SHARPE CAPITAL FINANCIAL MARKETS PROTOCOL & INVESTMENT PLATFORM

Whitepaper v2.1

Last updated: 2017-10-03

THE SHARPE PLATFORM TOKEN: SHP

CONTENTS

1 Introduction 5
1.1 Sharpe Capital Financial Markets Protocol 6
1.2 Sharpe Capital Investment Platform 6
1.3 Sharpe Capital Product Revenue Streams 6
2 Token Dynamics 8
2.1 SHP: Sharpe Platform Token 8
2.2 SHP Pre-Sale 9
2.3 KYC: Data Collection 10
2.4 Token Liquidity using the Bancor Network Protocol 10
3 Trustless Trading Ledger 11
3.1 Examples & Applications 12
4 Infrastructure 15
4.1 'On Chain' Features 15
4.2 'Off-Chain' Cloud Architecture 16
5 The Sharpe Investment Platform Quantitative Trading Model 20
5.1 Consistency Analysis: Data Triangulation 21
5.2 AI Portfolio Manager 22
5.3 Other Core Model Components 23
5.4 Initial ML-Driven Proof of Concept 31
6 Crowd-sourced Sentiment Platform 33
6.1 Incentive to Provide Sentiment 34
6.2 Proof-of-Reputation 34
6.3 Proof-of-Stake 35
6.4 Proof-of-Work 35
6.5 Calculating Service Fees for Sentiment Providers 35
6.6 The Sharpe Platform Application 36
7 Community and Corporate Governance 37
7.1 Democracy-based Governance: Motion Tabling and Voting . . 37
7.2 Consensus-based Community Fund Governance 38

7.3 Corporate Structure & Regulatory Oversight 39
8 Budgeting & Finance Considerations 40
9 The Sharpe Capital Team 42
9.1 Core Team 42
9.2 Advisory Board 43
10 Next Steps 46
11 Translation Notice 47
12 White Paper Purpose 47
13 Appendices 48
13.1 Backtesting Table of Trades 48
13.2 TLS Smart Contract 50

LIST OF FIGURES

Figure 1 SHP Token Exchange Dynamics 9
Figure 2 'Off-Chain' Cloud Architecture 17
Figure 3 Quantitative Trading Model Overview 21
Figure 4 Forecasting with Microeconomic Fundamentals . . . 25
Figure 5 Metaphorical Frames Impact Cognition 28
Figure 6 Network Analysis: S&P500 Network Topology 30
Figure 7 Network Clustering: Fund Diversification 30
Figure 8 2016 Q4 Backtesting Results 32
Figure 9 Proposed Budget 40

LIST OF TABLES

Table 1 Table of Pre-Sale Discounts 10
Table 2 Table 1/2 of Positions Taken by AI Portfolio Manager - 2016 Q4 Backtest 48
Table 3 Table 2/2 of Positions Taken by AI Portfolio Manager - 2016 Q4 Backtest 49

ABSTRACT

This white paper describes the *Sharpe Capital Financial Markets Protocol and Investment Platform*. The Sharpe Capital Investment Platform brings together a multitude of novel innovations in smart contracts, quantitative trading, machine learning, linguistic analysis and artificial intelligence. Principally, we are issuing Sharpe Platform Tokens (SHP). SHP provides a proof-of-stake that permits platform participants to earn service fees in ETH in exchange for providing sentiment toward global equities and blockchain assets through our web and mobile platforms. Users are rewarded with service fees in proportion to the accuracy of the sentiment they provide, utilising a proof-of-reputation mechanism.

The sentiment data collected is complemented with an array of natural language processing (NLP) strategies for performing linguistic analysis, including automated sentiment, emotional response, and contextual frame analysis. These novel approaches to understanding market dynamics, both in equity markets and blockchain assets, have been developed in collaboration with scientists at leading academic institutions. Together, these data provide a valuable source of insight for hedge funds, asset managers and private participants, and therefore a valuable revenue stream through the sale of these insights.

Sharpe Capital has developed a proprietary, automated quantitative trading algorithm driven by a hybrid machine learning and artificial intelligence model, bringing together microeconomic fundamentals, macroeconomic data, real-time world events, crowd-sourced market sentiment and NLP-driven linguistic analysis, into an overarching model capable of managing a robust, high alpha portfolio across various asset classes. Sharpe Capital will operate a proprietary investment fund operating much like an automated enhanced index fund to further generate revenue to support the SHP community economy.

The proof-of-stake metric allows us to infer the level of confidence that platform participants have in the sentiment they provide, which, when coupled with an immutable proof-of-reputation stored on the Ethereum blockchain, permits weighting of sentiment to determine both the size of service fees paid to each user, and the level of confidence to place upon each sentiment indication received. Through direct crowd-sourcing of participant sentiment, we can ensure our automated models continue to capture human, affect-driven & cognitive processes, in addition to microeconomic fundamentalist and linguistic analysis based asset value forecasting. This is unlike a prediction market - there are no losses for incorrect predictions, merely a reduction in the user's immutable reputation score, and consequently, the size of future payments. Likewise, consistently accurate users will increase their reputation, earning larger and larger payments in exchange for their insight. SHP also provides a mechanism for hedge funds and institutional participants to access our proprietary models, acting as a usage fee.

Utilising blockchain technology serves two additional purposes for Sharpe Capital: to create a decentralised, 'trustless' immutable trade ledger, such that any individual can view all our previous trades and fund performance with absolute confidence. This eliminates any dependency on trust for fund disbursement, and any possibility of fund manipulation at an institutional or individual level; to permit unrivalled community governance using both consensus-based and democratic vote-based governance models enabling the community of SHP owners to determine the direction of Sharpe Capital's future.

The Sharpe Financial Markets Protocol aims to describe a new gold standard for hedge fund management, leveraging blockchain technology to provide a low barrier to entry, continuous liquidity, anti-corruption protections, international access and optimal risk-adjusted returns. Our longer term goal, therefore, is to develop the Sharpe *Crypto-Derivative* (SCD) token, subject to necessary approvals. This first-in-class token creates the foundations for a solid link between blockchain assets and the global economy, and will ultimately provide possible payment of dividends to participants. Through our community governance structure, in which SHP holders have the right to table and vote on motions which guide the direction of Sharpe

Capital, we ultimately aim to deliver a suite of investment products with various risk profiles and across multiple asset classes, including blockchain assets. The timeline for SCD issuance is Q1 2019. An independent crowd-sale will be held for the issuance of SCD tokens, from which 100% of the proceeds will be directly invested using our proprietary trading technology.

In light of the many corrupt or ethically questionable practices in the industry, leading to disasters such as the collapse of Barings bank and the $2bn loss by UBS due to a single rogue trader, we are making the technology underpinning our Financial Markets Protocol freely available for any corporation or fund to use internationally. Our grand vision is to eliminate corruption in global financial markets while still protecting each individual corporation's proprietary information. This technology permits any institution to be instantly audited by any member of the public or regulatory body. Ultimately, the Financial Markets Protocol demonstrates how we, as a society, can eliminate financial malpractice in our lifetime through widespread adoption of Sharpe trustless ledger service technology, putting an end to 'man-made' economic disasters, and help stabilise global economies; ultimately, for the benefit of all people.

The SHP pre-sale begins on November 6th 2017 at 14:00, and the crowd-sale begins on November 13th 2017 at 14:00 UTC, with a maximum raise of $20MM equivalent in ETH. SHP provides platform participants with the right to earn ETH service fees for their market insight across many asset classes, in proportion to the amount of SHP held. SHP enables the community to vote on the creation of new funds, the development of new products, and on changes to other aspects such as service fees paid and capital allocation. SHP will be continuously liquid, directly convertible to ETH without the requirement of a counter-party.

1 INTRODUCTION

The stock market is fundamentally driven by two forces. The first is quantitative, dictated by the efficient market hypothesis, that an asset's trading market value is intrinsically linked to the microeconomic performance of the asset. The logic behind this driving force is so simple as to be almost trivial – that the collective action of every trader taking a position in an asset will encapsulate all information about that asset, thus driving it to an equilibrium reflecting its true value. However, traders do not take positions based on their belief of an asset's price. Traders do not take positions based on their belief of an asset's *future* price, either. Traders take positions based on their perception of what other traders believe about an asset's future performance. This nuance leads to the second force responsible for determining price action in stocks: participant sentiment.

Investor sentiment, whether rational or irrational, causes an asset's market trading value to deviate from its 'intrinsic' value as determined by fundamentalist microeconomic indicators. That is, trading is fundamentally driven by human decision-making based on individual traders' beliefs, perception and confidence levels. This concept, driven by concepts in social cognition and further informed by ideas from Andrew Lo's 'Adaptive Market Hypothesis', illustrates that traditional quantitative trading models *alone* are unlikely to provide a robust, high-yield portfolio. When viewed through the lens of traders' cognitive processes, behavioural observations often attributed to irrationality, such as 'loss aversion', can be seen as rational [1], learned emotive or affect-driven behaviour [2] that serves to prioritize the minimisation of losses at the cost of a proportion of potential gains [3].

Beyond the realm of fundamentalist economics, interpreting participant sentiment is rooted in human biology, described by cognitive & behavioural neuroscience, and in mathematics, described by decision theory. Investor sentiment reflects the emotive component of decision-making in trading, and causes market prices to deviate from their 'intrinsic' value according to participants' perception of what their competitors' beliefs and perceptions are of an asset. Quantitative and automated trading algorithms are not immune to sentiment, as these models are developed, parameterised and calibrated by human minds, containing implicit assumptions regarding acceptable risk profiles, implied volatility, and may contain irrational assumptions about how markets behave. To build and maintain a portfolio that is well-hedged with optimal risk-adjusted returns, both of these driving forces of price action have to be taken into consideration [4].

Therefore, our quantitative trading model aims to unite these two market forces in order to develop a truly robust portfolio management system that effectively hedges risk to provide risk-adjusted returns that consistently outperform the market. The recent leaps and bounds made in the fields of machine learning, artificial intelligence, quantitative modelling, and 'big data' create a new avenue for automated trading models that do not simply take advantage of short-term momentum in price action, opening and closing positions within minutes or even seconds, but rather, permit a sophisticated system of checks and balances, able to perceive of multiple potentialities as humans do and determine a level of confidence in a prediction and a position. The creation of such a system is the guiding principle behind our model development efforts, and the basis upon which we have established our fund.

The Sharpe Platform, supported by issuance of SHP cryptotokens, achieves this by developing models that capture the complex relationships between microeconomic asset data and its market value using state-of-the-art machine learning technology, together with the ability to assess its own perceptions of current market behaviour as a whole through independent incorporation of two sources of participant sentiment analysis. We are able to obtain a measure of participant sentiment using a proof-of-stake and proof-of-reputation crowdsourcing system that rewards users for providing useful sentiment by making service fee payments in ETH, exchanging their insight for cryptocurrency. SHP serves an additional, important role in providing access to our proprietary sentiment data and trading platform, which utilises SHP to charge fees for access to our 'off-chain' cloud-based components, similar to the concept of 'gas' on the Ethereum platform.

1.1 Sharpe Capital Financial Markets Protocol

The Sharpe Capital Financial Markets Protocol aims to be the new gold standard for hedge funds, providing for an entirely decentralised, community-governed organisation in which key governance decisions are made by the community through both consensus-based and democratic votes, managed entirely using blockchain technology in a trustless manner. The protocol includes the Sharpe Trustless Ledger Service (TLS), described in Section 3, and our community governance model, described in Section 7.

1.2 Sharpe Capital Investment Platform

The Sharpe Capital Investment Platform consists of a new form of quantitative trading model utilising the latest advancements in economic theory and cognitive science, integrated with advanced statistical and machine learning analysis of microeconomic, macroeconomic, participant sentiment and news/social media emotional content analysis. This permits the development of a high 'alpha' fund aiming to achieve significantly better risk-adjusted returns than tracker funds. The sections predominantly concerned with describing the Investment Platform are section 5, discussing development of the quantitative trading model itself, and section 6, describing how we will crowd-source sentiment using the Ethereum blockchain, and paying ETH to contributors in exchange for their insight.

1.3 Sharpe Capital Product Revenue Streams

Across the Sharpe Capital Investment Platform and Financial Markets Protocol, there are 5 key revenue streams, for redistribution to the community of Sharpe Platform users, and ultimately, for the issuance of securitized tokens. The key revenue streams for Sharpe Capital are as follows:

1. Return on investment from operating the Sharpe Capital Proprietary Investment Fund (Section 5).
2. Sale of crowd-sourced sentiment data-feed, obtained using a proof-of-stake, reputation and work, reward-based system (Section 6).
3. Profits derived from access to real-time data-feeds from our unique approach to linguistic analysis incorporating the latest insights from cognitive behavioural science, psychology, and decision theory (Section 5).

4. Consultancy fees for implementation of internal auditing tools and private enterprise-grade blockchain solutions to hedge funds and corporate clients, based on our 'Trustless Ledger Service' (Section 3).
5. Further diversification as permitted under our consensus-based and democratic vote-based 'Community Governance' model, outlined in Section 7.

2 TOKEN DYNAMICS

This section introduces the issuance & distribution dynamics of the Sharpe Platform Token, SHP. The main crowd-sale event will take place on November 13, 2017 at 14:00 UTC for the purchase of SHP. The total supply of SHP in circulation will remain constant following the initial crowd-sale: there will be no further minting (creation) of SHP once the token generation event ends.

2.1 SHP: Sharpe Platform Token

The Sharpe Platform Token, SHP, will be issued during our token crowd-sale event, starting on November 6, 2017, and available to purchase at a discount our pre-sale period, starting on November 6th, 2017 at 14:00 UTC. SHP will serve the primary purpose of supporting the Sharpe Capital crowd-sourced sentiment platform, described in Section 6. Owners of SHP will have the right, but not obligation, to provide sentiment analysis about a variety of financial instruments, via our web and mobile applications. Service fees will be paid to sentiment holders in ETH, in proportion to the quality of the sentiment they provide and the amount of SHP owned (Proof-of-Stake). This will be automatically managed by smart contracts. Additionally, the core Sharpe Capital investment platform will expose a set of APIs and services for analysing and developing quantitative trading models, which we plan to sell to the hedge fund industry in exchange for SHP.

In summary, SHP confers the following core rights to their owners:

- Access the many services and APIs that comprise the Sharpe Capital investment platform, for the purpose of analysing and developing quantitative trading models, paying access fees using SHP.
- Create and vote on motions regarding the creation of new funds, capital allocation, service fee payments through our community governance structure.
- Earn service fees, distributed in ETH, in exchange for market sentiment toward global equities and blockchain assets, provided via our web and mobile applications.

2.1.1 *SHP Exchange Dynamics: Crowd-sale*

Sharpe Platform Tokens (SHP) will be issued at the rate of 2,000 SHP per 1 ETH. This rate will remain constant throughout the duration of the token generation event, or 'crowd-sale'. Upon completion of the crowd-sale, no more SHP will be minted. The hard-cap for SHP generation, including pre-sale, is $20,000,000, to be pegged to ETH at a time prior to launch of the pre-sale event on November 6th. The public crowd-sale will begin on November 13th at 14:00 UTC. We have implemented 'dynamic ceilings' as developed by Status GmbH[1] to ensure decentralised distribution of SHP (maximum contribution per person not to be in excess of $75,000 during the crowd-sale). This is important to the integrity of the voting and consensus-based community governance models discussed in Section 7.

Further to the 2,000 SHP issued to the contributor per 1 ETH, an additional 2,000 SHP will be held in reserve for future fund raising, and 1,000

1 Readers are referred to the article by Status on 'dynamic ceilings' at https://goo.gl/pxwKaU

SHP distributed to Sharpe Capital and community members to fund future platform development. The reserve funds will be non-transferable for 6 months, and subsequently dynamically vest over a period of 18 months. We will be utilising reserves via the Bancor Protocol to provide continuous liquidity for both the purchase and sale of SHP, as discussed in section 2.4 below.

Figure 1: Sharpe Platform Token exchange dynamics: token sale participants will receive 2,000 SHP per ETH, with 1,000 allocated to Sharpe Capital, and 2,000 in a vesting reserve fund. This results in a total of 5,000 SHP being minted per ETH received during the token sale period

SHP will be distributed to three Ethereum addresses upon receipt of Ether; these are defined below:

- Contributor address (2,000 SHP per ETH)
- Trustee smart contract address (2,900 SHP per ETH)
- Sharpe Multisig smart contract address (100 SHP per ETH)

The Trustee smart contract is used to manage vesting grants. ETH with vesting periods applied is sent to the Trustee smart contract. Grants can be set up, which permit the recipient to claim ETH according to a defined cliff and vesting duration. The following vesting grants are applied to the Sharpe Capital Trustee smart contract (on a per ETH basis):

- Team & community: 1,000 SHP - 6 month cliff & 2 year vesting period
- Reserve fund: 1,900 SHP - 6 month cliff & 2 year vesting period

An allocation of 100 SHP per ETH will be immediately liquid to pay for bounty services received by Sharpe Capital prior to the crowd-sale. This SHP will be sent to the Sharpe multi-sig smart contract address.

2.2 SHP Pre-Sale

The SHP pre-sale will go live 7 days before the SHP crowd-sale. This will be open to participants making sizeable contributions, and discounts will be offered in proportion to the size of the contribution being made. The pre-sale will start on Monday 6th November 2017 at 14:00 UTC, with the discounts described in Table 1 automatically applied.

Payments of ETH must be submitted during our pre-sale window and SHP will be issued when the crowd-sale goes live. White-list registered contributors will be able to participate for the first 48 hours of the sale only.

Table 1: Table of Pre-Sale Discounts

Deposit (USD Nominal)		
Minimum	Maximum	Token Discount
10k	50k	10%
50k	250k	20%
250k	500k	30%

90% of ETH received during the pre-sale will be locked in a multi-signatory smart contract address until the token crowd-sale concludes, and 10% will be available upon conclusion of the pre-sale to permit rapid mobilisation prior to platform launch on 11th December 2017. We will enforce a hard-cap on the pre-sale of $8m USD, to be pegged to ETH prior to November 6th 2017 and published on the Sharpe Capital website. The maximum contribution from an individual participant during the pre-sale is $500k equivalent.

2.3 KYC: Data Collection

Prior to contribution, we will require participants to provide a name, address and telephone number. Participants will be asked to confirm they are not resident in Singapore or China. Additionally if participants are US citizens, or their primary domicile is within the US, they will be asked to confirm they are accredited participants.

2.4 Token Liquidity using the Bancor Network Protocol

We will issue tokens leveraging the Bancor Network to provide instant liquidity to SHP, without requiring the use of traditional exchange-based mechanisms. We are currently in discussions with the Bancor team, working to determine the optimal approach to determine the optimal currency reserve, virtual reserve and constant reserve ratio values to strike the ideal balance between liquidity and permitting SHP price movement via the Bancor protocol. Section 8 describes our budget allocations, including a 10% reserve allocated for continuous liquidity and user service fees. We are in the process of developing a Monte Carlo trading simulator with the aim of determining the optimal allocation of reserve funds and liquidity rate for SHP. This will take the form of a dynamic function, with the final value dependent on the raise achieved upon conclusion of the SHP crowd-sale.

The Bancor Protocol provides convertibility and built-in liquidity to tokens issued through the Bancor Network utilising a unique price discovery mechanism. This is achieved through the use of 'reserve' currency balances, permitting smart tokens, such as SHP, to be traded without requiring a counterparty, an exchange, or any of the risks associated with these token exchange mechanisms. The price discovery mechanism ensures a 'network effect' among tokens traded through the Bancor Protocol, collectively supporting the value of all tokens utilising the network. Read more about the Bancor Network Protocol at https://www.bancor.network/.

3 TRUSTLESS TRADING LEDGER

Smart Contract Enforced Transparency to Ensure Investor Confidence

A core component of the Sharpe Capital Financial Markets Protocol is our Trustless Ledger Service for investment and asset trading (TLS). This underpins our investment activity in global equity markets, and creates an immutable record of our profit and loss account in real-time. Our ledger is implemented with an Ethereum smart contract, which is detailed in this section and visible on our public GitHub page [2]. The Ethereum blockchain-hosted TLS may be adopted by anybody through a suite of methods in the smart contract, opening it up for use by any organisation involved in asset management and investment. We have extended use of the public TLS free of charge to drive further transparency in this sector, and expect that crowd-sale participants in other investment-related blockchain funds, and eventually participants in funds generally, will begin to demand its utilisation by other blockchain-related investment funds.

The key mechanism in maintaining a trading ledger service that is trustless, incorruptible and protects sensitive data, such as currently open positions, is the addition of trades to the public ledger when they are opened, with sensitive position details encrypted via RSA public key cryptography. The following information is encrypted upon addition to the public ledger:

1. Open Price
2. Stop Price
3. Limit Price
4. Ticker Symbol
5. Stock Exchange Symbol

The profit and loss of each position is updated daily, providing an up-to-date account of all investments in our portfolio, immutably recorded on the blockchain for anyone to see. When a position is closed, the RSA key pair used to encrypt the sensitive information will be released to the blockchain, such that anyone viewing our public record can decrypt the information and verify its integrity at the time of publishing. Current and historical equity points and leverage utilisation is recorded unencrypted to permit real-time monitoring of fund performance by any individual. If we were to edit position details after we publish them to the TLS, whether at an institutional level or due to a 'rogue' team member, this would be visible to anyone viewing the Ethereum blockchain, and the world would see any malpractice taking place. Maintaining a ledger of this nature creates a *truly* trustless relationship between the Sharpe Capital proprietary fund and its stakeholders. This trustlessness provides an important element to the intrinsic value of Sharpe Capital-issued tokens, as all participants can be entirely certain that we are operating legitimately and always in their best interests.

In light of such scandals and fraudulent activity, regulators, policy makers, researchers and economists have proposed various models for reducing risk. For example, Chang *et al* (2012) [5] proposed internal computerization systems for derivative management. As an alternative approach, Hornuf *et al* (2014) [6, 7] suggested that 'behavioural designs' could serve to mitigate risk through creation of social and physical environments that diminish propensity to commit fraudulent acts. The TLS proposed by Sharpe Capital incorporates elements from such proposals, relying on a blockchain-driven

2 https://github.com/sharpe-capital

computerised system to eliminate the possibility of undetected fraud, with the *public* aspect of the ledger enforcing social conditions not conducive to even attempting such acts of fraud and deception.

A detailed description of our TLS implementation can be found in Section 13.2.

3.1 Examples & Applications

To better illustrate the Platform of the Sharpe TLS, we will consider a series of real-life use cases and provide examples of how the application of this technology would have prevented various scandals and crises within financial institutions. We will consider examples of fraud and malpractice that have occurred in the past, for example, the collapse of Barings Bank [8], and we'll look at ethically questionable practises that are still in operation today [9]. Ultimately, we present a convincing argument that our platform would have prevented some of the biggest scandals in finance in the last 30 years, and paves the way toward improved ethical and moral standards within the finance sector in the future.

These use cases illustrate just a handful of near innumerable scandals in the sector, their prevalence being so significant that Wexler (2010) proposed a theory grounded in social psychology to explain the abundance of rogue trading activity in financial institutions, whereby individuals are incentivised by their environments to commit such acts [10]. The following use cases describe how such incentives are eliminated by our trustless ledger service.

Use Case 1: *Nick Leeson & Barings Bank*

Arguably one of the biggest scandals of the 20th Century was the collapse of Barings Bank – established in 1762 and brought to its knees by the unauthorised actions of a single individual. Barings Bank collapsed in 1995 as a result of unauthorised trading by Nick Leeson [11]. He was tasked with arbitraging futures, seeking to profit from differences in prices of the Nikkei 225 futures contract across the Osaka Securities Exchange and the Singapore International Monetary Exchange. However, instead of buying on one market and selling immediately on another, he held onto positions gambling on the future direction of the markets, against the knowledge of his superiors. By the time the bank collapsed, Leeson's unauthorised trading had amassed losses of £827 million (GBP), twice the bank's available trading capital. Leeson managed to hide the losses from his superiors in London by altering the branch's error account – known infamously as the "five-eights account".

It is simple to demonstrate how an immutable, decentralised trading ledger, accessible and verifiable by the entire world, on demand via a dedicated web platform, would have prevented this catastrophe from happening. Leeson would have been required to record his 'buy' orders with one exchange, and 'sell' orders with another, on the public trading ledger in real-time for anyone to see. He would only be required to release information about the trades post-execution, as all sensitive information would be encrypted. This would have rendered the dishonest maintenance of the '88888' account impossible, as accounting and reconciliation would have been decentralised and immune from unseen manipulation. Simply one anomaly in the accounting of his practise would have been flagged up to superiors within a matter of days of occurring. Instead Leeson managed to keep up his dishonest actions for a staggering duration of over two-years [12].

This is used as an extreme example, as it ultimately ended in the demise of a bank that had been in existence for over 2 centuries, but this is a problem that has yet to be solved by the capital markets sector. Only six years ago, in 2011, UBS announced that it had lost $2 billion (USD) in a scandal involving Kweku Adoboli's unauthorised trading [13].

Use Case 2: Dark Pools

Dark pools are private forums in which institutional participants buy and sell securities without recording transaction details until a while after the trade has been fulfilled. These are a form of market manipulation beneficial to institutional and accredited participants within the pool, as it permits execution of large securities transactions with a delayed impact on market movements [14]. However, this is to the detriment of the market as a whole, because order flow information is hidden and the market is no longer transparent – prices of execution are kept secret until much later, leading to increased market inefficiency and confounding the price discovery process [15].

The Sharpe Capital Financial Markets Protocol TLS represents a big leap forward in transparency within financial markets, and sets an excellent example of ethical practises and honesty within the sector. All trade execution prices are logged immutably and immediately on our public trading ledger in real-time, such that any changes in the future would be known to the entire world. This information is considered to be sensitive and encrypted whilst the position is 'open'. However, once the position is closed and the decryption keys are released automatically through a smart contract, observers of the Ethereum blockchain would be able to verify that the price of execution is consistent with public records for the security in question at the time of the trade. This discourages the use of dark pools as more funds and institutions adopt our platform, and the transparency it enforces becomes widespread throughout the financial services industry.

Any attempted use of dark pools by Sharpe Capital, or other adopters of our trustless ledger technology, is rendered impractical as it would be disadvantageous to log different prices on the ledger to those obtained within the dark pool, which cannot be modified or 'hacked' by the institution in question due to the immutable nature of blockchain records. Such discrepancies would be very quickly flagged by the trustless ledger platform, alerting participants to the manipulation.

Use Case 3: Insider Trading

This section addresses a longer-term goal of the trustless ledger service, achievable once a critical mass of adoption has taken place. Financial regulatory bodies such as the SEC (United States) or FCA (United Kingdom) would have a single point-of-entry to analyse trading patterns in real-time without any requirement for institutions to reveal sensitive information such as their currently open positions.

While certainly an ambitious proposal, we expect that clients and participants of the capital markets sector will ultimately demand the adoption of trustless ledger technology, as there are no good arguments *not* to do so. Indeed, as blockchain technology becomes better understood and embraced by regulatory authorities, it could one day be mandatory. As all sensitive and proprietary information is cryptographically protected, the only motivation to not adopt a trustless system would be to participate in fraudulent, manipulative or otherwise unethical trading practices that ultimately act to the detriment of retail participants and clients of financial institutions.

Use Case 4: *Private TLS for Internal Auditing*

In addition to the Sharpe TLS being available for investment companies, it is also possible for organisations such as hedge funds to adopt a *private* implementation for internal auditing and governance purposes. To raise additional capital for development of the Sharpe Capital core products and funds, we will offer private blockchain implementations for financial institutions with customised web-based front-ends with automated analytics and account reconciliation.

4 INFRASTRUCTURE

A key feature that makes the Sharpe Capital digital token sale unique, when compared to similar companies utilising the 'ICO' model ('Initial Coin Offerings') of the past year which generate revenue, in part, through proprietary investment activities, is our novel combination of off-chain and on-chain systems and architecture, redefining how capital investment funds should operate. We recognise that the power of blockchain-based applications is the trustlessness provided through decentralisation and having no 'central authority'. However, we also recognise that to truly leverage the power of the blockchain and decentralisation, start-ups must learn how to combine their smart contract implementations with more traditional "off-chain" systems and applications.

It is widely accepted that code must be succinct and efficient to be suitable for execution on a decentralised blockchain. The first implementation of blockchain technology had a very simple purpose – Bitcoin simply stores an immutable ledger of financial transactions with no central authority. With the advent of Ethereum, we have the power to perform more complex computation, and build sophisticated applications that leverage the decentralised nature of blockchain technology. However, we find ourselves constantly grappling with the trade-off between the power we can leverage through the fundamental benefits of blockchain technology (namely, trustlessness, immutability and decentralisation) and the significant cost of executing more complicated applications and algorithms on a distributed blockchain as compared to the utilisation of cloud architecture, such as the Amazon Web Service cloud platform.

At Sharpe Capital, we have been faced with particularly complex challenges in this area. The nature of our modelling and investment algorithms, is such that we require immense amounts of computational power to analyse millions of data points in order to determine optimal fund allocation through our predictive modelling tools. This cannot be efficiently performed on a blockchain application – least of all Ethereum, as the transaction costs would diminish or even negate any returns. This presented us with an interesting challenge, to determine which features of our investment platform would be suitable for smart contracts, and which should be built as traditional cloud-based applications. Integrating extremely intensive computational models with blockchain technology in a trustless, decentralised manner therefore presents particularly complex challenges.

The following section will provide an overview of the applications built 'on-chain' (executed on the Ethereum blockchain using Solidity) versus off-chain (implemented on our proprietary cloud architecture, with a 'trustless link' to on-chain activity), discussing the rationale behind our decisions and the benefits offered to the users of our platform. Finally, we will conclude with a high-level explanation of the verification process we plan to implement, which will prove that off-chain activity has occurred as stated – *beyond a reasonable level of doubt.*

4.1 'On Chain' Features

Token Minting & Distribution

As described in Section 2, the minting, distribution and transferring of tokens will be executed on the Ethereum blockchain.

Trustless Trading Ledger: Enforcing Honesty & Ensuring Investor Confidence

Our trustless trading ledger (TLS), described in Section 3, is built using an Ethereum smart contract. The service is available for use by anyone with access to the Ethereum blockchain through execution of smart contract methods, either manually as described in Section 13.2 or through a web-based tool presently under development. The TLS is not limited to proving the 'trustlessness' of just the Sharpe Capital fund - it is available to anyone who wishes to provide a trustless investment fund to both protect against 'rogue' activity and inspire confidence in participants. We will register trades on our public ledger from micro-services running within our cloud-hosted, off-chain infrastructure. We are developing a public 'ledger explorer', permitting any internet user to monitor, audit and verify the activity occurring on our TLS with ease.

Sentiment Crowdsourcing Platform

The crowd-sourced sentiment platform (described in Section 6) and reputation scoring that feeds into our investment modelling, and provides service fees to SHP holders in exchange for accurate sentiment indications, is implemented and maintained through a smart contract on the Ethereum blockchain as well. This creates an immutable track record of the sentiment provided by a specific SHP holder, such that reputation scores cannot be manipulated to 'trick' our reward-system into paying excess or undue service fees, nor is it possible for service fees due to *not* be paid - this too is enforced by the smart contract. As with the TLS, we are building a web-based platform and mobile app for the sentiment analysis & reputation scoring platform.

4.2 'Off-Chain' Cloud Architecture

We have chosen to leverage the power of Amazon Web Services (AWS) for our off-chain architecture. Amazon provides a large number of fully-managed, scalable, cost-effective services, enabling us to maintain a low-cost base and operate in highly agile manner. A key distinguishing factor between Sharpe Capital and many traditional investment funds, aside from our novel approach to quantitative, sentiment & cognitive-behavioural modelling of market movement, and our TLS, is that we will constantly strive to operate as leanly as possible. This tech-driven, highly cost effective operating model, combined with consistent market-beating returns, permits us to deliver the highest return possible to our participants, service fees to sentiment providers, and cost savings for users of our bespoke modelling tools (wherein usage costs are shared in SHP, operating in a similar manner to 'gas' in Ethereum for transaction processing/computation execution). The AWS cloud architecture we are utilising is outlined in Figure 2.

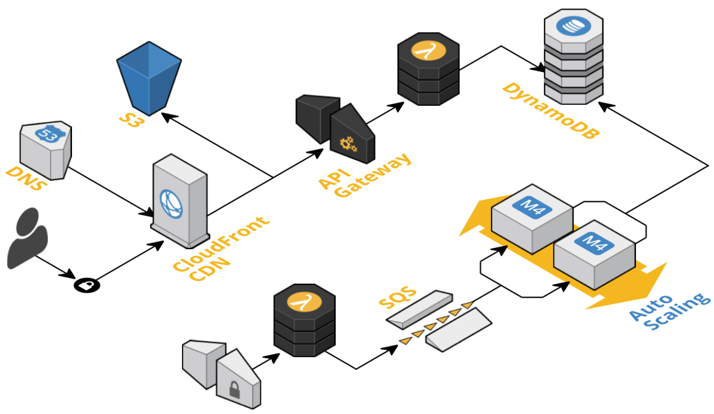

Figure 2: Overview of our highly cost efficient cloud architecture utilising AWS to provide sufficient computational resources to execute our complex trading models.

Auto-scaling Groups

Auto-scaling groups are clusters of servers, which share similar characteristics for the purposes of scaling and management. It is possible to scale groups according to time-schedules, memory utilization, CPU consumption and much more. They can be provisioned & scaled on-demand and destroyed when no longer needed. We are using auto-scaling groups to ensure our internal systems can handle an effectively infinite load, and to provide our financial modelling with extremely powerful yet highly cost-efficient computational resources.

Big Data Storage & Management: NOSQL Databases

The volume of data required for the Sharpe Capital investment modelling significantly exceeds amounts that can be cost-efficiently stored in traditional relational databases (RDBMS) such as MySQL. AWS provide a fully-managed NoSQL database called DynamoDB, which is low-cost, infinitely scalable, fully-redundant and extremely fast. We are leverage DynamoDB to store all of our historical macro and microeconomic data, going back over 25 years for 9,000 US assets. Currently we're storing approximately 25GB of data, with up to 10 million records per table. This represents only ~1% of the data we plan to store, and will need to efficiently query, when the Sharpe Capital Investment Platform is fully operational. By setting the right foundations from the start, we will avoid complex data migrations at a later date. As we expand our Investment Platform, we will incorporate data from all major asset-classes globally.

Asynchronous Messaging

We're using the Amazon Simple Queue Service (SQS) to transport data between our cloud-hosted microservices. SQS is a fully-managed message queueing system that makes it easy to decouple and scale applications and distributed systems. We run a number of important background processes to ensure our historical data is kept up-to-date, execute and test investment models and monitor our live investment strategies. Most of these processes can execute independently and asynchronously: SQS enables us to build highly scalable, massively parallel, non-blocking applications that satisfy our back-office requirements.

Serverless Architecture

Our externally hosted services, such as the Sharpe Capital public website and platform API are hosted using 'serverless' architecture. Serverless architecture enables microservices to be 'spun-up' on demand, meaning we only pay to host our applications when there is demand for them. Amazon monitors HTTP traffic and only deploys our applications for the duration needed to service incoming requests.

Ensuring Trustless Off-Chain Activity

Essential to the operation of our proprietary investment fund, is the fact that distribution of service fees to SHP holders, and eventually dividends to holders of our crypto-derivative tokens, are trustless and immune from manipulation. This is partially solved by logging our trades immutably on the Ethereum blockchain in real-time. However, because many of our investments are largely focussed on global equity markets, we will need to liquidate the ETH raised during crowd-sale to USD and other 'functional currency' (fiat) to carry out our investment activity. This creates an undesirable level of trust that the balances held in "off-chain" accounts are currently intact, and the funds have not been misappropriated. It would be possible for the Sharpe team to liquidate to USD, log a record of "fake" investments on the TLS, and take-off with the funds. Therefore, a trustless solution is necessary.

This trust can only be eliminated by verifying that 'off-chain' balances are consistent with the information stored on the blockchain. Let's assume that we hold $10m USD of ETH in 'Account A' (on the Ethereum blockchain) and $20m USD 'off-chain' in a bank account for investment purposes. After three months, we've recorded 20 open positions on the TLS, with a total leverage of 5x and collateral of $10m (at 10% margin). This means we can freely transfer $10m between functional currency and cryptocurrency, with $10m effectively locked away off-chain for our open positions (trades). Clearly, with traditional investment funds (and various blockchain-driven funds), there is a large degree of trust involved - the arbiter of 'off-chain capital' effectively operating as a Central Authority. We propose to remove the trust created in this situation by periodically moving the 'free' capital that is declared on our TLS into ETH and back out again. This will prove to the community that the available funds reconcile with the information stored on the ledger, and that they were available at any given point in time. As we begin to liquidate open positions, our profit and loss account will change, this will subsequently cause the capital to fluctuate over time, steadily increasing as we begin to turn a profit. This information will be stored on the blockchain, and the transactions of our 'free' capital in-and-out of the ETH will consistently reconcile with the TLS data.

This solution does not come without its drawbacks. Moving funds from 'functional currency' to ETH, and back again, will diminish returns slightly through the cost of 'gas' and exchange transaction fees. The proposed transactions will be large, and we will negotiate favourable rates with a reliable exchange. This will mitigate to some extent the effect it has on ROI of the fund. To mitigate this reduction in ROI, we will carry out this reconciliation action once per quarter, thereby proving the fund status and activity described on the TLS is genuine. We will include a mechanism in our smart contracts, permitting SHP holders to table a motion to carry out this reconciliation more frequently should additional 'proof-of-capital' be required. If the community votes with a simple majority to approve the ad-hoc reconciliation request, then we will reconcile the funds as described up to a maximum of once per month. We believe this approach balances our good intent to prove that we are using the 'off-chain' funds in an honest and transparent way, with the interests of the community held in the highest regard, to build and maintain 'trust-through-trustlessness' with Sharpe Capital token owners at all times.

5 THE SHARPE INVESTMENT PLATFORM QUANTITATIVE TRADING MODEL

This section provides an overview of the quantitative modelling approach Sharpe Capital will utilise to determine investments. We discuss the grounding of our model in economic theory and cognitive sciences, and provide a brief overview of the core technologies we are utilising and how we are combining them to derive accurate price forecasts and asset data streams. This platform was conceived by Sharpe Capital's Chief Investment Officer, James A. Butler PhD in collaboration with CEO Lewis M. Barber, with additional contributions across our core team and advisory board. J. A. Butler holds a position as Research Fellow at the University of Oxford, and has 7 years of doctoral and post-doctoral research in delivering novel approaches to the development of robust, fit-for-purpose predictive models at the intersection of complex systems analysis, traditional statistics, data science, and machine learning. Further development of the core concept, particularly relating to the AI Portfolio Manager, was contributed by L. M. Barber, who holds extensive experience consulting for global investment banks in developing fault-tolerant engineering solutions. The development of novel approaches to linguistic analysis, that goes beyond traditional 'sentiment analysis' but looks to evaluate expected 'participant response' to consumed media at a cognitive-behavioural level, is supported by leading linguist and anthropologist, Mieke Vandenbroucke PhD, a Fulbright Scholar and Visiting Researcher at the University of California, Berkeley. An overview of the entire quantitative trading model is provided in Figure 3; each of the components of this model are discussed in more detail below.

As introduced in Section 1, the investment strategy utilised by the model is underpinned by the notion that asset prices are driven by two forces. The first of these is the tendency for asset prices to move towards their 'intrinsic' value through the equilibria achieved from supply-demand curves, based on the collective opinions of very many traders or participants that, by and large, have access to the same microeconomic data relating to each asset. This market force is essentially that described by the Efficient Market Hypothesis [16], however, the efficient market hypothesis is not the whole story [17]: it fails to take participant sentiment, emotional responses, and decision theory, into account - these effects result in a market that is not entirely efficient, thus permitting utilisation of the understanding of these effects in investment strategies to develop an 'edge' over the market as a whole. The second of these two forces is therefore essentially an acknowledgement of the role of cognitive bias, behavioural psychology, emotion and sentiment in human decision-making. The latter force causes asset prices to fluctuate and is largely responsible for observed volatility in the stock market. As described by Shull (2011) [4], the "real game" in trading is not investing based on one's beliefs of an asset's price at some point in the future, but on one's beliefs of what other other traders, on average, believe the market will do. Belief formation has been shown to have a crucial role in market behaviour, and research shows that established beliefs in investment are resistant to change following reception of new data [18], leading to confirmatory bias. To put it another way, traders are constantly attempting to pre-empt each others' decisions, perceptions and beliefs when taking positions in the market, leading to volatility and short-term 'irrational' deviations in asset prices.

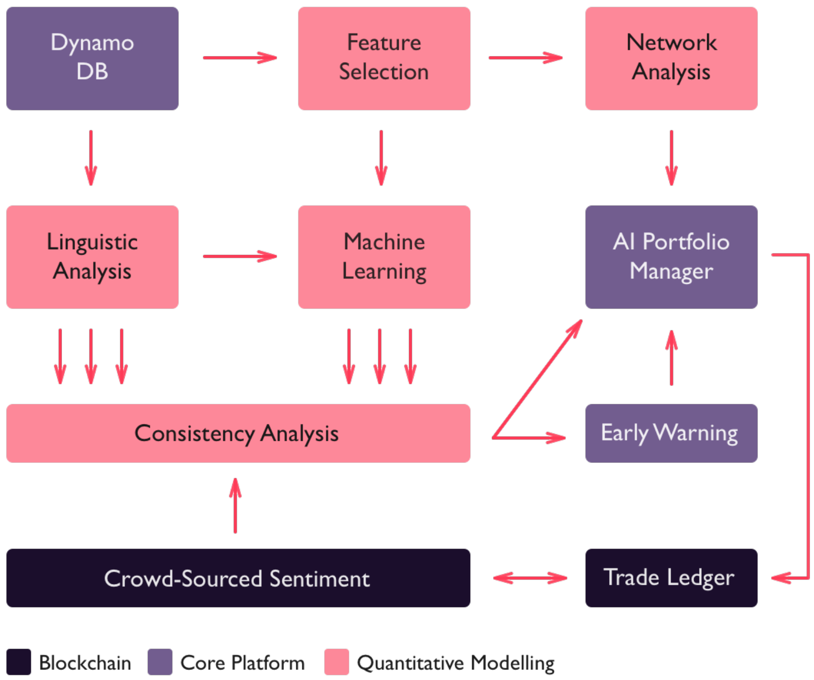

Figure 3: The Sharpe Capital quantitative trading model: a high-level summary of our next-generation model that incorporates fundamentalist analysis with multiple machine learning models, automated linguistic analysis (sentiment, emotion, contextual framing), integrated using our proprietary 'consistency analysis' algorithm to provide both forecasts and confidence levels. These are subsequently evaluated by an AI-driven portfolio manager to determine the optimal positions to take given available capital, using network analysis to ensure proper diversification across dissimilar assets.

It should be noted before moving forward that, in parts, this section is somewhat technical at times for the benefit of the interested reader. While generally described at a high level, aspects of this description are targeted at those with a background in investments, quantitative modelling and complex systems analysis.

5.1 Consistency Analysis: Data Triangulation

The core of the automated trading model we propose in this section is 'consistency analysis'. In essence, this takes forecasts derived from multiple data sources and multiple approaches to modelling price changes based on those data, evaluating the consistency of these predictions for each asset. The data sources utilised are principally:

- Microeconomic indicators, presently obtained from quarterly reports.
- Macroeconomic indicators, such as commodity prices, employment levels, currency exchange rates and market sector indices.
- Sentiment derived from analysis of traditional news and social media channels using natural language processing (NLP), leveraging insights into language perception from the latest research in linguistics and cognition.

- Crowd-sourced sentiment obtained via the Ethereum blockchain, for which we will pay participants service fees based on the accuracy of their indications

The resulting triangulation of these data sources allows the model to provide a measure of 'confidence' in each forecast made. That is, if notably different outcomes are predicted by each individual approach, relative to the differences of the set of predictions for all other assets, the consistency analysis will indicate a 'low' *relative confidence score* (RCS). If the multiple data sources and prediction methods are much more aligned, the analysis will indicate a 'high' RCS.

5.1.1 *MiDAS: Manifold Driven Asset Scoring*

The consistency analysis described in this section is a proprietary development of Sharpe Capital, the following represents a brief overview of the methodology utilised to produce an 'RCS' (relative confidence score), and discusses the reasoning behind its utilisation. The aim of consistency analysis is, in essence, to produce a number between 0 and 1 for every asset forecast produced by the model; this is the RCS. In order to determine *relative* confidence, the consistency of predictions for a specific asset must be considered relative to the consistency of predictions for all other assets. The dataset that the consistency analysis operates over can be described as an $n * m$ matrix where n are assets and m are predictions. Each prediction for each asset may then be represented as data points within an m-dimensional manifold. Use of dimensionality reduction techniques allow the relative similarity of predictions to be visualised in a 2D plane. Clustering analysis performed on this 2D representation of the manifold then permits us to rank the similarity of predictions for each asset based on how well clustered they are. Feature scaling can then be applied to convert these rankings into an RCS. We can further enhance the Platform of RCS through integration of our linguistic analysis of media sources, including: sentiment analysis, blockchain-driven crowd-sourced sourced sentiment, emotional response analysis, and contextual frame analysis to predict participant response to consumed media. This can provide not only a measure of confidence in the prediction, but an indication of whether a prediction is more likely to be too low or too high, effectively resulting in a prediction, a confidence score, and a 'direction' of prediction for each asset in the model.

The RCS can be viewed as an abstraction of participant 'gut instinct' or 'anxiety' with respect to specific forecasts, essentially emotional responses that drive participant decision-making. It is now well established that emotion, or affect, are a fundamental part of the human decision-making process [19], and that these serve to protect participants against loss [2, 20]. Therefore, through the MiDAS algorithm, we are able to recreate models of cognitive processes that shape human participants' decision-making.

5.2 AI Portfolio Manager

The artificially intelligent portfolio manager is responsible for managing the execution of trades in global equity markets. Our quantitative modelling ultimately provides a list of highly accurate future price predictions for thousands of assets over multiple time horizons. The AI portfolio manager (PM) optimises the execution of trades based on these predictions, whilst minimising risk and net exposure in order to maximise risk-adjusted returns.

In addition to future priced predictions, the portfolio manager relies on forward-looking volatility (also known as implied volatility), which is derived from current option prices over different time horizons. This volatility gives our PM a decent understanding of the magnitude of price movements to be expected on an asset over a specific future time period. To explain how this is useful information for our portfolio manager, consider the following example:

Let us assume today is July 1, 2016 and AAPL is trading at $95.98 with a 30-day implied volatility of 10%. Our PM has a price prediction for AAPL of $130.55 on September 10, 2016. A quick reconciliation of the risk-reward ratio on this position suggests that a good trading opportunity has been identified. The maximum downside over the next 30 days can be approximated at 10%, whilst the upside is over 30%. If the market were to correct within the 10% implied volatility first, we would still be looking a highly profitable trade, with good risk-reward, provided our price forecast is accurate to within a reasonable margin of error. As MiDAS provides a confidence score for predictions, we can be sure to preferentially use predictions with the model's confidence (RCS) in the prediction providing an indication.

Our AI portfolio manager constantly analyses implied volatility, current price and price predictions for thousands of assets every day, determining how to structure the most well-balanced portfolio of long & short positions, based on the confidence score of each individual price prediction and the size of the risk-reward ratio that has been identified.

5.3 Other Core Model Components

This section outlines the role of the additional components of the quantitative trading model outlined in Figure 3 in improving forecast accuracy, reducing model over-fitting, and ensuring forecasts are most appropriately leveraged using our artificially intelligent portfolio manager and network analysis.

Machine Learning

Machine learning (ML) at its most fundamental level, can be described as computer pattern recognition and classification algorithms that improve as they are exposed to further examples, in a process known as 'training' - these are known supervised ML approaches. That is to say, the algorithms 'learn'. More recent developments in 'unsupervised' machine learning aim to stratify data based on their relative similarity without a 'training' requirement: a key example being the sub-field of 'manifold learning', which contributes largely to the design of the MiDAS consistency analysis briefly outlined in Section 5.1.1. To describe the vast number of approaches, algorithms and strategies that now exist in the ML field is beyond the scope of this document. Therefore, this section is primarily concerned with evidencing previous success in applications of ML to market trading, and presenting an argument grounded in economic and scientific theory that leverages these existing tools in the manner described in Figure 3. We do plan, however, on preparing a series of blog posts describing our progress in this area and providing further insight into Sharpe's proprietary approaches to the maximum extent possible without infringing on the core intellectual property that makes the Investment Platform possible. Some of the technologies and algorithms that are key to the ML aspect of the quantitative trading model described herein include data flow graphs (conceptually similar to artificial neural networks) using the TensorFlow library [21], adaptive boosting [22], perceptron maps [23], and manifold learning-based techniques such as the t-stochastic nearest neighbours algorithm (tSNE) [24, 25]. In addition, to optimise model parameters, both for ML models and for the AI Portfolio Manager, we are developing multi-objective optimisation methods such as 'evolutionary algorithms' that aim to identify optimal values in a model to maximise a given set of 'fitness functions' [26, 27] (in our case, the fitness functions would primarily be most accurate forecasts and most effective use of these forecasts in terms of maintaining high alpha/maximising risk-adjusted returns).

The beneficial application of ML strategies to financial trading has long since been established, applied to optimising trade execution [28, 29], or in more specialised applications including Forex trading [30], futures trading [31], bankruptcy prediction [32], stock price forecasting [33], and capital asset pricing models [34]. There are innumerable examples in the scientific literature of machine learning models out-performing 'traditional' regression models, and certainly many additional proprietary, unpublished ML-driven quantitative trading models utilised by investment funds globally.

ML-driven strategies for asset pricing, forecasting and automated trading have been developed in accordance with various underlying economic strategies. These include principally models utilising 'technical analysis' (also known as 'chartism' [35], which posits that asset price trends tend to follow repeatable and therefore identifiable patterns and often pays little-to-no regard to underlying economic fundamentals [36], and models employing 'fundamentalist' analysis [37] that focus on economic indicators relating to the assets underlying performance as opposed to previous price action. Predictions derived from a chartist viewpoint are typically over very short time frames, with very many positions being opened and closed each day, often being open for mere minutes. The nature of fundamentalist analysis, assessing the longer term viability of an asset based on its business performance and other macroeconomic factors, is inherently more forward looking and returns greater ROI per trade [38]. The process of selecting the most appropriate subset of data to utilise in a ML-driven model is known as feature selection, this is a key process in ensuring that models do not 'overfit' to 'training' data and are capable of appropriately generalising to unseen datasets [39]. We leverage a variety of statistical and theoretical methods to filter, combine and transform 'feature spaces' for the optimisation of our models.

Sharpe Capital's position on these somewhat opposing market views is essentially that over longer time horizons, an assets price will generally reflect the microeconomic performance of that asset, and that short-term trends or 'momentum' in price action arises from what has been traditionally been called '(irrational) participant sentiment'. However, based on the various studies in cognitive science, decision theory, behavioural psychology, etc., outlined within this paper, we maintain that these local deviations in price from an asset's intrinsic value is not a result of (just) irrationality, but is, fundamentally, due to human cognitive processes in decision-making[3]. With that in mind, the Sharpe Capital Investment Platform does not incorporate historical price trends into its forecasting or trading algorithms. Instead, we opt to combine ML strategies founded upon the principles of fundamentalist analysis at the microeconomic and macroeconomic scales to forecast as-

3 Perhaps the best non-technical literature espousing this market view is Shull (2011) [4]

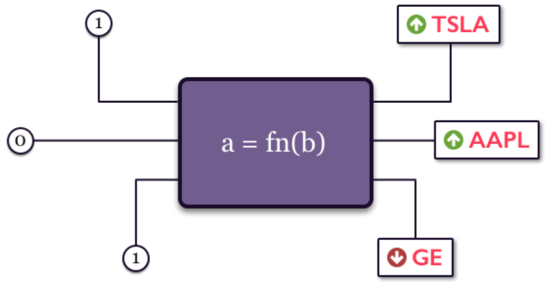

Figure 4: Machine Learning models can be employed to understand the relationship between economic indicators input into the model, identifying persistent relationships between an assets historical microeconomic indicators and asset price in the following quarter.

set price performance based on underlying economic indicators, augmented with natural language processing (NLP) and ML-driven linguistic analyses that aim to predict participant sentiment, emotional response, and cognitive-behavioural processes that are ultimately responsible for bullish or bearish trends, and examples of what would previously have been labelled as 'irrational' behaviour within the market.

Our ML-based approaches to asset price evaluation are not restricted to any specific technique. In our previous experience developing ML applications not just in finance, but for the biotechnology and pharmaceutical sectors, the best results are derived from employing a variety of techniques in a complementary fashion. This philosophy is what ultimately drove development of the overarching platform outlined in Figure 3. In that high-level schematic, multiple arrows can be seen to leave the 'Machine Learning' component, feeding into 'Consistency Analysis'. This reflects the utilisation of multiple modelling methodologies which utilise different feature sets (asset microeconomic and macroeconomic variables). Our research in this area to date suggests that simply averaging over multiple, distinct, ML-grounded asset forecasting models serves to reduce the mean average error in future price predictions/ This is because, in situations where one model predicts a price that is too high, we often find that alternative models under-predicted. Through the implementation of MiDAS and consistency analysis, we are seeking to take this several steps further, not merely taking average values, but evaluating how the similarity, and accuracy, of different forecasting approaches changes over time. By further incorporating results from automated, NLP-based linguistic analyses, discussed below, we expect to significantly improve the robustness of our models. Furthermore, as is described in the 'Early Warning System' section, this combination of multiple approaches (some well-established in the scientific literature, some cutting edge, and some entirely novel) permits a degree of 'market perception' within the model, enabling it to alert human analysts when it detects anomalous inconsistencies that it is unable to reconcile into a reliable trading decision.

Finally, it is worth noting that regional markets behave differently, particularly for emerging versus developed economies, in that the relationship between microeconomic fundamentals and stock market asset prices does not necessarily hold between exchanges in different countries [34, 40, 41, 42]. Our models are designed to take these differences into account, both through adaptation of the algorithms used based on observed differences in economic literature, and by ensuring that distinct markets and asset classes are separately processed when training models. We will also apply an integrative approach that feeds into the 'consistency analysis' to take advantage of observed commonalities [43] across international markets.

Linguistic Analysis

This section outlines the key forms of linguistic analysis that we are developing to complement our economic fundamentalist approach to quantitative trading models outlined above. The purpose of including linguistic analysis of news and social media feeds is two-fold:

- To anticipate market shocks that microeconomic indicators would not reveal in a time-efficient manner.
- To determine market movements driven by participant sentiment, emotional response, and the contextual framing that shapes participant cognitive processes, in order to pre-empt deviations of an asset's market value from its 'intrinsic' fundamentalism-informed value.

Sentiment Analysis

Sentiment analysis describes a wide range of techniques that utilise NLP, ML and rules-based algorithms to infer 'sentiment' in social media (e.g. Twitter, Facebook forum discussions,), traditional news media (e.g. newspapers) and 'new' media (e.g. blog posts). There are a wide variety of existing sentiment analysis libraries available [44], and there exists significant scientific literature on the successful application of these for anticipating market movements [45, 46, 47].

Emotional Content Analysis

Emotional content analysis, as its name implies, seeks to infer the emotional content of texts. It has been established that emotional response drives decision-making [19]. It therefore follows that decisions can, to some extent, be inferred from the likely emotional response following media consumption in relation to a potential investment or trade. The principles underpinning 'neuromarketing' are largely based on determining emotional responses to advertisements [48, 49], optimising marketing campaigns to produce the desired emotional response in a target audience. A corollary of these findings is that prediction of emotional response from language use in social and other media can form the basis for predicting the internal emotional state of participants and could be predictive of the decisions they subsequently take.

Emotional content analysis differs from sentiment analysis in that the seeks to discover viewpoints or opinions generally from a text, whereas the latter seeks only to infer the emotions expressed by the text. We believe that in combination, these two techniques have the capacity to significantly improve the accuracy of well-established sentiment analysis methodologies. Much of the scientific literature on this area is based on neuromarketing, however some research does exist in understanding the role of emotion and the stock market [50]. This is an emerging research area; MIT have very recently developed an algorithm capable of inferring emotional content (specifically, sentiment, emotion and sarcasm) of twitter messages utilising artificial neural network-based 'deep learning' approaches [51].

Contextual Frame Analysis: Predicting Investor Decision Making

Contextual frame analysis, as defined here, is a concept developed by Sharpe Capital that aims to provide an enormous leap forward in understanding participant decision-making at a cognitive level, going far beyond negative/positive word association-based sentiment analysis and emotional content analysis. This section provides a very brief background on frame analysis in linguistic theory, and goes on to describe a novel hypothesis for predicting, and even pre-empting, participant decisions.

Frames are cognitive constructs that relates networks of related concepts or ideas. Frame analysis can therefore be defined as an approach to cognition that focuses upon understanding the organisation of experience [52]. In the sub-field of frame semantics, it is often that the meaning of a word to a reader can only be understood in the context of how the word forms part of a cognitive conceptual system [53]. The concept of NLP-driven 'contextual frame analysis' is a hypothesis developed based on existing linguistic research into frame analysis and how frames shape cognitive processes. The justification for developing such an analytic method is driven by a desire to improve the accuracy of existing sentiment analysis, by instead attempting to infer participant response based on contextual frames contained within media consumed when investment decisions are made, such that we may predict how an participant will respond to market sentiment at a cognitive level. There is no existing literature on development of such a process nor its application to investment activity. This section briefly outlines the core concept underlying this hypothesis. Sharpe Capital is in the process of creating a proof-of-concept for the development application of contextual frame analysis, which will eventually be shared on our blog and included in a later version of this document.

In seminal work by Thibodeau and Boroditsky (2011) [54], it was demonstrated that 'metaphor frames' shape the way people think at a pre-conscious level. To very briefly summarise this work, it was found that changing one word in an article about a fictional crime wave drastically changed the manner in which two test groups would deal with the perpetrators of these crimes. Specifically, one group received an article in which the criminal activity was described as a 'beast' and in the other, the word was changed to 'virus'. A majority of participants exposed to the 'beast' metaphor voted for 'enforcement' (punitive) measures, while a similar majority of those exposed to the 'virus' metaphor instead opted for 'reformation' (rehabilitation). When asked to justify these reasons, justifications were provided from other elements of the article. The metaphor used was not mentioned by any participants. That is, the groups on average *came to opposing conclusions and believed they did so due to their exposure to data, despite both groups being exposed to identical articles, other than the one word change*. When the metaphor frame was removed, and participants were instead exposed to a 'lexical prime', in which words relating either to beasts (such as 'caging' or 'invader') or viruses (such as 'sickness', 'infection') in a context not relating to the article, these observed differences in opinion disappeared entirely. The results of this experiment can be seen in Figure 5. The conclusion therefore, is that the different frames contained within the articles had a major effect in shaping the decisions made by the participants, and that they were unaware of the manipulation - justifying their decisions after the fact.

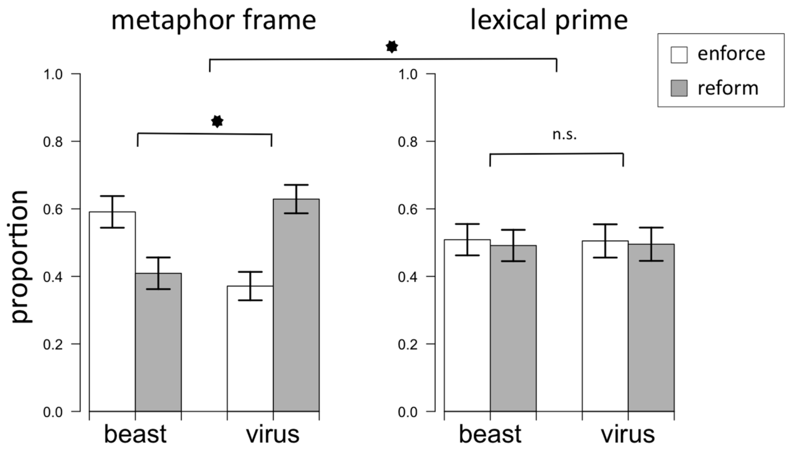

Figure 5: Data taken from [54]

A logical extension of this conclusion is that contextual frames contained within media articles and corporate reports relating to stocks and shares are likely to have a similar impact on participants or traders that consumed them. Identifying the frames contained within news articles, training an ML algorithm on price trends in the time following article release, with a sufficiently large number of previously published articles, should permit identification of 'predicted participant response' due to the presence of contextual frames within the media they consume. That is, in a manner similar to identifying 'reform' or 'enforce' in Figure 5, we aim to identify 'long' and 'short' responses to contextual frames within the financial corpus. We are working to utilise NLP techniques to isolate and identify frames, and deep learning strategies to correlate these to market outcomes. This would permit us to go far beyond the realm of sentiment analysis, and understand participant decisions at a fundamental, cognitive level. If our evidence-based hypothesis proves correct, this approach will offer unrivalled insights into market movements. We are working with leading linguists and computer scientists to develop this highly novel approach to la

Early-Warning System

It is an inevitability that circumstances relating to an asset will sometimes change following our forecasting and taking of a position on that asset. This could be due to microeconomic factors, such as reputation damage, public relations issues, revelation of a scandal, or pre-report announcements of disappointing profits. There could additionally be market shocks affecting positions taken due to macroeconomic events; for example, the collapse of mortgage-backed securities in 2008.

As described in Section 5.1.1, our model has an abstraction of cognitive and affect-driven processes that drive participant decision-making, defining a measure of relative confidence in predictions utilising 'manifold learning' techniques - defined as the relative confidence score, or RCS, generated for each forecast made for each asset. We have developed the concept of an 'Early Warning System' (EWS). The EWS monitors the change (ΔRCS), and rate of change (differential with respect to time) of the RCS ($\frac{\delta RCS}{\delta t}$), detecting rapid changes in the consistency of predictions driven by sudden shifts in sentiment, emotional content and contextual frames detected by our linguistic analysis model. For 'long' positions, when the RCS reaches a minimum threshold, or the differential reaches a certain negative level, the EWS will instruct AI Portfolio Manager to close the position if it can do so without making a loss. In the event that it cannot, the EWS will alert all Sharpe Capital executives and analysts via SMS (cell phone) alerts, push notifications and email, permitting manual intervention. For 'short' positions, the opposite applies and the EWS monitors for increases in the RCS or its differential.

When an analyst deems manual intervention is necessary, several options are open to them:

- Close the position immediately.
- Monitor the position and close it at a later date.
- Choose to take an opposing position (e.g. go 'short' on a previously 'long' position).
- Allow the AI Portfolio Manager to continue normal operation.
- Utilise options, futures, or other instruments to hedge against additional risk identified.

Improving AI Portfolio Manager by Training on Human Responses

To improve the ability of AI Portfolio Manager to handle exceptional or changing circumstances, we will develop a graph-based data flow model (much like an Artificial Neural Network), that 'trains' on human activity taken in response to changes in the RCS. Once sufficient experience has been gathered, the AI Portfolio Manager will be equipped to deal with many difficult circumstances under its own volition.

Crowd-sourced Sentiment Platform

The crowd-sourced sentiment platform, described in Section 6, will provide a valuable additional input into our forecasting model. By obtaining real sentiment indications from users of the platform, using a combination proof-of-reputation and proof-of-stake system to determine 'confidence weightings' to be applied to predictions, we are able to augment our model with genuine human insight into the market. Essentially, this allows us to add 'organic' intelligence into an artificial intelligence-driven quantitative trading model. Contributors of sentiment are paid a service fee based on the accuracy of the sentiment they provide and its utilisation in the Sharpe Capital investment portfolio.

Network Analysis

Network analysis allows us to visualise the market in terms of similarity between assets. To generate such a network, we first compute a 'distance matrix' from all of the microeconomic fundamentals for each asset, providing a measure of distance or 'similarity' between each asset. This matrix contains each asset in each column and row, with the distance score between each asset indicating their similarity. From this, we can construct a network, illustrated for S&P500 assets in 2015, as demonstrated in Figure 6.

We can project any microeconomic indicator onto each node (asset) by scaling their size and colour accordingly, to aid in visualisation. The Platform of network analysis is derived from applying clustering analysis to this network, as demonstrated for the 2015 S&P500 assets in Figure 7, in which

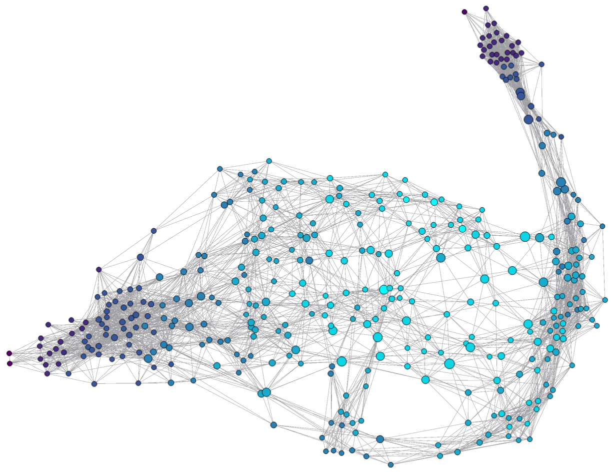

Figure 6: Network Analysis: S&P500 Network Topology:Each node (circle) is a specific asset, and each edge (line) describes a connection between them based on the distance matrix, and boundary conditions on generation of the network (number edges, number of neighbours, distance metric used, etc.)

each node identified as belonging to a specific cluster, or 'microeconomic sector' is coloured appropriately. This algorithm groups assets together based on how similar they are at a fundamental level, rather than relying on arbitrary sector labels applied based on the industry of the asset. This is a much more dynamic approach than diversification across sectors in the traditional manner, allowing our diversification strategy to evolve as assets diverge or converge in their similarity. Spreading our positions across dissimilar assets partially insulates against sector-based market shocks and corrections, and furthermore, by identifying under-performing and over-performing sectors with respect to the market as a whole, the AI Portfolio Manager is able to intelligently and dynamically determine the best assets to go long and short on.

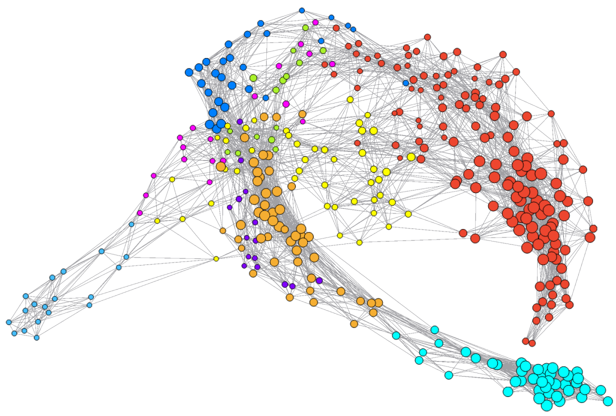

Figure 7: Network Clustering: Clustering based on asset similarity permits funds to be distributed optimally based on a dynamic, constantly updating, identification of sectors and their relative performance.

Outputs from the Network Analysis module of the Sharpe Capital Investment fund are fed into the AI Portfolio Manager as illustrated in Figure 3. This permits it to best leverage predictions it receives from machine learning and linguistic analysis, via consistency analysis (MiDAS), selecting assets based not only on risk-reward ratio (based on expected returns given predicted growth and implied volatility) but based on the optimal allocation amongst network analysis-derived microeconomic-inferred sectors.

5.4 Initial ML-Driven Proof of Concept

Much of the Sharpe Capital Investment Fund is currently under development. In particular, the linguistic analysis and automated implementation of MiDAS is still in progress. As described in Section 8, a proportion of funds from the token crowd-sale have been allocated to permit development and deployment of an operational automated quantitative trading model by Q1 2019. We have, however, implemented a proof-of-concept using a restricted subset of the tools and techniques described in this section. In particular, we used a hybrid approach combining adaptive boosting [55] and other 'non-linear regression' techniques, network analysis, statistical analysis for feature selection, and a simplified implementation of AI Portfolio Manager that takes only 'long' and 'short' positions. AI Portfolio Manager is being extended to be capable of utilising options and other more complex financial instruments, but has been restricted to taking long and short positions only during the prototyping phase.

With this proof-of-concept, we built a machine learning model that trains over microeconomic and macroeconomic indicators from the first three quarters of any year to forecast prices on the release date of each assets Q4 quarterly report, restricted for prototyping purposes to assets in the S&P500, reducing computational requirements. We cross-validated this model to develop an early understanding of how well it would generalise to other datasets, and built a proprietary back-testing algorithm to evaluate the performance of this model. The model was permitted to utilise leverage dynamically to the extent the AI Portfolio Manager judged it safe to do so. The equity curve when back-testing this model over 2016 performance produced the following results, shown in Figure 8. Each position taken is listed in a table that can be found in the Appendix (Section 13.1).

We are aware of over-fitting issues with machine learning-driven quantitative trading models. The consistency analysis and MiDAS algorithm described above is designed to mitigate over-fitting by combining many different modelling techniques and data sources. Furthermore, by ensuring that the AI Portfolio Manager properly diversifies assets according to their microeconomic similarity, and appropriately hedges positions using a combination of 'long' and 'short' instruments, we can effectively insulate the fund against a degree of uncertainty. The model that resulted in the performance shown in Figure 8 had a mean average error (MAE) of approximately $10 USD in either direction, and still provided an annualised ROI of 85% thanks to this strategy of utilising distinct ML and AI components in co-operation. Incorporating linguistic analysis, crowd-sourced sentiment analysis, and more sophisticated network analysis methods into the trading platform, utilising consistency analysis to evaluate confidence levels, with expansion to many more assets, is expected to improve model performance significantly. It must however be stated that back-tests nor past performance

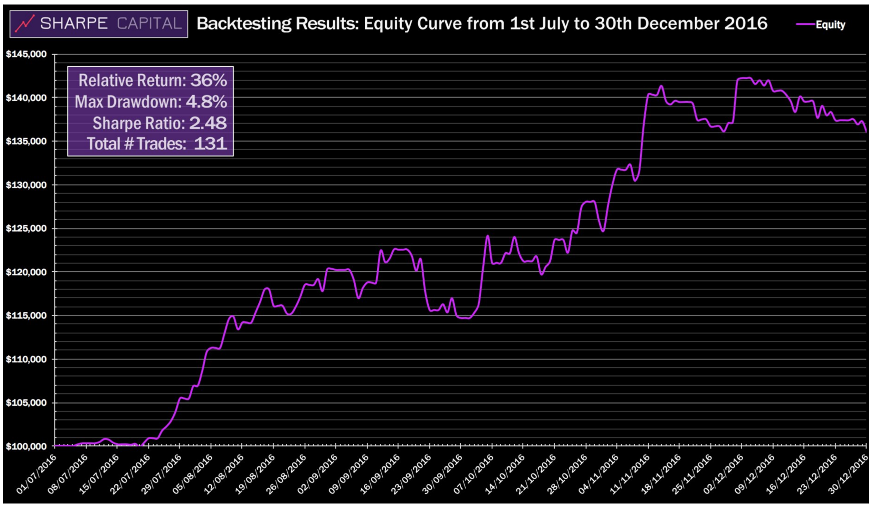

Figure 8: Proof-of-concept ML model (without sentiment or other linguistic analysis), showing the fund equity curve over a six month period. The risk-adjusted return from this period annualises to 84%. Utilising four models, each designed to predict a specific quarter from the previous three, will likely improve this ROI further.

are necessarily indicative of future performance, this in large part depends on general movements in the markets.

6 CROWD-SOURCED SENTIMENT PLATFORM

This section describes how we will crowd-source real-time sentiment from SHP owners utilising a proof-of-stake, proof-of-reputation, and proof-of-work system that allows us to determine the weight to give to each prediction when incorporating user sentiment indications into the quantitative model described in Section 5. This platform provides an additional revenue stream as we will sell access to the raw sentiment data for use by hedge funds and other interested parties. 40% of profits derived from this will be distributed into our reserves, from which service fees will be distributed to users on a quarterly basis. This will comprise a 20% ETH reserve and a 20% USD(T) reserve. The remaining 60% will be used to drive further growth of Sharpe Capital according to the budget outlined in Section 8.

The origins of crowd-sourced sentiment date back to the earliest days of hedge funds. The first hedge fund was created in 1949 by A. W. Jones, whose contributions to the field foreshadowed modern portfolio theory. Jones successfully 'hedged' his positions by taking long and short positions to insulate his fund, and utilised leverage (borrowed capital) to increase his returns. Arguably one of his greatest achievements, however, was foreseeing what cognitive science and behavioural psychology would only come to understand half a century later: that participant or trader *emotions* created trends in asset price action [4]. Jones argued that price increases generated optimism (or confidence) in an asset, fuelling further price increases, resulting in a positive feedback loop causing an asset's price to deviate from its intrinsic value, such that these trends could be identified from following price action [56]. Jones quickly realised that to determine the best long and short positions, he should not rely on chart watching alone, but should source sentiment from those with a stake in market performance. Therefore, he created a system whereby he invited brokers to select their favourite long and short positions, simulating running live portfolios from these stock picks. Jones utilised these predictions to select stocks and positions for his hedge fund, rewarding brokers in proportion to the value of the predictions they provided. Through this method, Jones was able to achieve substantially improved returns on his investments.

With the advent of blockchain technology and a large user base interested in the stock market, we are now in a position to leverage this technology to replicate Jones' method on a large scale. We are developing a web and mobile platform that allows SHP owners to indicate their current sentiment toward a vast number of assets. Users of the platform will be paid service fees with ETH for correct sentiment indications (or 'predictions'), using smart contracts to enforce this automatically.

The following sections describe how we use proof-of-stake and proof-of-reputation to ensure users providing sentiment are invested in making successful predictions, and proof-of-work to provide a disincentive to users from creating new accounts and transferring SHP to new Ethereum addresses if their reputation level falls below the initial reputation score of new users, 0.5. ETH service fees paid are a function of the number of correct sentiment indications provided, the users' 'reputation score' and the amount of participation with the platform the user provides. These are known as the 'three proofs', and are stored on the blockchain associated with the user's Ethereum address: proof-of-reputation, proof-of-stake, and proof-of-work.

6.1 Incentive to Provide Sentiment

Crowd-sourcing of high-quality asset sentiment is crucial to the performance of the Sharpe Capital proprietary investment fund. Quantitative modelling is useful for predicting asset prices up-to a certain point, beyond which we must rely on the opinions of 'the crowd' to identify and explain anomalies. In order to ensure the sentiment analysis that we rely on is of the highest possible quality, we have designed a system of incentivisation, in which holders of SHP can build up a 'reputation score' (R) based on the quality of the asset sentiment they provide on our platform. Their reputation score is maintained on the Ethereum blockchain using our smart contracts and is linked to an Ethereum address. Bi-annual payments of ETH will be distributed to providers of sentiment, in proportion to their reputation score, and the amount of SHP they own. This immutable track-record of reputation scores, linked to a proof-of-stake model for SHP, creates a mechanism that provides the highest possible incentive to provide the very best sentiment, and thus the most predictive of future asset prices. It is not possible to 'lose' SHP, and therefore this is *not* a form of 'prediction market' or other form of gambling.

6.2 Proof-of-Reputation

It is important to maintain a high reputation score by providing the best forward looking asset sentiment possible, as this will be used to determine ETH service fees owed to platform users. The track-record for each Ethereum address is immutably stored on the blockchain. The reputation score may be calculated according to the following formula, presently planned to be updated on a per prediction basis. The quarterly ETH service fee payments will use the user's reputation score as it is *at the time the service fees are paid*.

$$R_{\text{current}} = \left(R_{\text{previous}} + \left(\sum_{n=0}^{p_{\text{correct}}} R_{\text{previous}} \cdot m \right) \right) - \left(\sum_{n=0}^{p_{\text{incorrect}}} R_{\text{previous}} \cdot m \right) \quad (1)$$

$$0 \leqslant m \leqslant 0.1 \quad (2)$$

$$\text{Reputation s.t.} \quad 0 \leqslant R \leqslant 1 \quad (3)$$

In equation 1, R_{current} is the user's current reputation score, calculated weekly based on how accurate the provided sentiment was. R is linearly modulated by a scalar constant, m (equation 2) positively for each correct prediction and negatively for each incorrect prediction. That is, for each correct prediction the reputation score is increased or decreased by multiplying the user's previous reputation score, R_{previous}, by m, and adding this to R_{current}. This is repeated for incorrect predictions, but by reducing R by $R \cdot m$ instead of increasing it. The maximum value of R is 1, and the minimum value is 0, as stated in equation 3.

The impact of predictions on a user's reputation score is perhaps best demonstrated by example (let us assume that $m = 0.05$):

- A user has $R = 0.5$, and provides 6 correct and 3 incorrect predictions. Their new R is $(0.5 + (6 * 0.5 * 0.05)) - (3 * 0.5 * 0.05) = 0.575$. This represents an *increase* of 7.5 percentage points for all service fees received by this user for correct predictions.

- A user has $R = 0.2$, and provides 6 correct and 3 incorrect predictions. Their new R is $(0.2 + (6 * 0.2 * 0.05)) - (3 * 0.2 * 0.05) = 0.23$. This represents an *increase* of 3 percentage points for all service fees received by this user for correct predictions.
- A user has $R = 0.5$, and provides 3 correct and 6 incorrect predictions. Their new R is $(0.5 + (3 * 0.5 * 0.05)) - (6 * 0.5 * 0.05) = 0.425$. This represents a *decrease* of 7.5 percentage points for all service fees received by this user for correct predictions.
- A user has $R = 0.2$, and provides 3 correct and 6 incorrect predictions. Their new R is $(0.2 + (6 * 0.2 * 0.05)) - (3 * 0.2 * 0.05) = 0.17$. This represents a *decrease* of 3 percentage points for all service fees received by this user for correct predictions.

As the example above illustrates, changes in a user's R are linear - that is, the relative change in R is dependent only on the accuracy of the user's provided sentiment, and not the previous value of R.

6.3 Proof-of-Stake

Requiring ownership of SHP for sentiment provision is important because it indicates a level of confidence in the reputation score associated with a specific Ethereum address. We call this a 'lossless' proof-of-stake, as the quantity of SHP owned cannot diminish due to incorrect sentiment indications (or 'predictions'). If an address is holding a lot of SHP, with a high reputation score, then it's reasonable to assume the forward-looking asset sentiment has a higher probability of being of correct.

6.4 Proof-of-Work

A simple proof-of-work scheme is utilised for two purposes. Firstly, to incentivise continued participation with the platform, and secondly, to remove incentive from users whose reputation score, R, has dropped below the starting point of 0.5 from transferring SHP to a new Ethereum address and creating a new account on the Sharpe crowd-sourced sentiment platform in an attempt to 'game' the system and 'reset' their reputation. A user's proof-of-work score is multiplied by the service fees due, calculated as described in Section 6.5.

Users will begin with a proof-of-work score, W_u, of 0.5. This will increase as a function of the number of predictions provided and their regularity. We are in the process of evaluating various proof-of-work mechanisms to determine the most optimal incentivisation mechanism. The precise time horizons and number of trades required is yet to be determined through simulation of platform use, and will further be refined upon the alpha (testnet) and beta releases of the crowd-sourced sentiment platform.

6.5 Calculating Service Fees for Sentiment Providers

Users of the Sharpe sentiment crowd-sourcing platform will receive a service fee for every correct sentiment indication that they provide. These will accrue quarterly, recorded against the user's Ethereum address on the blockchain, and will be distributed automatically by smart contract from our Ether/liquidity reserve pool in a trustless manner. Service fees will begin accruing from the moment the web application is available on the Ethereum mainnet. We will allocate a discretionary service fee pool each quarter, and the entirety of this will be distributed to users in proportion to the relative value of their insight compared to all other users - calculated as a function of number of correct predictions, the user's reputation score, the user's proof-of-work score (level of participation), and the user's proof-of-stake (amount of SHP owned). This provides an incentive to improve accuracy by more carefully considering future predictions, creating an incentivisation feedback loop that will improve the quality of sentiment provided over time. The formula to calculate payments due to users is described in Equation 4, based on the number of correct sentiment indications ('predictions')and the 'three proofs' captured in Equation 5:

$$\text{User Service Fee} = \frac{\Gamma_u}{\sum_{n=1}^{N} \Gamma_n} \cdot Q \tag{4}$$

$$\Gamma_u = P_c \cdot R_{(u,t)} \cdot W_{(u,t)} \cdot \frac{S_u}{S_T} \tag{5}$$

Equation 5 defines a user's overall 'success' metric, which captures their accuracy, reputation, work and stake, which we call their Γ or 'gamma' score. P_c is the number of correct sentiment indications provided by the user in the previous quarter, $R_{(u,t)}$ is their reputation score as defined in Section 6.2, $W_{(u,t)}$ is their proof-of-work score, and $\frac{S_u}{S_T}$ is the proof-of-stake, defined as the proportion of SHP held compared to the total amount of SHP used to provide sentiment.

From the user's Γ, we can now determine the proportion of the service fee pool, Q, they are entitled to. This is described in Equation 4, in which the proportion of a user's Γ score is compared to the cumulative total Γ of all users. When the Γ is added together for every participant, it will always equal 1, and therefore, the entirety of the service fee pool, Q, is distributed amongst participants in each quarter. Service fees are distributed trustlessly via smart contracts on the Ethereum blockchain, from an allocation within the Sharpe Capital ETH liquidity reserves.

Service Fee Summary

The service fee distribution mechanism described above is designed to ensure platform participants are rewarded fairly for the sentiment they provide. Through our unique combination of the three proofs of reputation, work, and stake, we can ensure that holders of large proofs-of-stake (SHP) do not receive service fees disproportionate to their participation nor do they dilute the fees paid to hard working participants with a smaller stake. The proof-of-work requirement disincentives attempts to 'game' the system by using multiple accounts, as any increase achieved in reputation through these means will be offset by the reduction in the proof-of-work associated with that address.

6.6 The Sharpe Platform Application

We have published designs for the Sharpe Platform mobile application, including the crowd-sourced sentiment platform, the trustless trading ledger described in Section 3, and the community governance structure described in Section 7. These can be found at `https://goo.gl/f3sqyf`.

7 COMMUNITY AND CORPORATE GOVERNANCE

This section introduces a core component of the Sharpe Capital Financial Markets Protocol, describing how the community of SHP owners are able to govern management of Sharpe Capital in a decentralised, trustless manner via blockchain technology. Firstly, we utilise a 'consensus-based' model, in which new fund creation, or other ventures proposed by Sharpe Capital, are determined by the will of SHP owners, such that capital allocation for such ventures is dictated by the degree of community consensus. Secondly, we have developed a 'democracy-based' model in which any owner of SHP may table a motion for approval by a majority of SHP holders. We have developed a novel solution that permits voting without 'locking' SHP during voting periods through creation of single-use 'vote tokens', solving a key problem in previous attempts at creating decentralised autonomous organisations (e.g. [57]).

In conjunction with the trustless ledger service outlined in Section 3, this forms a new gold standard for management, in which community participants are able to dictate Sharpe Capital's operational direction, and in other matters relating to governance of Sharpe Capital's ventures, with enforcement managed through smart contracts on the Ethereum blockchain. All of the governance structures outlined herein will be in place prior to the SHP crowd-sale event, ensuring that the community can help shape the future of Sharpe Capital from day one.

7.1 Democracy-based Governance: Motion Tabling and Voting

The principal method for community governance is the consensus-based model described in Section 7.2. In addition to this, we are implementing a 'majority-rule' system for community governance relating to matters beyond new fund creation. This is achieved through the following mechanism:

1. An SHP owner tables a motion utilising our 'Democratic Community Governance' smart contract, for which a user interface will be developed.
2. This motion must be seconded by at least 10% of SHP within 7 days. Sharpe Capital and its employees are forbidden from taking part in community votes, but may table motions.
3. Once the motion has been seconded, it goes to a full vote open for 7 days.
4. A simple majority of SHP must vote in favour of the motion for it to succeed.
5. The motion will be executed by Sharpe Capital.
6. The Sharpe Capital Advisory Board will act as an intermediary with the power to delay implementation of the motion, permitting additional time for the core team to engage the community on the issues at hand.
7. If the Advisory Board reaches a majority consensus (excluding abstentions) against the motion, in the event that a motion is deemed infeasible or would potentially cause harm to the community. A position paper will subsequently be published to further engage the community prior to a second round of voting. This governance model is adopted from a parliamentary model of governance.

8. Motions may be rejected by unanimous consensus (excluding abstentions) of the Advisory Board if a second voting round passes, and the Advisory Board deem that the motion would cause significant harm to the community of SHP owners.

Sharpe Capital owned reserve SHP will not be permitted to participate in votes. Submitted motions will be subject to review before any voting may take place. Illegal or frivolous motions will be rejected and not displayed to the community, as will motions that would cause the Board of Directors to violate their fiduciary responsibility to Sharpe Capital.

7.1.1 *Collaborative Development of Democratic Governance*

The proposed democracy-based governance model described above should be considered prototypical of our aims with regard to community governance. We will actively engage with the community and continue to work with our legal counsel to ensure that the final implemented structure reflects the values of the SHP community, and is not subject to inadequate definitions that may lead to difficult-to-foresee 'edge cases'.

Voting Without Restricting Token Transfer

In previously proposed voting mechanisms, such as that proposed by the DAO, a significant issue was that during voting periods, participants were unable to transfer their tokens as this could permit a user to vote multiple times by moving SHP between new accounts and voting from them. This provided an incentive for users to not vote, undermining the democratic nature of the system [57]. We have solved this problem through the creation of 'single-use' ERC20 voting tokens, SVT (Sharpe Voting Tokens), issued to every SHP owner at 1-to-1 parity whenever a motion has been seconded. These tokens cannot be transferred to any address other than two addresses created by Sharpe specifically for the purpose of vote counting. One of these addresses represents 'PASS' votes, and the other represents 'FAIL' votes. This permits users to freely transfer SHP during the 7 day period from a motion being tabled to the voting period ending, as only a fixed supply of SVT will have been created and these cannot be transferred. Once the voting period has ended, users will be able to trustlessly verify the result by viewing the amount of SVT in the 'PASS' and 'FAIL' addresses. The smart contract will destroy all SVT at the end of the voting period. Should more than one motion be tabled at once, the Smart Contract will create multiple single use tokens, each usable only for a specific motion.

We will create a web and mobile platform to simplify the voting process, such that users merely need to select the motion and indicate their vote. This application will allow previous motions and the results to be viewed on the Ethereum blockchain.

7.2 Consensus-based Community Fund Governance

We envisage the creation of a range of funds that serve different purposes,focused on different asset classes - including global equity markets and blockchain assets, with their issuance governed by community consensus.These tokenised funds will be created through independent crowd-sales, and 100

7.3 Corporate Structure & Regulatory Oversight

Sharpe Capital is in discussion with various financial and regulatory authorities with regard to tax-efficient structuring of the primary vehicle through which the proprietary investment fund will operate. Sharpe Capital Ltd will remain a UK entity, responsible for the development of the Sharpe Financial Markets Protocol, Investment Platform and associated products (including private ledger service and blockchain solutions for 3rd party hedge fund auditing systems, the crowd-sourced sentiment platform and the quantitative trading model).

The SHP token generation event will take place through a Gibraltar registered entity due to the clarity of DLT (Distributed Ledger Technology) financial regulations. The England-registered Sharpe Capital Ltd will become a wholly owned subsidiary of this primary vehicle, and operate as a software and quantitative model development company.

The community at large has been resistant to regulatory oversight, due to the *ideally* trustless, decentralised, and therefore self-governing, nature of cryptocurrencies. This made perfect sense in the early days of the development of a technology whose full scope and potential was far from realised. The recent SEC declaration signals a new era in the history of blockchain technology development, as it begins to penetrate the mainstream economic space with increasing pace. As such, regulatory oversight is an inevitability, and serves to protect both participants and the creators of the many innovative blockchain products both in the FinTech space and other areas.

8 BUDGETING & FINANCE CONSIDERATIONS

We are proposing the following utilisation of funds following completion of the SHP crowd-sale, illustrated in Figure 9. This section assumes a crowd-sale of $20m USD, in which $2m is deposited as reserve currency across BNT, ETH and USD(T), to both provide continuous liquidity via the Bancor Protocol (Section 2.4), and from which we will allocate quarterly pools for distribution as service fees to sentiment platform participants. The largest allocation of funds, 55%, is earmarked for proprietary investment using the state-of-the-art trading model described in Section 5. Prior to deployment of the complete modelling and investment platform in Q1 2019, we will distribute these funds across a variety of low-risk tracker funds such as Vanguard and S&P500, blockchain asset baskets, and other financial instruments, to attempt to prevent deprecation.

This fund distribution will be recorded on the TLS. 20% of capital raised during the token crowd-sale will be utilised for development and operations: essentially building the Sharpe Capital Investment Platform from its present stage to a deployable trading algorithm. This includes the expansion of the Sharpe team, including hiring full-stack software engineers, machine learning developers and financial analysts, and all associated overheads. Marketing is allocated 10% of raised capital, this is crucial for developing our additional revenue streams such as the sale of sentiment data and implementation of internal auditing software based on our TLS for hedge funds and other financial institutions. Finally, 5% of the funds are allocated for handling any legal and regulatory requirements to ensure we are in compliance with local laws in countries in which we operate, as appropriate. Any remaining funds following expenditure on development, marketing and legal costs will be diverted to the capital investment fund.

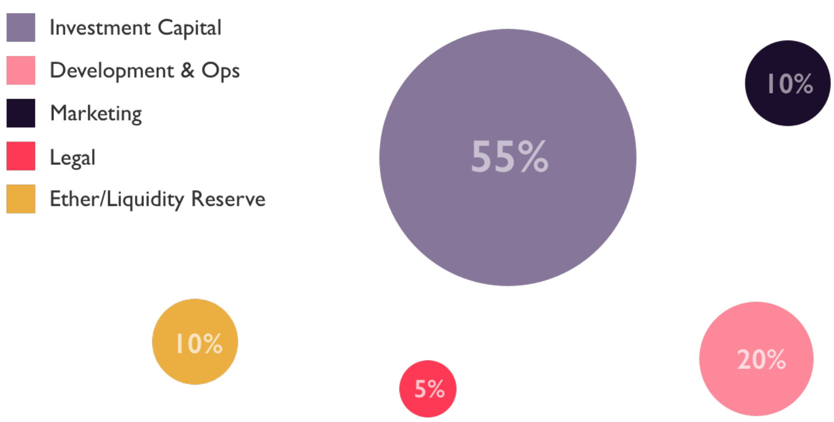

Figure 9: Our proposed budget allocations assuming a crowd-sale raise of $20m USD is realised.

This budget allocation was developed and agreed by the core team members to permit us to achieve the following key goals in delivering the protocol and platform:

- Grow our team to support our delivery objectives and establish our London HQ

- Support on-going development of the early livenet Beta web platform, slated to be available upon crowd-sale conclusion on 11th December 2017. Users can begin accruing service fees upon launch of this platform.
- Deploy an alpha (testnet) version of the Sharpe Platform mobile application by Q1 2018
- Migrate this to the mainnet toward the end of Q2.
- Market our sentiment analysis platform and build strategic relationships with hedge funds
- Build our core platform, external APIs, sentiment platform & smart contracts
- Develop, test & deploy our quantitative investment models, by Q2 2019
- Obtain any necessary approvals required to issue SCD tokens during an independent crowd-sale with a target date of Q1 2019

9 THE SHARPE CAPITAL TEAM

The Sharpe Capital team is comprised of a diverse group of experts across the fields of quantitative modelling, financial engineering, linguistic analysis, international law & regulatory requirements. Our team is expanding rapidly, with new members joining both the core team and advisory board as scale up our operations.

9.1 Core Team

Lewis M. Barber
Chief Executive Officer
Lewis graduated with a Bachelor of Engineering degree with Honours from the University of Leeds in 2011. He has since developed extensive experience in full-stack software engineering, and in consultancy with global investment banks on their IT architecture and the development of fault tolerant transactional systems. At Sharpe Capital, Lewis is responsible for driving the strategic development of the busines and overseeing implementation of our advanced cloud and blockchain architecture, including the Sharpe Capital Financial Markets Protocol.

Dr James A. Butler
Chief Investment Officer
James graduated with Lewis in 2011 and went on to complete a PhD in complex systems modelling at the University of York in 2015. He subsequently worked consulting with academic and commercial partners in the pharmaceutical sector. He was invited to a one year appointment as a Visiting Research Scholar at the University of California, Berkeley (EECS) in 2016. He is currently a Research Fellow at the University of Oxford, applying systems analysis to accelerate the process of drug development using advanced quantitative modelling and machine learning techniques. At Sharpe Capital, James is responsible for overseeing the development of the Sharpe Capital Investment Platform.

Israel Colomer
Chief Technology Officer
Israel graduated with a Bachelor of Engineering degree from the Universidad Politecnica de Valencia in 2007. He has since worked with a myriad of companies, both in the enterprise and startup spaces, with a dedicated focus on implementing resilient and robust back-end systems. Having experience in all areas of software delivery, he has handled teams both small and larger during his career, evolving his leadership skills. At Sharpe Capital, Israel is responsible for driving the technological development of the company's IT infrastructure and software solutions.

Ali Bros
Lead Developer
Ali graduated with a Bachelor of Engineering degree in Electronic and Computer Engineering from the University of Leeds in 2011. As an R&D Engineer, he has since worked for a diverse array of startups and long-established companies in the technology and media sectors. Ali has an excellent knowledge of myriad software platforms and stacks with a special interest in Android and iOS platforms.

Dan Pilch
Cloud Architect & Systems Engineer
Dan graduated with a Bachelor degree in Computer Network Engineering from the University of Portsmouth in 2013. Since then he has worked in progressing DevOps culture, system operations, Cloud architecture and engineering highly-available, fault-tolerant and massively scalable systems. He has experience in managing operations teams & contributing to progressing a startup to an SME (Small and medium-sized enterprises). At Sharpe Capital, Dan is responsible for architecting cloud infrastructure on the AWS platform and Site Reliability Engineering.

9.2 Advisory Board

Our Advisory Board have a critical role, along with community participants, in shaping the direction and growth of Sharpe Capital. Our advisors have a diverse range of backgrounds and are all highly experienced in what they do, whether they're from from a finance & trading, blockchain, modelling, linguistics, or commerce background.

Dimitri Chupryna
Entrepreneur and Investor
Dimitri graduated with BA in Economics from University of San Francisco. Dimitri is an expert in business development and portfolio management. His interests primarily focus on rapid-growth technologies. He is a Managing Partner and co-founder of TaaS (Token-as-a-Service). TaaS is the first tokenized closed-end fund that allows its investors to capitalize on the rise of Blockchain markets, and produced a 61% ROI for its first fully-operational quarter.

Dr Mieke Vandenbroucke
Advisor in Linguistic Analysis
Mieke completed her MA (Magna Cum Laude) in 2010 and PhD in 2015 at the University of Ghent. In 2016 she was awarded a Fulbright Scholarship and took up an appointment as a Visiting Researcher at UC Berkeley specialising in fraud investigations. Her research interests lie at the intersection of discourse analysis and forensic linguistics. As a Sharpe Capital Advisory Board Member, Mieke is responsible for advising on development of the linguistic analyses that power our trading models.

Dr German Leonov
Advisor in Quantitative Modelling
German completed his BSc (Hons) in 2011 and in 2015 earned his PhD in complex systems analysis from the University of York, where he is currently a post-doctoral research associate. German's research interests focus on the applications of quantitative modelling to understanding complex systems, currently focused on virology. As a Sharpe Capital Advisory Board Member, German is responsible for providing insight into the development of robust quantitative models, and for providing Russian language support to token crowd-sale participants. German also provides Russian language support to the community.

Jonas Sevel Karlberg
Commercial & TGE advisor
Jonas has more than 20 years' commercial experience working with some of the worlds largest consumer brands. He has significant management and leadership experience as well as experience in building and scaling organizations, primarily within the retail sector in Denmark. Jonas co-founded the Nordic Blockchain Association in the beginning of 2017 and later operated as a freelance community manager during the Bancor TGE. Jonas then founded the AmaZix team (https://www.amazix.com) and went on to work with the Stox TGE, and bitJob ICO, before joining the advisory board of Sharpe Capital.

Lexi Gao
Advisor in International Law
In 2010, Lexi completed a Bachelor's degree in Chinese Law from ECUST (East China University of Science and Technology) and a Bachelor's degree in Applied Psychology from East China Normal University. In 2011, Lexi completed a Master of Laws degree in International Corporate and Commercial Law at the University of York, followed by a Graduate Diploma in Law from the University of Law. Based in Shanghai, Lexi is responsible for providing insight into international legal and regulatory issues, and Chinese language support for crowd-sale participants from East Asian markets.

Barnaby Mannerings
Blockchain and Financial Markets Specialist
Barnaby completed a BSc (Hons) in Computer science in 2004 from Lancaster University's Cartmel College. He is a Capital Markets industry veteran with deep experience including design and implementation of electronic trading platforms and algorithms for equities, fixed income derivatives and FX products. From 2011, Barnaby was a consultant in Financial Services specialising in innovative technology and approaches. He has been involved with cryptocurrencies since 2012 in both a personal and professional capacity, including early cryptocurrency investments, solution design, training and outreach, and as a blockchain strategist. Barnaby is currently the CEO of Pik, a payments startup for online publishing.

10 NEXT STEPS

This white paper should be considered a 'living document'. That is, in response to feedback from the community and advisory board, we will update, grow, adapt, improve and extend both the Sharpe Capital Investment Platform and the Financial Markets Protocol we've described in this paper. Many of the smart contracts required for deployment of this platform have been created and discussed within this document and on our 'Medium' blog. As we move forward in developing these over the coming weeks, we will publish our progress on our blog, Telegram and Twitter to ensure maximum community engagement. As more substantial developments are made, we will update this white paper with further details, screenshots and mock-ups describing how the platform and protocol will operate. We are currently developing proof-of-concept platforms for the TLS and crowd-sourced sentiment analysis, while further refining our quantitative trading model approach. These will be included in future versions of our white paper.

In the interests of transparency, we will retain each major release of our white paper using version control, such that the community can see how the concepts, products and protocol described within this document have adapted and grown in response to the feedback we receive. Our Telegram group will be staffed around the clock leading up to the pre-sale and crowd-sale dates of November 6th and November 13th 2017 (both beginning at 14:00 UTC).

This document has been written at a level that tries to strike a careful balance between providing a high-level overview of the Sharpe Protocol & Platform, while providing sufficient technical detail and further references to the scientific literature for the interested reader. In addition to this, we are working on a series of blog posts and other documents, to be made available on our website, that describe the various aspects of the platform at both a more general and a more technical audience, to satisfy interested parties from various backgrounds. In particular, we will soon be releasing a 'summary paper' that outlines the Protocol & Platform in a much shorter document, as we appreciate that this document runs to considerable length, in an attempt to cover as much material as is feasible. Should technical audiences have additional questions relating to any aspect of this paper, we would invite them to contact the team members in our Slack[4], in which we have dedicated channels for machine learning discussion, enquiries regarding our token crowd-sale event, and channels for enquiries in Chinese and Russian.

4 https://sharpe.capital/slack

11 TRANSLATION NOTICE

In the event of any discrepancy, ambiguity, disagreement or any other confusion arising from comparison between the English language version of this white paper and any official or unofficial translations into any other language, the English language version is to be always treated as the correct version of the document. If any discrepancies are noted between this document in English and other languages, please bring it to the attention of the Sharpe Capital team at hello@sharpe.capital.

12 WHITE PAPER PURPOSE

The purpose of this white paper is to provide an overarching vision of the Sharpe Platform, with particular emphasis on the Sharpe Platform Token, SHP. Nothing in this white paper constitutes a contract, agreement or memorandum of understanding between any parties, of any kind. Any participation in the pre-sale or crowd-sale is subject to the terms of agreement that will be provided before participants are provided with the Ethereum contribution address from which participants' ETH will be sent to Sharpe Capital, and to which SHP will be received.

13 APPENDICES

13.1 Backtesting Table of Trades

Table 2: Table 1/2 of Positions Taken by AI Portfolio Manager - 2016 Q4 Backtest

Ticker	Size	ProfitLoss	Direction	ClosePrice	OpenPrice	OpenDate
CAG	1.19	1238.70	SHORT	36.55	46.92	Fri Jul 01
JWN	0.82	1237.55	LONG	52.12	37.05	Fri Jul 01
LVLT	0.83	-430.66	SHORT	56.74	51.53	Fri Jul 01
MHK	0.23	406.46	LONG	N/A	185.48	Wed Jul 06
ISRG	0.06	-313.01	SHORT	724.00	668.99	Thu Jul 07
NVDA	0.68	1266.67	LONG	66.46	47.75	Thu Jul 07
PPG	0.46	317.90	SHORT	N/A	101.68	Thu Jul 07
PSX	0.45	496.84	LONG	N/A	74.81	Thu Jul 07
UDR	1.41	-338.09	LONG	33.52	35.91	Thu Jul 07
COF	0.54	995.84	LONG	N/A	67.89	Thu Jul 14
CSCO	2.28	196.36	LONG	N/A	29.18	Fri Jul 15
EQT	0.44	457.61	SHORT	N/A	76.08	Fri Jul 15
JEC	0.81	-508.63	SHORT	58.76	52.50	Tue Jul 19
MO	1.14	-48.89	SHORT	N/A	67.00	Tue Jul 19
NFLX	0.37	-314.60	SHORT	94.00	85.43	Tue Jul 19
TAP	0.39	11.71	LONG	N/A	97.61	Tue Jul 19
WDC	0.49	-350.55	SHORT	58.22	51.01	Tue Jul 19
HRB	2.08	51.94	SHORT	N/A	23.28	Wed Jul 20
KMB	0.45	813.80	SHORT	N/A	132.26	Wed Jul 20
MNST	0.65	486.06	SHORT	N/A	52.54	Wed Jul 20
PH	0.42	1366.96	LONG	143.11	110.47	Wed Jul 20
TJX	0.62	-301.99	LONG	73.57	78.47	Wed Jul 20
BCR	0.22	162.07	SHORT	N/A	231.14	Thu Jul 21
DE	0.50	-325.26	SHORT	88.00	81.54	Thu Jul 21
DLPH	0.55	8.18	SHORT	N/A	67.47	Thu Jul 21
FLR	0.75	-338.68	SHORT	57.28	52.76	Thu Jul 21
OKE	0.74	-313.57	SHORT	49.23	45.00	Thu Jul 21
RHT	0.69	-600.57	SHORT	82.05	73.38	Thu Jul 21
TGT	0.64	-337.81	LONG	67.26	72.54	Thu Jul 21
AAL	0.87	970.11	LONG	N/A	36.09	Fri Jul 22
GRMN	0.83	357.89	LONG	N/A	44.35	Fri Jul 22
IRM	1.32	875.36	SHORT	N/A	38.55	Fri Jul 22
NLSN	0.99	1206.86	SHORT	42.01	54.24	Fri Jul 22
PBI	1.64	489.05	SHORT	N/A	18.12	Fri Jul 22
SO	1.39	435.55	SHORT	N/A	51.45	Fri Jul 22
UPS	0.63	-386.72	SHORT	112.73	106.57	Fri Jul 22
CME	0.58	1246.51	LONG	118.46	97.09	Mon Jul 25
DGX	0.86	-311.96	SHORT	87.10	83.47	Mon Jul 25
IR	0.74	-323.88	SHORT	70.85	66.45	Mon Jul 25
MLM	0.18	-305.21	LONG	180.26	197.03	Mon Jul 25
SE	1.39	-493.41	SHORT	38.74	35.19	Mon Jul 25
TRIP	0.32	737.10	SHORT	N/A	70.02	Mon Jul 25
AON	0.52	-324.48	SHORT	115.48	109.24	Tue Jul 26
AVGO	0.23	-308.82	SHORT	172.51	159.29	Tue Jul 26
CINF	0.82	-308.52	SHORT	77.24	73.46	Tue Jul 26
CSRA	1.00	-409.08	SHORT	29.78	25.69	Tue Jul 26
DHI	1.47	-332.33	LONG	31.10	33.36	Tue Jul 26
HD	0.45	-25.01	SHORT	N/A	134.26	Tue Jul 26
MRK	1.11	-411.36	SHORT	61.13	57.41	Tue Jul 26
ORLY	0.15	-39.94	SHORT	N/A	280.00	Tue Jul 26
PVH	0.37	-351.60	SHORT	109.92	100.43	Tue Jul 26
AAPL	0.62	832.68	LONG	N/A	102.31	Wed Jul 27
AMZN	0.04	-326.52	SHORT	818.00	737.97	Wed Jul 27
BA	0.48	-384.71	LONG	123.87	131.83	Wed Jul 27
CFG	1.89	1343.47	LONG	29.20	22.08	Wed Jul 27
FE	1.44	553.17	SHORT	N/A	34.27	Wed Jul 27
FRT	0.37	-334.23	LONG	154.24	163.31	Wed Jul 27
HPE	1.40	396.82	LONG	N/A	20.33	Wed Jul 27
LNC	0.79	1360.31	LONG	60.42	43.10	Wed Jul 27
MSI	0.60	916.65	LONG	N/A	67.31	Wed Jul 27
MU	1.70	-351.15	SHORT	16.92	14.85	Wed Jul 27
RF	4.44	1313.65	LONG	12.00	9.04	Wed Jul 27
URI	0.41	-343.76	SHORT	89.04	80.60	Wed Jul 27
VIAB	0.59	564.25	SHORT	N/A	44.81	Wed Jul 27
AMGN	0.39	-316.95	SHORT	173.88	165.69	Thu Jul 28
BHI	0.72	1279.67	LONG	62.71	44.82	Thu Jul 28
C	1.19	1272.32	LONG	54.30	43.61	Thu Jul 28
CAH	0.61	-478.00	LONG	73.80	81.68	Thu Jul 28
CI	0.32	68.70	SHORT	N/A	137.70	Thu Jul 28

Table 3: Table 2/2 of Positions Taken by AI Portfolio Manager - 2016 Q4 Backtest

Ticker	Size	ProfitLoss	Direction	ClosePrice	OpenPrice	OpenDate
DO	0.94	-309.61	LONG	18.62	21.92	Thu Jul 28
EBAY	1.58	142.55	SHORT	N/A	31.28	Thu Jul 28
EXR	0.60	-353.03	LONG	80.03	85.87	Thu Jul 28
FTR	7.36	1235.94	SHORT	2.95	4.63	Thu Jul 28
HCN	0.71	-421.78	LONG	67.50	73.47	Thu Jul 28
HSIC	0.32	818.43	SHORT	N/A	178.87	Thu Jul 28
KEY	3.97	1546.94	LONG	15.35	11.45	Thu Jul 28
KR	1.37	46.42	SHORT	N/A	34.71	Thu Jul 28
MDLZ	1.32	-190.66	SHORT	N/A	43.33	Thu Jul 28
MNK	0.35	-333.58	SHORT	77.15	67.63	Thu Jul 28
NDAQ	0.92	261.60	SHORT	N/A	69.85	Thu Jul 28
NSC	0.59	-336.26	SHORT	93.36	87.68	Thu Jul 28
PAYX	1.34	-312.60	LONG	55.94	58.28	Thu Jul 28
PX	0.60	198.69	LONG	N/A	113.15	Thu Jul 28
SIG	0.34	-407.29	SHORT	98.21	86.28	Thu Jul 28
SYY	1.20	-316.11	LONG	47.93	50.57	Thu Jul 28
TGNA	2.04	149.14	SHORT	N/A	22.23	Thu Jul 28
AGN	0.15	738.12	SHORT	N/A	254.94	Fri Jul 29
AMT	0.64	548.74	SHORT	N/A	114.49	Fri Jul 29
AN	0.81	286.35	SHORT	N/A	52.28	Fri Jul 29
BBY	0.94	1255.58	LONG	45.93	32.55	Fri Jul 29
HCA	0.72	-313.83	SHORT	81.11	76.73	Fri Jul 29
HST	2.36	-340.26	SHORT	18.06	16.62	Fri Jul 29
IP	1.31	1281.28	LONG	54.12	44.31	Fri Jul 29
NUE	0.91	-361.46	LONG	48.66	52.64	Fri Jul 29
PWR	1.34	-319.82	SHORT	27.95	25.57	Fri Jul 29
TXT	1.34	1271.57	LONG	48.57	39.08	Fri Jul 29
AWK	0.65	583.22	SHORT	N/A	80.83	Mon Aug 01
CF	0.99	-382.80	SHORT	27.29	23.41	Mon Aug 01
CLX	0.48	412.75	SHORT	N/A	128.19	Mon Aug 01
ADI	0.61	-325.29	SHORT	67.99	62.66	Tue Aug 02
AFL	1.05	-82.97	LONG	N/A	70.39	Tue Aug 02
CMG	0.09	335.39	SHORT	N/A	412.70	Tue Aug 02
CTAS	0.70	-604.18	SHORT	114.09	105.47	Tue Aug 02
CVX	0.55	-376.68	SHORT	103.48	96.59	Tue Aug 02
DAL	0.83	898.88	LONG	N/A	38.28	Tue Aug 02
GLW	2.41	-345.12	SHORT	23.31	21.88	Tue Aug 02
HOG	0.59	431.28	LONG	N/A	50.73	Tue Aug 02
KLAC	0.93	-548.30	LONG	68.56	74.45	Tue Aug 02
MAT	1.54	599.19	SHORT	N/A	31.22	Tue Aug 02
PBCT	3.64	1292.84	LONG	18.05	14.50	Tue Aug 02
F	4.05	145.66	LONG	N/A	11.53	Wed Aug 03
LEG	1.17	94.99	SHORT	N/A	50.25	Wed Aug 03
NEM	0.63	640.26	SHORT	N/A	45.51	Wed Aug 03
NKE	1.09	-335.87	LONG	50.96	54.04	Wed Aug 03
SBUX	1.18	-356.96	LONG	52.70	55.72	Wed Aug 03
BBBY	1.01	-321.68	LONG	40.04	43.23	Thu Aug 04
CBG	1.46	-426.75	SHORT	31.75	28.82	Thu Aug 04
CHD	1.06	415.59	SHORT	N/A	48.47	Thu Aug 04
DVA	0.78	1397.09	SHORT	58.27	76.09	Thu Aug 04
EA	0.63	138.36	LONG	N/A	78.36	Thu Aug 04
EXPE	0.41	-323.03	SHORT	118.78	110.88	Thu Aug 04
FCX	1.82	-327.08	LONG	10.43	12.23	Thu Aug 04
FIS	0.86	-321.90	LONG	73.72	77.46	Thu Aug 04
GPC	0.73	-321.20	SHORT	101.44	97.04	Thu Aug 04
HPQ	2.62	237.98	LONG	N/A	13.98	Thu Aug 04
HRL	1.13	-364.99	SHORT	39.05	35.82	Thu Aug 04
RSG	1.51	-338.51	SHORT	52.24	50.00	Thu Aug 04
SNA	0.46	-358.59	SHORT	160.44	152.71	Thu Aug 04
SWK	0.63	-422.28	LONG	113.52	120.24	Thu Aug 04
SYMC	2.27	775.67	LONG	N/A	20.76	Thu Aug 04
ADM	1.32	-421.06	SHORT	46.09	42.91	Fri Aug 05
ADSK	0.65	1281.29	LONG	78.21	58.44	Fri Aug 05
EQIX	0.15	148.14	SHORT	N/A	366.26	Fri Aug 05
GT	1.69	-377.84	SHORT	30.31	28.07	Fri Aug 05
IPG	2.67	-341.45	SHORT	23.70	22.42	Fri Aug 05
WBA	0.85	318.86	LONG	N/A	79.34	Mon Aug 08
BXP	0.52	-378.59	LONG	132.53	139.85	Tue Aug 09
EMR	1.22	-401.87	LONG	49.35	52.64	Tue Aug 09
NWSA	3.53	-547.15	LONG	11.06	12.61	Tue Aug 09
PCAR	0.88	-378.51	LONG	52.74	57.03	Tue Aug 09
FAST	1.53	-350.80	SHORT	44.18	41.88	Thu Aug 11
CERN	0.84	1368.84	SHORT	50.22	66.58	Tue Aug 16
ALLE	0.87	-426.89	LONG	66.35	71.28	Thu Aug 25
RCL	0.62	-343.83	SHORT	74.42	68.86	Wed Aug 31
MTB	0.57	1346.82	LONG	140.35	116.63	Fri Sep 02
JBHT	0.67	-420.88	SHORT	87.52	81.21	Tue Sep 20

13.2 TLS Smart Contract

In this section, the smart contract behind our trading ledger is broken down into key methods, giving a full overview of the functionality it currently exposes. The code itself will be frequently updated as development is ongoing, with source code publicly available on our GitHub page[5].

Add Account

Any individual or institution can add an account to the public trade ledger, with a specified initial balance, provided the account ID does not already exist.

Get Account

Anybody can call the *getAccount* function on the TradeLedger smart contract, provided the ID provided is valid. An array of key values from the account struct will be returned to the caller.

Update Account Leverage

The account owner can update the leverage stored against an account at any time by providing a valid account ID. The leverage is stored scaled to be a factor of 100 greater than the leverage value, as Ethereum's Solidity language for smart contracts does not support floating points. This facilitates accuracy of leverage recorded to two decimal places.

Add Position

The owner of an account can add new positions to it, provided the given account ID is valid and the position ID is unique. The *addPosition* method ensures the required parameters are provided, before adding the position and updating the leverage value for the associated account.

Update Position

The position owner can update a position's profit and loss, provided the position ID is valid and the position is currently open. The profit and loss is also updated against the account struct and a new equity point is recorded.

Close Position

The position owner can close a position, provided the position ID is valid and the position is open. When closing a position, the final profit and loss must be provided, plus the closing price and date (including time component and time zone). After closing a position the account leverage is updated and a new equity point is recorded.

Get Position

Anybody can get a position by ID, provided the ID is valid. Key values from the position will be returned in an array.

5 https://github.com/sharpe-capital

Get Position Key

Anybody can get the RSA keys for a position, provided the ID is valid. If the position's RSA keys have not been released yet by the position owner, 'TBC' will be returned.

Fetch Positions

Unfortunately, Solidity does not support returning arrays from functions. Therefore, in order to fetch all positions for an account, the consumer of the smart contract must first count the positions using function *countAccount-Positions* and then call the *getPositionByIndex* method for each valid index value. Both the *countAccountPositions* and *getPositionByIndex* methods can be called by anyone, provided the account ID is valid.

Release Key Pair

The position owner can call the *releaseKeyPair* method at any time to make the RSA encryption keys available on the blockchain, enabling consumers of the trading ledger smart contract to view sensitive position details.

Add Equity Point

New equity points can only be added by the account owner, and the method is only callable from within the smart contract – it is not a public method. This method is called from other public methods (such as function *updatePo-sition*) and is intended to store a snapshot of the account struct at a given point in time.

Fetch Equity Points

As with fetching positions, the consumer must first count the number of eq-uity points for an account and then call the getEquityPointByIndex method for each valid index value. Both the countAccountEquityPoints and getE-quityPointByIndex methods can be called by anyone, provided the account ID is valid.

REFERENCES

[1] Alex Frino, David Johnstone, and Hui Zheng. The propensity for local traders in futures markets to ride losses: Evidence of irrational or rational behavior? *Journal of Banking & Finance*, 28(2):353–372, 2004.

[2] Benedetto De Martino, Dharshan Kumaran, Ben Seymour, and Raymond J Dolan. Frames, biases, and rational decision-making in the human brain. *Science*, 313(5787):684–687, 2006.

[3] Arjen Siegmann and Andre Lucas. Explaining hedge fund investment styles by loss aversion: a rational alternative. 2002.

[4] Denise Shull. *Market mind games: A radical psychology of investing, trading and risk*. McGraw Hill Professional, 2011.

[5] Seokjoo Andrew Chang, Guy Fernando, and Mohamed E Hussein. Internal control computerization for derivatives. 2012.

[6] Lars Hornuf and Georg Haas. Regulating fraud in financial markets: can behavioural designs prevent future criminal offences? *Journal of Risk Management in Financial Institutions*, 7(2):192–201, 2014.

[7] Georg Haas and Lars Hornuf. Regulating fraud in financial markets: Can behavioural designs prevent future criminal offences? *European Business Law Review*, 28(1):41–54, 2017.

[8] Adam Tickell. Making a melodrama out of a crisis: reinterpreting the collapse of barings bank. *Environment and Planning D: Society and Space*, 14(1):5–33, 1996.

[9] Eleanor Stead and Clive Smallman. Understanding business failure: learning and un-learning from industrial crises. *Journal of contingencies and crisis management*, 7(1):1–18, 1999.

[10] Mark N Wexler. Financial edgework and the persistence of rogue traders. *Business and Society Review*, 115(1):1–25, 2010.

[11] Ian Greener. Nick leeson and the collapse of barings bank: socio-technical networks and the 'rogue trader'. *Organization*, 13(3):421–441, 2006.

[12] Helga Drummond. Living in a fool's paradise: The collapse of barings' bank. *Management Decision*, 40(3):232–238, 2002.

[13] Robert Adamson. Corporate governance reforms and our regulatory future. *Business Horizons*, 55(6):551–555, 2012.

[14] Kuzman Ganchev, Yuriy Nevmyvaka, Michael Kearns, and Jennifer Wortman Vaughan. Censored exploration and the dark pool problem. *Communications of the ACM*, 53(5):99–107, 2010.

[15] Carole Comerton-Forde and Tālis J Putniņš. Dark trading and price discovery. *Journal of Financial Economics*, 118(1):70–92, 2015.

[16] Eugene F Fama. Market efficiency, long-term returns, and behavioral finance. *Journal of financial economics*, 49(3):283–306, 1998.

[17] Burton G Malkiel. The efficient market hypothesis and its critics. *The Journal of Economic Perspectives*, 17(1):59–82, 2003.

[18] Sebastien Pouget, Julien Sauvagnat, and Stephane Villeneuve. A mind is a terrible thing to change: confirmatory bias in financial markets. *The Review of Financial Studies*, 30(6):2066–2109, 2017.

[19] Antoine Bechara. The role of emotion in decision-making: evidence from neurological patients with orbitofrontal damage. *Brain and cognition*, 55(1):30–40, 2004.

[20] Brian M Lucey and Michael Dowling. The role of feelings in investor decision-making. *Journal of economic surveys*, 19(2):211–237, 2005.

[21] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *OSDI*, volume 16, pages 265–283, 2016.

[22] Nangang District. A combination of models for financial crisis prediction: Integrating probabilistic neural network with back-propagation based on adaptive boosting. *International Journal of Computational Intelligence Systems*, 10:507–520, 2017.

[23] Pastukhov Aleksey and Prokofiev Alexander. Representative sample formation with the use of kohonen self-organizing map and lipschitz constant estimation in the training of a multilayer perceptron. In *Young Researchers in Electrical and Electronic Engineering (EIConRus), 2017 IEEE Conference of Russian*, pages 712–714. IEEE, 2017.

[24] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008.

[25] Ingrida Vaiciulyte, Zivile Kalsyte, Leonidas Sakalauskas, and Darius Plikynas. Assessment of market reaction on the share performance on the basis of its visualization in 2d space. *Journal of Business Economics and Management*, 18(2):309–318, 2017.

[26] Eckart Zitzler and Lothar Thiele. Multiobjective evolutionary algorithms: a comparative case study and the strength pareto approach. *IEEE transactions on Evolutionary Computation*, 3(4):257–271, 1999.

[27] Luís Lobato Macedo, Pedro Godinho, and Maria João Alves. Mean-semivariance portfolio optimization with multiobjective evolutionary algorithms and technical analysis rules. *Expert Systems with Applications*, 79:33–43, 2017.

[28] Yuriy Nevmyvaka, Yi Feng, and Michael Kearns. Reinforcement learning for optimized trade execution. In *Proceedings of the 23rd international conference on Machine learning*, pages 673–680. ACM, 2006.

[29] Michael AH Dempster and Chris M Jones. A real-time adaptive trading system using genetic programming. *Quantitative Finance*, 1(4):397–413, 2001.

[30] Michael AH Dempster and Vasco Leemans. An automated fx trading system using adaptive reinforcement learning. *Expert Systems with Applications*, 30(3):543–552, 2006.

[31] Jun Wang. Trading and hedging in s&p 500 spot and futures markets using genetic programming. *Journal of Futures Markets*, 20(10):911–942, 2000.

[32] Kyung-Shik Shin, Taik Soo Lee, and Hyun-jung Kim. An application of support vector machines in bankruptcy prediction model. *Expert Systems with Applications*, 28(1):127–135, 2005.

[33] Wei Huang, Yoshiteru Nakamori, and Shou-Yang Wang. Forecasting stock market movement direction with support vector machine. *Computers & Operations Research*, 32(10):2513–2522, 2005.

[34] Qing Cao, Karyl B Leggio, and Marc J Schniederjans. A comparison between fama and french's model and artificial neural networks in predicting the chinese stock market. *Computers & Operations Research*, 32(10):2499–2512, 2005.

[35] Jigar Patel, Sahil Shah, Priyank Thakkar, and K Kotecha. Predicting stock and stock price index movement using trend deterministic data preparation and machine learning techniques. *Expert Systems with Applications*, 42(1):259–268, 2015.

[36] Mark P Taylor and Helen Allen. The use of technical analysis in the foreign exchange market. *Journal of international Money and Finance*, 11(3):304–314, 1992.

[37] JG Agrawal, VS Chourasia, and AK Mittra. State-of-the-art in stock prediction techniques. *International Journal of Advanced Research in Electrical, Electronics and Instrumentation Engineering*, 2(4):1360–1366, 2013.

[38] Fred Block and Margaret R Somers. *The power of market fundamentalism*. Harvard University Press, 2014.

[39] Avrim L Blum and Pat Langley. Selection of relevant features and examples in machine learning. *Artificial intelligence*, 97(1):245–271, 1997.

[40] Randall Morck, Bernard Yeung, and Wayne Yu. The information content of stock markets: why do emerging markets have synchronous stock price movements? *Journal of financial economics*, 58(1):215–260, 2000.

[41] Graciela Kaminsky and Sergio L Schmukler. Emerging market instability: do sovereign ratings affect country risk and stock returns? *The World Bank Economic Review*, 16(2):171–195, 2002.

[42] Kam C Chan, Benton E Gup, and Ming-Shiun Pan. International stock market efficiency and integration: A study of eighteen nations. *Journal of business finance & accounting*, 24(6):803–813, 1997.

[43] Kenneth Kasa. Common stochastic trends in international stock markets. *Journal of monetary Economics*, 29(1):95–124, 1992.

[44] Ahmed Abbasi, Ammar Hassan, and Milan Dhar. Benchmarking twitter sentiment analysis tools. In *LREC*, volume 14, pages 26–31, 2014.

[45] Johan Bollen, Huina Mao, and Xiaojun Zeng. Twitter mood predicts the stock market. *Journal of computational science*, 2(1):1–8, 2011.

[46] Bo Pang, Lillian Lee, et al. Opinion mining and sentiment analysis. *Foundations and Trends® in Information Retrieval*, 2(1–2):1–135, 2008.

[47] Gaowei Zhang, Lingyu Xu, and Yunlan Xue. Model and forecast stock market behavior integrating investor sentiment analysis and transaction data. *Cluster Computing*, 20(1):789–803, 2017.

[48] Ramón A Feenstra and Daniel Pallarés-Domínguez. Ethical debates on political neuromarketing: the technological advance and its potential impact on the formation of public opinion. 2017.

[49] Nick Lee, Leif Brandes, Laura Chamberlain, and Carl Senior. This is your brain on neuromarketing: reflections on a decade of research. *Journal of Marketing Management*, pages 1–15, 2017.

[50] Zhenkun Zhou, Ke Xu, and Jichang Zhao. Tales of emotion and stock in china: Volatility, causality and prediction. *arXiv preprint arXiv:1705.00294*, 2017.

[51] Bjarke Felbo, Alan Mislove, Anders Søgaard, Iyad Rahwan, and Sune Lehmann. Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm. *arXiv preprint arXiv:1708.00524*, 2017.

[52] Branca Telles Ribeiro and Susan M Hoyle. Frame analysis. *Grammar, meaning and pragmatics*, 5:74, 2009.

[53] Charles J Fillmore. Frame semantics and the nature of language. *Annals of the New York Academy of Sciences*, 280(1):20–32, 1976.

[54] Paul H Thibodeau and Lera Boroditsky. Metaphors we think with: The role of metaphor in reasoning. *PloS one*, 6(2):e16782, 2011.

[55] Jane Elith, John R Leathwick, and Trevor Hastie. A working guide to boosted regression trees. *Journal of Animal Ecology*, 77(4):802–813, 2008.

[56] Sebastian Mallaby. *More Money Than God: Hedge Funds and the Making of a New Elite*. The Penguin Press, 2010.

[57] Quinn DuPont. Experiments in algorithmic governance: A history and ethnography of "the dao," a failed decentralized autonomous organization. *Bitcoin and Beyond: Cryptocurrencies, Blockchains and Global Governance. Routledge*, 2017.

[58] SEC. Report of investigation pursuant to section 21(a) of the securities exchange act of 1934: The dao. *Release No. 81207*, 2017.