3):499-516, 1986.

[3] S. J. Grossman, M. H. Miller. Liquidity and Market Structure. *The Journal of Finance*. 43(3):617–633. 1988.

[4] J. Muranaga, M. Ohsawa. Measurement of liquidity risk in the context of market risk calculation. *Institute for Monetary and Economic Studies Bank of Japan*, 1997.

[5] P. Artzner, F. Delbaen, J.-M. Eber, D. Heath. Coherent Measures of Risk, *Mathematical Finance*, 9(3), 1999.

[6] M. Castro, B. Liskov. Practical Byzantine Fault Tolerance. *Proceedings of the Third Symposium on Operating Systems Design and Implementation*. 1999.

[7] S. P. Jones, J.-M. Eber, J. Seward. Composing contracts: an adventure in financial engineering. In: Oliveira J.N., Zave P. (eds) FME 2001: Formal Methods for Increasing Software Productivity. Springer 2001.

[8] J. Hull. *Options, Futures and Other Derivatives*. Prentice Hall, 2002.

[9] P. Glasserman. *Monte Carlo Methods in Financial Engineering*. Springer 2004.

[10] G. C. Chacko, J. W. Jurek, E. Stafford. Pricing Liquidity: The Quantity Structure of Immediacy Prices. *Harvard Business School Working Paper*. 2006.

[11] F. Delbaen, W. Schachermayer. *The mathematics of arbitrage*. Birkhäuser, 2006.

[12] H. Föllmer, A. Schied. *Convex and coherent risk measures*. https://www.math.hu-berlin.de/∼foellmer/papers/CCRM.pdf, 2008.

[13] T. Björk. Arbitrage Theory in Continous Time. Oxford University Press, 2009.

[14] D. Chaum, M. Jakobsson, R. L. Rivest, P. Y. A. Ryan, J. Benaloh, M. Kutylowski, B Adida (eds). *Towards Trustworthy Elections: New Directions in Electronic Voting*. Springer, 2010.

[15] A. Kiayias, M. Yung. Tree-Homomorphic Encryption and Scalable Hierarchical Secret-Ballot Elections. *Proceedings of the International Conference on Financial Cryptography and Data Security*. 257–271. 2010.

[16] S. V. Stoyanov, S. Rachev, B. Racheva-Iotova, F. J. Fabozzi. Fat-tailed Models for Risk Estimation.*Journal of Portfolio Management*, 37(2), 2011.

[17] ISDA. Standard Initial Margin Model for Non-Cleared Derivatives. https://www.isda.org/a/cgDDE/simm-for-non-cleared-20131210.pdf, 2013.

[18] J. Kwon. Tendermint: Consensus without Mining. https://cosmos.network/resources/whitepaper, 2014.

[19] E. Budish, P. Crampton, J. Shim. The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response. *Q. J. Econ.*, 130(4), 1547–1621, 2015.

[20] M. Broadie, Y. Du and C. C. Moallemi. Risk Estimation via Regression. *Operations Research*, 63(5), 1077–1097, 2015.

[21] K. Malinova, A. Park. Subsidizing Liquidity: The Impact of Make/Take Fees on Market Quality. *The Journal of Finance*. 70(2):509–536, 2015.

[22] M. Musiela, M. Rutkowski. *Martingale Methods in Financial Modelling*. Springer, 2015.

[23] F. Abergel, M. Anane, A. Chakraborti, A. Jedidi, I. M. Toke. *Limit Order Books*. Cambridge University Press, 2016.

[24] ISDA. SIMM Methodology, version 2.0. https://www.isda.org/a/cgDDE/simm-for-non-cleared-20131210.pdf, 2017.

[25] D. Šiška. *Risk-neutral asset pricing*. http://www.maths.ed.ac.uk/∼dsiska/RNAP-Notes.pdf, 2017.

[26] T. Zhang, L. Wang. A decentralized dark pool exchange providing atomic swaps for Ethereum-based assets and Bitcoin. *Republic Protocol white paper*. https://releases.republicprotocol.com/whitepaper/1.0.0/whitepaper_1.0.0.pdf, 2017.

[27] E. Posner, E. Weyl. *Radical Markets: Uprooting Capitalism and Democracy for a Just Society*, Princeton University Press, 2018.

[28] B. Mannerings. Vega Token: Peer-To-Peer Tradable Liquidity. *Vega white paper*, To be released 2019.

[29] D. Šiška. Market risk management on a distributed anonymous exchange. *Vega research paper*, To be released 2019.