

A Jump Start to Stock Trading Research for the Uninitiated Control Scientist: A Tutorial

B. Ross Barmish, Simone Formentin, Chung-Han Hsieh, Anton V. Proskurnikov and Sean Warnick

Abstract—The target audience for the line of research to be described in this tutorial paper is control system researchers with an interest in algorithmic stock trading but without a substantial background in finance and economics. To this end, we focus our attention on just a few hand-picked problem areas to illustrate how algorithmic trading research might be carried out from a control-theoretic perspective and refer the reader to a number of references where extensive survey-style material can be found. The paper begins with the exposition of some basics associated with opening a brokerage account and mathematical modelling of common order types. Subsequently, we consider a number of trading scenarios involving feedback control design, optimization problems arising in portfolio management, the theory of Kelly Betting in a stock trading context and interaction with the Limit Order Book which is crucial for smooth market operation. Given the control-theoretic point of view taken in this paper, many of our basic tools come into play; e.g., standard results from areas such as convex optimization, discrete probability theory and Markov processes, to name a few. One of the salient features of this tutorial is our use of idealizing assumptions and simplistic models whenever convenient for pedagogical and motivational purposes. In the conclusion section, we mention some challenging new research opportunities involving more general models and relaxation some of our simplifying assumptions.

I. INTRODUCTION

As a starting point for this tutorial, it is convenient to imagine a small stock trader who is operating from a desktop computer at home with a high-speed internet connection and a limited budget. Orders are being submitted in discrete time and the trader is positioned within a feedback loop, as depicted in Figure 1. The restriction to this narrow framework, in lieu of a general problem formulation involving hedge funds, investment professionals, extensive computing resources and the like, is deliberate. Whereas in this paper, for simplicity, we confine attention to the use of historical stock price data, in a more general framework, feeding additional data such as news, earnings and social media to the controller would also be of interest.

Throughout this paper, to communicate basic ideas about stock trading, whenever convenient for pedagogical and

B. Ross Barmish is an emeritus professor in ECE at University of Wisconsin, Madison, and CEO of Robust Trading Solutions, LLC; bob.barmish@gmail.com. Simone Formentin is a faculty member at Politecnico di Milano in the Department of Electronics, Information and Bioengineering; simone.formentin@polimi.it. Chung-Han Hsieh is a faculty member at National Tsing Hua University in the Department of Quantitative Finance; ch.hsieh@mx.nthu.edu.tw. Anton Proskurnikov is a faculty member at Politecnico di Torino in the Department of Electronics and Telecommunications; anton.p.1982@ieee.org. Sean Warnick is a faculty member at Brigham Young University in Computer Science; sean.warnick@gmail.com.




Fig. 1. Control-Theoretic Setup for Stock Trading

motivational purposes, we make idealizing assumptions and work with some of the simplest models available in the literature. Then, in the conclusion, we mention some more general modelling possibilities and revisit some of our assumptions which are not straightforward to relax. This serves as a stepping stone to research on some significant new problems which can be easily understood within the control-theoretic framework of this paper.

In the exposition to follow, we provide citations which can be used to supplement the material which we cover. Consistent with the financial engineering perspective taken in this paper, some good starting textbooks for the reader to consider are [1]–[7]. Moreover, given that the topics to follow are described in a control-theoretic setting, the reader may wish to consult [8] which provides a survey of many papers taking this same point of view. Unlike a survey paper which is aimed at giving the authors’ perspective as to which papers are important from a research contribution point of view, our choice of tutorial citations are mainly limited to the topics which we are using to illustrate how to easily conduct stock trading research for attendees from the control community.

With these considerations in mind, the plan for the remainder of this paper is as follows: In Section II, we begin by considering a number of preliminaries associated with becoming a real-world trader. This includes but is not limited to considerations related to selecting a reputable broker, account insurance and regulations by governmental agencies, deciding on account type (margin versus cash) and utilizing an Application Program Interface (API). Subsequently, in Section III, we describe what might appropriately be called the “rules of the road.” To this end, we first carry out some mathematical modelling of the mechanics associated with some of the basic types of orders such as *market*, *limit* and *stop*. This opens the door to using many standard tools from systems and control; e.g., results on stochastic systems and

Markov processes, convexity, mathematical programming and adaptive control, to name a few.

In Section IV, stock trading is considered as a feedback control design problem. We highlight the use of linear feedback and the so-called *Simultaneous Long-Short* (SLS) controller as we describe advantages and limitations of this approach hopefully shedding light on its suitability across various market conditions. This is followed by Section V where the spotlight is on optimization problems in trading an individual stock and more generally a multi-stock portfolio. Then, in Section VI, we provide an introduction to Kelly Betting, and its eventual use in the stock market pioneered by Thorpe and others. Following this, in Section VII, we cover the so-called *Limit Order Book* (LOB) which can be viewed as the “brain center” for a fully electronic marketplace where buyers and sellers are matched up to consummate trades. Given the nanosecond time stamping of transactions in the LOB, this topic is seen to be a convenient starting point for the study of high-frequency trading. Then in Section VIII, a number of simulations are given, serving as a simple example of “backtesting.” This involves the assessment of the performance of various linear feedback stock trading controllers applied to the Grayscale Bitcoin Trust (GBTC), an exchange-traded fund, over a two-year period of high price volatility. Finally, in Section IX, the conclusion section, the reader is reminded about our objectives and some important problems for future research, motivated by our use of simplified models and idealizing assumptions, are discussed.

II. ON OPENING A BROKERAGE ACCOUNT

This preliminary section is entirely nontechnical and provides a brief description of most of the main issues to be addressed by readers who are determined to take some of the technical ideas in this paper and put them into practice. It can be skipped by readers solely interested in algorithmic trading issues that arise once an appropriate brokerage account is fully established.

Large brokerage companies offer a comprehensive array of financial services to cater to the diverse needs of investors. As a preliminary step in screening a broker under consideration, aspiring traders are well served by checking the range of financial instruments and services being offered. This includes, but is not limited to, stocks, options, bonds, exchange-traded funds, educational resources, various trading platforms, market data and news. Once a “candidate” brokerage company has been identified, it is arguable that an uninitiated trader’s highest priority should be to take every reasonable step to ensure that the firm being considered is trustworthy and adequately regulated. There are typically many internet reviews of the broker and the website should be highly transparent in its disclosure of costs to the trader such as commissions, broker-assisted trades, mutual funds, margin interest rates and account maintenance fees. Suffice it to say, when seeking a broker, protection and accessibility of the trader’s assets should also be assigned high priority.

Further to this “protection and accessibility,” in most countries, oversight of brokers is carried out by various

government and regulatory agencies. For example, in the U.S., agencies such as the Securities and Exchange Commission (SEC), Federal Reserve Board (FRB), Federal Deposit Insurance Corporation (FDIC) and the Financial Industry Regulatory Authority (FINRA) play an important protective role for the trader. These agencies address many critical issues, which include account insurance requirements, limitations on margin loans and many other restrictions aimed at guaranteeing account safety and integrity of all parties involved in trading. To provide further examples of such agencies, in the UK, the Financial Conduct Authority (FCA) regulates financial services, in Germany, the Federal Financial Supervisory Authority, known as BaFin, oversees financial institutions, in France, a similar role is played by *Autorité des Marchés Financiers* (AMF) and in Japan, it is the *Japan Financial Services Agency* (JFSA). Over and above governmental protections, it is a good idea to determine the extent to which a broker carries significant supplemental private insurance and whether any complaints have been filed. For example, in the United States, such complaints can be seen on the website of the Better Business Bureau and many other countries have similar organizations providing this type of information.

One final important consideration we mention pertains to traders planning to use an algorithm which potentially involves a trading frequency whose realization is “unimaginable” by clicking a mouse and entering orders by hand on one’s desktop: It would be important to consider whether the brokerage company offers an Application Programming Interface (API) and if so, the details of its use. Assuming an internet connection with adequate speed, this enables a small trader’s software to communicate with the brokerage’s servers thereby enabling order placement at a rate up to several trades per second.

Most large brokerage companies offer a wide variety of account types to meet various customer needs. For the purpose of this tutorial, we focus our main attention on the most basic of account types — so-called *cash account*. On few occasions in the sequel, for example, when short selling is involved, a *margin account* is required. Having such an account, a trader is given certain “lending privileges” by the broker; this makes it possible to carry out a variety of “leveraged” trades which are not available in a standard brokerage cash account; see Section IX where margin accounts are further discussed in the context of future research.

Finally, there is one important account setup detail to be mentioned: the specification of beneficiaries which would be in play in the event of the account owner’s death. Then, to activate the account, a trader’s initial funds come into the account usually by a wire transfer, Automated Clearing House (ACH) transfer or by a check. For a cash account, whereas some brokers require several thousand dollars to get started, in most cases the required initial funds can be very small — even as little as a few dollars. However, to open a margin account, it is typical to see regulatory authorities requiring significantly more; e.g., in the United States, the minimum initial deposit is two thousand dollars.

III. ON MATHEMATICAL MODELLING OF STOCK PRICES, ORDER TYPES AND DYNAMICS

In the sequel, we carry out some mathematical modelling of some of the most common types of orders used by traders. To this end, we first introduce some notation: We use index $k = 0, 1, 2, \dots, n - 1$ to enumerate n trading intervals each having time duration of $\Delta t > 0$, which is arbitrary and assumed to be known. This can accommodate various different investment styles. For instance, many money managers purchase or sell shares on a quarterly basis while high-frequency traders may be working with time stamps at the nanosecond level. We begin at $t_0 = 0$ and the k -th such interval is

$$T_k \doteq [t_k, t_k + \Delta t].$$

In practice, there are often “time gaps” between the end of one interval and the beginning of the next; e.g., to model stoppage of trading overnight and on holidays, for some values of k we may have $t_{k+1} > t_k + \Delta t$.

As far as the stock price is concerned, associated with each T_k interval, it is assumed that the following information is available: The *opening price* O_k , the *low* L_k , the *high* H_k and the *closing price* C_k satisfying the obvious inequalities

$$L_k \leq O_k \leq H_k; \quad L_k \leq C_k \leq H_k.$$

In the subsections to follow, we use the notation N_k to denote the number of shares under control by the trader at instant t_k . Note that $N_k < 0$ represents a short position, in which case, per earlier discussion, a margin account is required. When a short sale occurs, consistent with the standard practice of brokers, we require the proceeds to be held aside and not be available for use by the trader. We also introduce the notation M_k to represent the amount of money, say cash equivalents, held at the end of interval T_k . Finally, to complete our description of the notation, we let V_k denote the market value of the account at the end of interval T_k , and, for consistency with most authors who work with only one stock price rather than our quadruple (O_k, L_k, H_k, C_k) , we take $S_k \equiv C_k$. Then, trading is initialized with $V_0 > 0$ prespecified, $N_0 = 0$ and $M_0 = V_0$ and the account value at the end of interval T_k is found as

$$V_k = N_k S_k + M_k.$$

A. Simplifying Assumptions

For pedagogical purposes, we now impose a number of assumptions to avoid making the mathematical analysis unduly complicated for the uninitiated trader seeking to understand the main ideas driving this line of research. Per discussion in the Introduction, in the conclusion section, we revisit some of the assumptions below noting that relaxing some of them is highly nontrivial and suggest new research opportunities.

1) *Trading at the Daily Market Close Assumption:* In the sequel, for simplicity, we take the interval T_k corresponding to the stock market’s day session which typically lasts 6–8 hours. We further assume that orders are submitted at the

market closing time but may not necessarily be filled until the next session associated with T_{k+1} ; e.g., see the limit order scenario described later in this section.¹

2) *No-Leverage Assumption:* The value of the control variable u_k corresponds to the total desired market value of the shares, both new and old, associated with an order placed at the end of T_k . For simplicity, it is assumed that the *no-leverage condition* $|u_k| \leq V_k$ is satisfied. This simplifying assumption helps to avoid dealing with a number of technical issues associated with a margin account. For example, by restricting attention to a cash account, we avoid the need to address maintenance requirements, margin interest and fees for borrowed shares; see Section IX for further discussion.

3) *Cash Settlement at Closing Assumption:* To avoid technicalities associated with the trader incurring a so-called *good faith violation*, we assume that all trades are “cash settled” by the end of each trading day. Without such an assumption in place, the cash proceeds associated with a sale on Day k may not be available for use on Day $k + 1$.²

4) *No-Dividend or Interest Assumption:* For simplicity, it is assumed that the stock being traded does not pay dividends. It should be noted that one standard method for circumventing this assumption involves working with *adjusted closing prices* and taking account of the fact that a short seller is responsible for the payment of such dividends. As far as cash equivalents represented by M_k are concerned, in this section, for simplicity, it is assumed that no interest is accrued; see the examples in Sections IV and VIII where the analysis is expanded to include interest as well.

5) *Fractional Shares Assumption:* In the analysis to follow, whenever convenient, we do not require the number of shares being transacted to be an integer. In fact, some brokerage companies allow this practice while others do not. In practice, if the number of shares being traded is significant, from a practitioner’s point of view, “rounding off” a prescribed trade to an integer number of shares should not have a material effect on performance.

6) *Frictionless Market Assumption:* It is assumed that trading is being conducted in a *frictionless market*; e.g., see [9]. We note that this concept, instrumental to many models in the literature, takes on a number of different forms; see Section IX. We now elaborate on this rather technical requirement, and, to keep the exposition as simple as possible, we contextualize the discussion to the modelling to follow in the remainder of this section: Indeed, for each of the three order types to follow, assuming submission at the close of interval T_k , there is an associated number of desired shares ΔN_k to be bought or sold and a trigger price S'_k which is not necessarily the closing price, $S_k = C_k$. In this setting, the frictionless market assumption tells us the following. If the price S'_k is reached in the market, be it during T_k or T_{k+1} , the desired transaction

¹The analysis to follow is readily extended to allow for trading at the open as well.

²We note that the so-called cash settlement rule in the United States was updated on May 28, 2024 from $T + 2$ to $T + 1$; see Section IX where relaxation of this assumption is discussed.

will occur instantaneously. This type of instantaneous order-fill requirement is tantamount to a zero-latency assumption. That is the time for the order information to travel between the trader and exchange server is zero. Furthermore, the following *adequate liquidity condition* will be satisfied: If the trader is a buyer, sellers on the “ask side” of the market will be offering at least ΔN_k at price S'_k so as to guarantee that the order is filled. On the other hand, if the trader is a seller, buyers on the “bid side” of the market will be seeking at least ΔN_k so as to guarantee again that the order is filled. Whereas an assumption of adequate liquidity is typically reasonable for small desktop traders, it can easily fail to be satisfied for a large hedge fund trading millions of shares per day. For example, if the number of shares being sought is suitably large, it may be impossible to buy them without bidding up the price above S'_k . Finally, it is important to make the following comment which bears upon the *reliability* of the trigger price S'_k . If during T_{k+1} , we see $O_{k+1} \leq S'_k \leq H_{k+1}$, the continuous price variation assumption guarantees, by an intermediate value argument, that the order will be triggered.

B. Basic Order Types

In this section, we provide some illustrations of the modelling of some of the most common order types. Recalling that u_k denotes the desired total market value of the shares, both new and old, associated with an order placed at the end of T_k . It is important to note that this value which may not necessarily be realized. For example, as seen below, a limit order with an acceptable price $S'_k < C_k$ may never be filled on T_{k+1} .

1) *Modelling of Market Orders*: Suppose the trader is holding $N_k > 0$ shares (long) at the close of T_k and is seeking to increase the stock position from its current value $N_k S_k$ to desired value $u_k > N_k S_k$. Then, at the current market price $S_k = C_k$, the desired number of shares to acquire is

$$\Delta N_k \doteq \frac{u_k - N_k S_k}{S_k}$$

and the frictionless market assumption assures that a market order to buy these shares will be instantaneously filled. Hence, the number of shares being held over the next interval T_{k+1} is given by

$$N_{k+1} = N_k + \Delta N_k,$$

the new money balance is

$$\begin{aligned} M_{k+1} &= M_k - \Delta N_k S_k \\ &= M_k + N_k S_k - u_k \\ &= V_k - u_k \end{aligned}$$

and the updated account value is

$$V_{k+1} = N_{k+1} S_{k+1} + M_{k+1}.$$

Notice that, in accordance with the no-leverage assumption, the account balance remains nonnegative; i.e., $M_{k+1} \geq 0$.

Given the lack of contingencies for this most basic of all order types, and consistent with many existing papers, the update in the account value can also be expressed in terms of the single period return $X_k \doteq (S_{k+1} - S_k)/S_k$. That is,

$$V_{k+1} = V_k + u_k X_k.$$

Finally, it is also noted that a similar formula can readily be obtained for a market sell order or market short sale order.

2) *Modelling of Limit Orders*: At the end of T_k , at price S_k , suppose the trader is holding N_k shares (long) and wants to increase this position from its current value $N_k S_k$ to desired value $u_k > N_k S_k$ but only if the new shares are acquired at the discounted limit price $S'_k < S_k$. In this case, a *buy limit order* can be submitted with the desired number of shares being

$$\Delta N_k \doteq \frac{u_k - N_k S_k}{S'_k}.$$

Now, there are three mutually exclusive cases to consider. The first case, a windfall of sorts, occurs if $O_{k+1} \leq S'_k$. Given the frictionless nature of the market, the order is filled at the open of T_{k+1} at price O_{k+1} , and, with the same update equation for N_{k+1} as given above, the account money update is

$$\begin{aligned} M_{k+1} &= M_k - \Delta N_k O_{k+1} \\ &= M_k - \frac{u_k - N_k S_k}{S'_k} O_{k+1} \end{aligned}$$

and the account value update is the same as given above.

The second case occurs if $O_{k+1} > S'_k$ and $L_{k+1} \leq S'_k$. In this situation, the frictionless market assumption assures that limit price S'_k is reached during T_{k+1} . Then, again the same update N_{k+1} is used and the account money update is

$$\begin{aligned} M_{k+1} &= M_k - \Delta N_k S'_k \\ &= M_k - \frac{u_k - N_k S_k}{S'_k} S'_k \\ &= M_k - u_k + N_k S_k \end{aligned}$$

The third case occurs if $S'_k < L_{k+1}$ which implies that the limit order remains unfilled during T_{k+1} leading to

$$N_{k+1} = N_k; \quad M_{k+1} = M_k.$$

The outcome for this third case is consistent with the fact that if a trader's limit price is too ambitious in the sense that S'_k is too far below S_k , the probability of an order fill may be quite low. Finally, we note that a nearly identical analysis can also be carried out for a *sell limit order* working with H_{k+1} rather than L_{k+1} .

3) *Modelling of Stop Orders*: This type of order is triggered and automatically converted to a market order if its so-called *stop price* is reached or surpassed. More precisely, for the case of a *buy stop order* shares with stop price $S'_k > S_k$ submitted at the close of T_k , the formula for ΔN_k is the same as in the limit order case and if S'_k is reached or surpassed during T_{k+1} , a market order for these shares is triggered at such time. Then, our standing assumption that the market is

frictionless guarantees a fill at price S'_k . The analysis for this order has a similar flavor to the one used for the limit order and again involves three mutually exclusive cases and we omit some details for brevity. For the case when $O_{k+1} \geq S'_k$, this order is filled at the open for T_{k+1} at price O_{k+1} and we obtain

$$\begin{aligned} M_{k+1} &= M_k - \Delta N_k O_{k+1} \\ &= M_k - \frac{u_k - N_k S_k}{S'_k} O_{k+1}. \end{aligned}$$

If $O_{k+1} < S'_k$ and $H_{k+1} \geq S'_k$, the stop price S'_k is reached during T_{k+1} and we obtain

$$\begin{aligned} M_{k+1} &= M_k - \Delta N_k S'_k \\ &= M_k - \frac{u_k - N_k S_k}{S'_k} S'_k \\ &= M_k - u_k + N_k S_k. \end{aligned}$$

For the third and final case, occurring if $S'_k > H_{k+1}$, the stop order is unfilled during T_{k+1} leading to $N_{k+1} = N_k$ and $M_{k+1} = M_k$ as in the limit order case. Finally, analogous to the limit order case, a nearly identical analysis can also be carried out for a *sell stop order* working with L_{k+1} rather than H_{k+1} playing the pivotal role.

C. Causality Requirement

At the end of stage k , the desired value of shares under control u_k , in practice, must be a causal function of past history variables. Subsequently, with u_k specified and the trader indicating the order type, the ideas in the previous subsections can be used to update the account value. To provide an example, at stage k , suppose u_k is a linear time-invariant feedback on the s most recent account values; i.e., say

$$u_k = K_0 V_k + K_1 V_{k-1} + \dots + K_s V_{k-s}$$

with associated feedback gain vector $K \doteq (K_0, K_1, \dots, K_s)$ leading to satisfaction of the no-leverage condition. The trader can enter a buy order of one of the aforementioned types, and the previously described update equations can be immediately applied.

IV. STOCK TRADING AS A FEEDBACK CONTROL DESIGN PROBLEM

The example above with its gain vector K illustrates how a feedback control mechanism can be used to define a suitable trading strategy or enhance an existing one. Also recalling the Introduction where further motivation for control-inspired trading methods and the related literature is discussed, this section provides some illustrations of this type of integration of automatic control principles into trading schemes. In this context, we first consider an approach which we describe as “model free” and later introduce both model-based and data-driven analyses into the discussion.

To begin, we consider the case of trading a single stock using a linear time-invariant feedback control to determine the desired investment level I_k corresponding to controller u_k . In this section, since only market orders are assumed, consistent with the frictionless market assumption in the previous




Fig. 2. Stock Trading via Linear Feedback

section, we are assured that I_k is instantaneously realized. Instead of using data or a model to make predictions of the future stock prices as a basis for adjustment of I_k as is typically done by practitioners of classical technical analysis,³ we simply introduce a “garden variety” linear time-invariant feedback controller as depicted in Figure 2. This serves a starting point for the more general framework in [8] in next subsection.

Now, with I_0 representing the initial investment, notice that the controller I_k is “model free” in the sense that this investment level at stage k , given by

$$I_k = I_0 + K g_k$$

is based *solely* on the cumulative gain or loss to date g_k with no assumed parameterization of the price dynamics S_k or stock returns X_k . Consistent with this, in the exposition to follow, the only assumption on X_k is that it arises from underlying stochastic process with mean

$$\mu \doteq \mathbb{E}[X_k]$$

for all k . Finally, we note that with $K > 0$ and $I_0 > 0$, this corresponds to going long, which leads to profit when the price S_k is increasing; equivalently, when $X_k > 0$.

For the accounting block in the figure, the cumulative gain-loss is updated as

$$g_{k+1} = g_k + X_k I_k.$$

Furthermore, since g_k and the associated account value V_k satisfy $g_k = V_k - V_0$, we also have update equation

$$V_{k+1} = V_k + X_k I_k.$$

As in stochastic control, the returns X_k represent the primary uncertainty influencing the dynamics of g_k .

A. Simultaneous Long-Short Controller

The takeoff point for this subsection is that fact the linear feedback controller $I_k = I_0 + K g_k$ above can potentially lead to large losses. For example, in a rapid market decline with both I_0 and K being positive, if K is not sufficiently large, the investment level I_k may not be decreasing quickly

³Prior to the last couple of decades, technical analysis and its foundations were strongly criticized by academics based on considerations of market efficiency. More recently, however, these methods have been receiving increasing attention in the literature; e.g., see [51]–[55].

enough to attenuate the fast decay in g_k . With this scenario and others serving as motivation, in [8], a hedging strategy of sorts is introduced. To this end, the so-called *Simultaneous Long-Short (SLS) controller*, capable of both long and short positions, is now described.

Indeed, the SLS controller employs a pair of decoupled linear feedback controllers in parallel. Namely, for fixed gain $K > 0$, using two different investments, $I_{k,L}$ being long and $I_{k,S}$ being short, we define investment levels

$$I_{k,L} = I_0 + Kg_{k,L}; \quad I_{k,S} = -I_0 - KI_{k,S}$$

and overall investment

$$I_k = I_{k,L} + I_{k,S}.$$

Corresponding to these levels are the individual cumulative gains and losses

$$g_{k+1,L} = g_{k,L} + X_k I_{k,L}; \quad g_{k+1,S} = g_{k,S} + X_k I_{k,S},$$

and overall gain-loss function

$$g_k = g_{k,L} + g_{k,S}.$$

In [8], this SLS scheme is shown to have an interesting arbitrage-like property. Namely, the following theorem is established.

Robust Positive Expectation Theorem: *In an idealized frictionless market, the trading gain-loss function resulting from the SLS controller is*

$$\mathbb{E}[g_k] = \frac{I_0}{K} [(1 + K\mu)^k + (1 - K\mu)^k - 2],$$

which is positive for $k > 1$ and all μ non-zero.

B. Remarks on SLS Control and Related Work

The theorem above demonstrates the effectiveness of feedback control in a hedging context and it is also noted that performance of the SLS controller depends strongly on the selection of the feedback gain $K > 0$; e.g., see [11]. For the advanced reader, it is noted that there is an analogy to be made between the behavior of $\mathbb{E}[g_k]$ in the theorem as a function of μ and the behavior of a straddle in options theory. That is, if there is a large price movement either up or down, the SLS controller will typically perform well.

Despite its hedging potential, real-world use of an SLS controller has some negatives: For sideways price movement, even with $\mathbb{E}[X_k]$ non-zero and positive, along some “unlucky” sample paths, significant losses can occur. Another negative associated with the SLS controller relates the fact that real-world stock price dynamics may be highly nonstationary; a fixed gain K while performing well on one price regime may be inadequate for another. This consideration motivates research on control methods other than the simplest SLS scheme as described above. To this end, we point to some initial work along these lines.

One approach in the literature involves beginning with an a priori bound on the X_k obtained from historical data and then determining what SLS gain K to use by robust control




Fig. 3. Empirical Return Distributions for Eight American Stocks

methods; see [15]. By way of illustration, in the box plot in Figure 3, the empirical return distributions of eight well-known American stocks for the trading years 2010–2017 are given. They show the median, quartiles, and outliers, providing a concise summary of data spread, central tendency, and anomalies. Since a typical robust control method places heavy emphasis on the “extreme” returns associated with the dashed min and max lines in the figure, it stands to reason that an “optimal” robust stock-trading controller may be “overly conservative” because the large positive and negative returns which are emphasized may be highly improbable.

Another approach aimed at obtaining a suitable feedback gain K involves the use of data-driven adaptive methods in combination model-based price dynamics. The key ideas underlying such a scheme are depicted in Figure 4 as it might apply to the simple case of finding a standard time-varying linear feedback gain K_k which can be viewed as one of the two legs of an SLS controller. Consistent with the Internal Model Principle [50], the controller identifies local price behaviors to counteract them through backward optimization over a finite-length past window; for further details, see [16]. Finally, we mention the initial work in [14] looking at the use *extremum seeking* techniques aimed at optimizing the output of a dynamic closed-loop system represent the stock-trading algorithm. To conclude, we note that delving deeper into new trading schemes as illustrated by the ones described above is a possible direction for further work. In particular, the adaptive tuning of trading parameters such as the feedback gain K , is a subject currently under investigation. In the next section, instead of concentrating on single stock, we consider a portfolio consisting of multiple stocks and study optimization problems in this context.




Fig. 4. The Adaptive SLS Control Architecture of [16]

V. OPTIMIZATION PROBLEMS ARISING IN QUANTITATIVE PORTFOLIO MANAGEMENT

Optimization methods play a pivotal role in modern quantitative portfolio management. Beginning with this premise, we discuss how optimization theory and related algorithms can be used to efficiently solve various portfolio optimization problems. To this end, we shall delve into classical mean-variance models and their various extensions. These include the Black-Litterman approach, which integrates market equilibrium and investors' views into the analysis, and more recent developments such as stochastic and robust portfolio theory. We aim to equip theoreticians and practitioners with valuable insights and toolkits for optimizing portfolios by examining key models and theories; see [1] and [19].

A. Mean-Variance Model

Central to modern portfolio theory, the Mean-Variance (MV) Optimization formulated by Harry Markowitz in his seminal paper [13] provides a framework for constructing portfolios that maximize expected return for a given level of risk. Specifically, consider a portfolio of m assets, which may include equities, bonds, currencies, and other assets. For $i = 1, 2, \dots, m$, let $X = [X_1, X_2, \dots, X_m]^T$ with X_i being the rate of return on asset i . Take $\mu \doteq \mathbb{E}[X]$ to be the mean return vector.

Let $V_0 > 0$ be the initial account value of the portfolio and V_1 be the subsequent account value, which can be viewed as a system *state*. Consider the investment policy, or a *controller*, for each asset i given by $u_i \doteq K_i V_0$ with a constant feedback gain K_i . Then the *single-period* account value update equation is given by

$$V_1 = V_0 + \sum_{i=1}^m u_i X_i = V_0(1 + K^T X)$$

where $K \doteq [K_1, K_2, \dots, K_m]^T \in \mathbb{R}^m$ is the *feedback gain vector* that corresponds to the portfolio weights. The portfolio return is $r_p(K) \doteq (V_1 - V_0)/V_0 = X^T K$ and the corresponding portfolio mean return is $\mathbb{E}[r_p(K)] = \mu^T K$ and $\text{var}(r_p(K)) = K^T \Sigma K$ where

$$\Sigma \doteq \mathbb{E}[X X^T] - \mu \mu^T$$

is the covariance matrix of the asset returns.

A version of Markowitz's MV optimization model is given by

$$\max_{K \in \mathcal{K}} \{J(K) = \mu^T K - \rho K^T \Sigma K\} \quad (1)$$

where $\rho \geq 0$ is a given *risk aversion constant*⁴ and \mathcal{K} is the admissible set. A typical choice for \mathcal{K} is given by the *unit simplex* constraint

$$\mathcal{K} = \{K \in \mathbb{R}^m : K^T \mathbf{1} = 1, K_i \geq 0, i = 1, \dots, m\} \quad (2)$$

with $\mathbf{1} = [1, 1, \dots, 1]^T \in \mathbb{R}^m$ being the one-vector. The constraint set \mathcal{K} represents that the trades are *long-only*, $K_i \geq 0$ for all i , and *cash-financed*, $K^T \mathbf{1} = 1$. This framework is deeply connected to the *convex quadratic program*, offering efficient computational approaches for finding optimal portfolios; see [37] and [40].

In particular, it is well-known that, if short selling is allowed, that is, by dropping the nonnegativity constraints on K , then obtains

$$\begin{aligned} \max_K \quad & \mu^T K - \rho K^T \Sigma K \\ \text{s.t.} \quad & K^T \mathbf{1} = 1. \end{aligned}$$

That is, Problem (1) reduces to a convex quadratic program with only one linear constraint. For $\rho > 0$, applying the standard Lagrange multiplier technique yields a closed-form solution $K = K^*$ given by

$$K^* = \frac{1}{2\rho} \Sigma^{-1} (\mu + \lambda^* \mathbf{1})$$

where λ^* is the optimal dual variable

$$\lambda^* = \frac{2\rho - \mathbf{1}^T \Sigma^{-1} \mu}{\mathbf{1}^T \Sigma^{-1} \mathbf{1}}.$$

However, portfolio optimization problems always have additional constraints in practice due to market regulations, friction, and investor preferences; see [17] and [20]. For example, if one considers the unit simplex constraint (2), there is no closed-form solution, and one must employ a numerical approach. The next example illustrates the case with a simple three-asset portfolio.

B. Example: Three Asset Portfolio

Consider a portfolio with $n = 2$ risky assets: Nvidia Corporation (Ticker: NVDA) and Advanced Micro Devices, Inc. (Ticker: AMD), along with a risk-free treasury bond yielding a 5% interest rate covering the period from January 2, 2023 to January 2, 2024 with a total of 250 trading days; see Figure 5 for the price trajectories of the two risky assets.

Now, taking the vector $K \doteq [K_0, K_1, K_2]^T$, where $K_0 = 1 - K_1 - K_2$ represents the weight for the risk-free asset, and K_1, K_2 are the weights for NVDA and AMD, respectively, Figure 6 depicts the concave quadratic objective $J(K)$ with respect to (K_1, K_2) using $\rho \doteq 4$. The optimal weight

⁴Choosing the value of ρ depends on the investor's aversion to risk. A higher ρ means prioritizing the minimization of variance over achieving higher returns, while a lower ρ focuses more on maximizing returns at the risk of higher variance.

is obtained at $(K_1^*, K_2^*) \approx (0.676, 0.054)$, occurring on the interior of the unit simplex constraint set, which leads to corresponding optim $K_0^* = 1 - K_1^* - K_2^* \approx 0.269$.




Fig. 5. NVDA and AMD Prices: January 02, 2023 to January 02, 2024




Fig. 6. Three Asset MV Portfolio and Objective Value $J(K)$

It should be noted that the MV portfolio tends to overweight assets with large expected returns and low variance while underweighting assets with low expected returns and high variance. Indeed, taking the risk-aversion constant to be $\rho = 2$, then the optimum $K = (0, 1, 0)$, an overly concentrated weight on NVDA, is seen. Thus, MV portfolios may not be well diversified; e.g., see [35]. Moreover, the estimation error may largely impact the performance of the optimal MV portfolio. The uncertainty of returns tends to exert more influence than the associated risk in MV optimization, a point further elaborated upon in [36].

C. The Black-Litterman Model

Extending the MV model, the Black-Litterman (BL) model, introduced by [24], incorporates investors' *subjective views* and market equilibrium into portfolio optimization.

This model addresses practical concerns with the traditional MV model, particularly its sensitivity to input estimates. The BL model connects the optimal weighted least-square estimate and a convex quadratic program.

The basic idea of the BL model is to tilt the market equilibrium returns to incorporate an investor's views via a *Bayesian* approach.⁵ For example, consider $\mu \in \mathbb{R}^m$ as the unknown true expected return. The BL model proposes an estimate of μ using the market equilibrium return such as the capital asset pricing model (CAPM) or factor models, denoted by π , see [20], [24] and [44], as follows:

$$\pi = \mu + \varepsilon_\pi; \quad \varepsilon_\pi \sim \mathcal{N}(\mathbf{0}, \gamma\Sigma)$$

for some small parameter $\gamma \ll 1$, e.g., see [24] and $\gamma \in [0.01, 0.05]$. In particular, one can think of $\gamma\Sigma$ as our confidence in estimating the equilibrium expected returns. A small γ implies high confidence in our equilibrium estimates, and vice versa.

Expressing Investor's Views. Investors often hold specific *views* on the performance of certain assets, either in absolute terms or relative to other assets. The BL model formalizes these views, denoted by $q \in \mathbb{R}^d$, as:

$$q = P\mu + \varepsilon_q; \quad \varepsilon_q \sim \mathcal{N}(\mathbf{0}, \Omega)$$

where $P \in \mathbb{R}^{d \times m}$ is the parameter matrix of views, ε_q represents the normally distributed degree of confidence in the views, and $\Omega \in \mathbb{R}^{d \times d}$ expressing the confidence in the views, which is typically set to be a diagonal matrix, indicating independent error terms across individual views. In extreme cases, a view can have zero variance, signifying absolute confidence.

Example: Assigning the Investor's Views. Consider the prior example regarding an asset allocation problem with three assets: Risk-free treasury bonds, NVDA, and AMD. Suppose we have two *subjective* views for the current ongoing year:

- Return on treasury bond will be 4%.
- The NVDA will outperform the AMD by 2%.

Let $\mu = [\mu_0, \mu_1, \mu_2]^\top$ be the mean returns for the three assets unknown to the trader. Then, the two views above can be expressed as

$$\begin{aligned} 0.04 &= \mu_0 + \varepsilon_1; \\ 0.02 &= \mu_1 - \mu_2 + \varepsilon_2. \end{aligned}$$

One can write the above to the matrix-vector form as $q = P\mu + \varepsilon_q$ where $q \doteq [0.04, 0.02]^\top$, $\varepsilon_q = (\varepsilon_1, \varepsilon_2)^\top$, and

$$P = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & -1 \end{bmatrix}.$$

Note that the error terms ε_1 and ε_2 do not explicitly enter into the BL model—but their variances do.

⁵The classical approach to estimating future expected returns assumes that the "true" expected returns and covariances of returns are unknown and *fixed*. In contrast, the Bayesian approach assumes that the "true" expected returns are unknown and *random*.

Merging Investor's View and Market Equilibrium. To properly integrate the investor's views with the prior market equilibrium, we recall that

$$\begin{cases} \pi = \mu + \varepsilon_\pi, & \varepsilon_\pi \sim \mathcal{N}(0, \gamma\Sigma) \\ q = P\mu + \varepsilon_q, & \varepsilon_q \sim \mathcal{N}(0, \Omega). \end{cases}$$

Let $y \doteq [\pi \ q]^\top$, $M \doteq \begin{bmatrix} I_{m \times m} \\ P \end{bmatrix}$ and $V \doteq \begin{bmatrix} Q & \mathbf{0} \\ \mathbf{0} & \Omega \end{bmatrix}$ with $Q = \gamma\Sigma$. Then, it follows that

$$y = M\mu + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, V)$$

where $I_{m \times m}$ is an identity matrix and $P \in \mathbb{R}^{d \times m}$. Once we have merged the market equilibrium and the investors' views, we formulate a weighted least-squares (WLS) problem

$$\min_{\mu} (y - M\mu)^\top V^{-1} (y - M\mu).$$

whose solution, denoted by $\hat{\mu}$, is given by

$$\hat{\mu} = \pi + QP^\top (PQP^\top + \Omega)^{-1} (q - P\pi). \quad (3)$$

Moreover, via the Bayesian approach, e.g., [38], one can also obtain the BL estimate for the covariance matrix as follows

$$\hat{\Sigma} = \Sigma + (Q^{-1} + P^\top \Omega^{-1} P)^{-1}. \quad (4)$$

It should be noted that the BL model's reliance on normally distributed error terms is a critical assumption for deriving closed-form solutions like Equation (3). However, in cases where this assumption does not hold, alternative numerical methods, including nonlinear programming algorithms, may be required to find suitable solutions; see [19]. Having obtained the view-updated mean $\hat{\mu}$ and covariance matrix $\hat{\Sigma}$, one can then resolve the problem characterized by Equation (1) using these new estimates. However, it should be noted that the result here is entirely an in-sample result intended to demonstrate the construction of the Markowitz model and its possible extension. Therefore, in practice, out-of-sample testing should be the next step.

Remark on Further Generalization. The MV and BL model can be recast as a more general model under the umbrella of *stochastic optimization* (SO); see [21], which incorporates uncertainty via randomness within the optimization model itself. This is typically through the use of chance constraints or stochastic objectives. Specifically, let $U_k(\cdot)$ be a *risk-averse utility* function, which is concave and strictly increasing. Then, the general SO problem for a *finite-horizon* can be expressed as

$$\sup_{\{u_{k,i}\}} \sum_{k=0}^{n-1} \mathbb{E}[U_k(V_k)] + \mathbb{E}[U_n(V_n)] \quad (5)$$

where⁶ $V_{k+1} = V_k + \sum_{i=1}^m X_{k,i} u_{k,i}$ for $k = 0, \dots, n-1$ and $i = 1, \dots, m$. That is, the account value at the next time step V_{k+1} evolves based on the current account value V_k and the decision variable u_k influenced by returns X_k .

⁶Here, admitting some abuse of notation, we use $X_{k,i}$ to denote the return of the i th asset over period k .

VI. ON KELLY'S CRITERION FOR STOCK TRADING

In this section, we focus on a specific instance of optimization problem characterized by Equation (5), where only one stock is traded,⁷ and the performance index, viewed in the context of the previous section has running utilities

$$U_0 = U_1 = \dots = U_{n-1} \equiv 0$$

and terminal utility being the *expected logarithmic growth* (ELG) of the trader's account value, to be maximized and defined as

$$\text{ELG} \doteq \frac{1}{n} \mathbb{E} \left[\log \frac{V_n}{V_0} \right].$$

The use of the ELG is standard in investment and gambling literature; see [1], [9], [23] and [25]. By maximizing this quantity, a trader strikes a balance between wealth growth and risk mitigation; e.g., no trading strategy with a positive chance of ruining can be optimal. Another reason for the use of ELG, as noted in [25], is that in gambling and stock trading one typically deals with *reinvestments*, and the account value after n periods is the product of several factors, one per trading period. While dealing with the expected value of such products is inconvenient, the expected logarithm of a product naturally decomposes into the sum of expected logarithms, simplifying the determination of the optimal strategy.

A. Kelly's Criterion for a Known Distribution of Returns

Consider first a hypothetical scenario where the returns $X_k = (S_{k+1} - S_k)/S_k$ in the account update equation

$$V_{k+1} = V_k + u_k X_k \quad (6)$$

are i.i.d. with a known distribution. This distribution is assumed to be *discrete*⁸ with s atoms x_1, \dots, x_s and $p_i > 0 \forall i$ denoting the probability that $X_k = x_i$ and $p_1 + \dots + p_s = 1$.

The idea of a trading strategy maximizing the ELG traces back to the seminal work by John Larry Kelly [22], a scientist at Bell Labs with an interest in gambling theory, who uncovered a relationship between gambling and information theory. Initially developed for betting in sports and games of chance, Kelly's strategy, essentially a simple feedback controller, has since been applied in stock trading and portfolio management [23] and [26].

1) Kelly's Criterion: An Optimal Linear Controller:

Consider first the problem of finding the optimal trading strategy among all static proportional controllers $u_k = KV_k$, satisfying the no-leverage assumption $|u_k| \leq V_k$; equivalently $|K| \leq 1$. Note that $K > 0$ corresponds to a long trading strategy, whereas $K < 0$ represents a short trading strategy.

Substituting $u_k = KV_k$ into the account value update equation, one finds that $V_n = V_0(1 + KX_0) \dots (1 + KX_{n-1})$, and hence, the ELG can be written as follows

$$\text{ELG}(K) = p_1 \log(1 + Kx_1) + \dots + p_s \log(1 + Kx_s).$$

⁷Kelly's criterion can easily be extended to portfolio optimization; for further details, interested readers can refer to [25] and [26].

⁸Kelly's criterion can be given for the continuous distributions as well. However, in practice, a stock trading strategy must be data-driven, relying on discrete empirical distributions estimated from observed sample paths.

In practice, the returns are usually small⁹ with $|x_i| < 1$ for all i , and hence $\text{ELG}(K)$ is well-defined under the no-leverage condition $|K| \leq 1$. By noticing that \log is strictly concave, the maximizer exists and is unique

$$K^* = \arg \max_{K \in [-1, 1]} \sum_{i=1}^s p_i \log(1 + Kx_i). \quad (7)$$

The ELG-optimal linear time-invariant feedback controller $u_k = K^*V_k$ is known as *Kelly's strategy*, and the optimality condition (7) is referred to as *Kelly's criterion*. In practice, finding a maximizer for $\text{ELG}(K)$ is straightforward; e.g., one can generate a plot of $\text{ELG}(K)$ in appropriate software and visually identify the point of maximum. However, for more general versions of a problem involving a portfolio and a vector K more sophisticated methods, such as convex programming are typically used. Finally, note that Kelly's criterion generalizes to the long-only (respectively, short-only) trading, replacing $[-1, 1]$ in (7) with $[0, 1]$ (respectively, $[-1, 0]$).

2) *Kelly's Criterion for the Two-Atom Distribution*: To illustrate Kelly's criterion, we examine an idealized scenario where X_k is constrained to only two values ($s = 2$): $x_1 > 0$, occurring with a probability of $p_1 = p > 0$, and $x_2 < 0$, occurring with a probability of $p_2 = 1 - p$. Similar to a straightforward coin-flipping game falling under the umbrella of [22] and [27], in this scenario, for any $K \neq 0$, one of these two outcomes results in a loss for the trader (if $Kx_i < 0$), while the other outcome increases the trader's account value (if $Kx_i > 0$).

Recalling Kelly's criterion tells us that the optimal gambling strategy $u_k = K^*V_k$ is found by maximizing $\text{ELG}(K) = p \log(1 + Kx_1) + (1 - p) \log(1 + Kx_2)$ over $K \in [-1, 1]$, using concavity of the ELG, Kelly's controller K^* can be found as the solution of the equation $\text{ELG}'(K) = 0$, i.e.,

$$K^* = K_p^* \doteq -\frac{p}{x_2} - \frac{1-p}{x_1} = \frac{p}{|x_2|} - \frac{1-p}{x_1},$$

provided that K_p^* belongs to $[-1, 1]$. Otherwise, $K^* = 1$ if $K_p^* > 1$ and $K^* = -1$ if $K_p^* < -1$.

A special case of this example arises in a coin-flipping game with an even payoff, where the gambler stakes a portion of wealth $|u_k| \leq V_k$ at each round $k = 0, \dots, n - 1$. We interpret values $u_k > 0$ as betting on heads, whereas $u_k < 0$ means betting on tails; the gambler skips the k -th round of betting if $u_k = 0$. The wager is added to the gambler's wealth if the coin side is correctly guessed ($x_1 = 1$) and forfeited otherwise ($x_2 = -1$). In this scenario, we arrive at the celebrated Kelly's formula $K^* = 2p - 1$ for the optimum in [22]. The gambler should bet on heads if $p > 1/2$ and on tails when $p < 1/2$, with the proportion of wealth wagered being $|2p - 1|$.

In this subsection, we briefly outline important properties and some generalizations of Kelly's criterion.

⁹Of particular note, positivity of prices $S_k > 0$ for all k guarantees that the associated returns $X_k = (S_{k+1} - S_k)/S_k > -1$.

3) *More Advanced Kelly Betting Topics*: In this subsection, we first describe two important properties of Kelly's ELG maximizing strategy and then discuss a possible refinement of the theory with potential appeal to stock market practitioners: a data driven implementation of the optimal betting scheme. Finally, we describe a recent Kelly betting result by two of the coauthors of this tutorial paper. This involves a robustness scenario and conditions under which nonlinear controller, in lieu of Kelly's linear one, is preferred.

One key question regarding the Kelly optimum is whether the ELG performance level associated with the linear feedback scheme described above can be further increased by allowing nonlinear and time-varying trading strategies. Surprisingly, the answer is negative. Kelly's strategy cannot be outperformed by any admissible (causal) trading strategy, such as, e.g., a time-varying nonlinear controller $u_k = \tilde{K}_k(V_k)V_k$ or a more general controller with memory $u_k = \tilde{K}_k(V_0, \dots, V_k)V_k$. The mathematical proof (addressing a more general case of Kelly's portfolio trading strategies) is available, e.g., in [25]; this proof is inspired by Kelly's pioneering work [22].

The second important property of Kelly's solution pertains to asymptotic behavior. It can be proven that for long trading intervals ($n \rightarrow \infty$) Kelly's controller $u_k = K^*V_k$ outperforms any other admissible trading strategy *almost surely*. More formally, comparing the final account value V_n under any admissible strategy with the final value V_n^* delivered by Kelly's strategy for the same sample path, it is shown in [25, Theorem 15.3.1] that

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log(V_n/V_n^*) \leq 0$$

with probability one.

As far as the stock market practitioner is concerned, one undesirable feature of the Kelly theory is the requirement for information about the return variable distribution, an assumption that is obviously unrealistic in trading and needs to be discarded. A practitioner's approach, as illustrated in the simulation example in Section VIII, is to work with an empirical distribution for X_k in the optimization problem (7) (essentially, rendering the controller nonlinear and time-varying). Namely, one fixes $s \geq 1$ as the *lookback parameter* (the retrospective observation window). When deciding on the control on interval T_k , we allocate probability mass $1/s$ to each of the observed returns $X_{k-1}, X_{k-2}, \dots, X_{k-s}$ (in practice, it is almost impossible to see the exactly the same return twice). After that, we compute Kelly's controller $u_k = K_k^*V_k$ for interval T_k , using the constructed discrete distribution instead of the actual distribution for X_k , that is, we solve (7) with $x_i \doteq X_{k-i}$ and $p_i = 1/s$, $i = 1, \dots, s$. Note that the optimal level of investment K_k^* will not be constant as it depends on the historical data.

A modification of this data-driven approach is the method of adaptive Kelly betting described in [29], which shares similarities with the adaptive trading strategies discussed in Section IV. In this approach, it is assumed that X_k are i.i.d. random variables containing a parametric uncertainty; e.g.,

we play a coin-flipping game with uncertain parameter being the probability of heads p . Using statistical inference techniques such as Bayesian and maximum likelihood estimates, one can then estimate the uncertain parameters based on the full observed sample path X_0, \dots, X_{k-1} and use the obtained estimates in Kelly's criterion. One can expect that, as $n \rightarrow \infty$, the resulting controller approximates the "exact" Kelly's controller corresponding to the true parameter values.

An alternative *robust control* approach to the ELG-optimal trading can be used based on the theory developed in the recent work [28] and references cited therein. Whereas this approach, motivated by the theory of distributional robustness, deals with parametric uncertainty in the return distribution, the unknown parameters are not iteratively estimated. Instead, a causal controller $u_k = K_k(X)V_k$ is found that maximizes an integral ELG performance measure. Unlike the classical Kelly criterion, this optimal controller proves to be nonlinear even for the simplest of coin-flipping games.

VII. AN INTRODUCTION TO THE LIMIT ORDER BOOK

For any given stock, the Limit Order Book (LOB) is essentially a collection of all open limit orders and maintained on a server of the stock exchange. As limit orders either arrive or are filled dynamically in time, the LOB is updated to reflect these changes. As seen in the sequel, the Limit Order Book gives a clear picture of the supply and demand for a given stock and other insights into market depth and liquidity. As far as the control community work is concerned, this section is motivated by the brief introductions to the LOB in [41], [42] and the recent research papers [43] and [45]. Current research, be it in the control community or elsewhere is motivated by the following consideration: Given the nanosecond time stamping of transactions and the voluminous data sets associated with keeping track of orders with this time scale, the LOB is a mecca of sorts for both high-frequency traders and researchers. For example, relatively recent surveys such as [34] and [46] motivate research on the price movement prediction and potential profitability of various high-frequency trading algorithms exploiting the nanosecond-level time stamping as orders enter into the LOB.

This section offers a detailed model of the market mechanism that translates messages from traders, including offers to buy and sell, into the data structure called the Limit Order Book (LOB) and the time series we call "price." Given the target group of attendees, we take a control-theoretic point of view and develop modelling equations in a state space.

A. Market Mechanisms and the LOB

Market mechanisms are discrete event dynamic systems that receive messages from prospective traders and use them to update the LOB [30], [32], [33]. The LOB is the state of the market, a snapshot at any particular time of active orders in the market from investors collectively wishing to buy and sell a certain number of shares of a given stock at a particular price. The orders to buy are called *bids* while the orders to sell are called *asks*, and the LOB for any particular asset

traded on the market can be visualized as a histogram of the aggregate number of shares being "bid" or "asked" at various prices. This is depicted in Figure 7 for the case of General Electric (GE) on February 4, 2014 at time 9 : 55 : 27am with the bids in red and asks in blue. Notice the offer to buy a large volume (over 5000) of shares at the very low price of \$24.42 (red bar on the far left) in spite of similar offers at higher prices. All these offers contribute to the overall *shape* of the book.

Although actual markets need to differentiate between multiple orders to buy or sell at the same price in order to 1) follow specific precedence rules about whose order gets filled first, and 2) match filled orders to specific traders, here we will simplify the discussion by only considering the aggregate number of orders at any specific price.




Fig. 7. The LOB for GE; see Text for Details

B. LOB State-Space Modelling

We begin by defining the range of discrete prices, \mathbb{P} , available for admissible trades. For example, these could be \$0.01, \$0.02, ..., \$100,000.00, or any other increment and upper bound deemed acceptable to model the potential price of a single share of stock. We let $\ell \doteq |\mathbb{P}|$, and note that the LOB at any given time k will be composed of two nonnegative integer vectors, $z^b(k) \in \mathbb{Z}^{\ell, (0,+)}$ and $z^a(k) \in \mathbb{Z}^{\ell, (0,+)}$, respectively, representing in Figure 7 the number of shares being sought on the bid side (red) and the number of shares being offered on the ask side (blue) of the LOB at each price point. Hence, to describe the full state of the LOB, we define

$$z(k) \doteq \begin{bmatrix} z^b(k) & z^a(k) \end{bmatrix} \in \mathbb{Z}^{\ell \times 2, (0,+)}.$$

Note that, although conceptually simple, in practice this representation is not efficient because $z(k)$ will, in general, be very sparse, since most stocks do not receive offers to buy or sell at every conceivable price. Nevertheless, we will adopt it here for pedagogical purposes.

It turns out that characterizing LOB dynamics is easier when using *cumulative* representations of the vectors $z^{(a,b)}(k)$. For this, we define the $\ell \times \ell$ matrices:

$$U \doteq \begin{bmatrix} 1 & 1 & \dots & 1 & 1 \\ 0 & 1 & \dots & 1 & 1 \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & \dots & 1 & 1 \\ 0 & 0 & \dots & 0 & 1 \end{bmatrix}, \quad L \doteq \begin{bmatrix} 1 & 0 & \dots & 0 & 0 \\ 1 & 1 & \dots & 0 & 0 \\ \vdots & & \ddots & & \vdots \\ 1 & 1 & \dots & 1 & 0 \\ 1 & 1 & \dots & 1 & 1 \end{bmatrix}$$

and change bases to define the LOB state variables:

$$x^b(k) \doteq Uz^b(k); \quad x^a(k) \doteq Lz^a(k)$$

which are combined to obtain

$$x(k) \doteq \begin{bmatrix} x^b(k) & x^a(k) \end{bmatrix} \in \mathbb{Z}^{\ell \times 2, (0,+)}.$$

We say the book is *balanced* when $x_p^b \neq 0$ implies $x_p^a = 0$ and $x_p^a \neq 0$ implies $x_p^b = 0$ for all $p \in \mathbb{P}$.

To characterize order book dynamics, we adopt the following convention for the min and max operations: Given a real vector x , we use $\min(x)$ and $\max(x)$ to denote the minimum and maximum of its components, respectively. Given $c \in \mathbb{R}$ and a matrix A , let $\min(A, c)$ denote a matrix the same size as A with every entry given by $\min(a_{i,j}, c)$; the max operation is defined similarly. We use $\mathbf{1}^{s \times t}$ to denote an $s \times t$ matrix of all ones and also define a mapping, $f: \mathbb{Z}^{\ell \times 2} \rightarrow \mathbb{Z}^{\ell \times 2}$, called the *balancing function*, by

$$f(y) \doteq \max(y - \max(\min(y^b, y^a)\mathbf{1}^{\ell \times 2}, 0)) \quad (8)$$

with argument, $y = \begin{bmatrix} y^b & y^a \end{bmatrix} \in \mathbb{Z}^{\ell \times 2}$. Notice that f is identity when its argument is nonnegative and balanced, since $\max(\min(y^b, y^a)\mathbf{1}^{\ell \times 2}, 0) = 0 \times \mathbf{1}^{\ell \times 2} = \mathbf{0}^{\ell \times 2}$ in that case because every row of a balanced y has at least one zero element.

With this balancing function defined, our discrete-time state equation characterizing the order book dynamics, is given by

$$x(k+1) = f(\max(x(k) + u(k), 0)) \quad (9)$$

where $u(k) = \begin{bmatrix} u^b(k) & u^a(k) \end{bmatrix} \in \mathbb{Z}^{\ell \times 2}$, $u^b(k), u^a(k) \in \mathbb{Z}^{\ell}$ is an encoding of the order type received by the market at time step k . Loosely speaking, these order book dynamics (9) simply add incoming orders to the current cumulative state of the book and then balances using the balancing function, (8), to fill any overlapping orders and remove them from the book. The maximizations with zero simply keep the state representation, $x(k)$, nonnegative, as illustrated through the example to follow.

C. Messages and Level II ITCH Data

The LOB for a given stock reflects the aggregation of all the individual orders a market receives for that stock. These individual orders, however, are communicated as *messages* using a very particular data format and exchange protocol called the *ITCH* protocol.

NASDAQ introduced the ITCH protocol in January 2000. It was developed, and evolved, to become a lightweight data format to efficiently use limited bandwidth resources between traders and markets. Today, the ITCH Protocol is used across both Nasdaq-owned venues as well as other, non-Nasdaq exchanges, and it has become one of the industry's de facto standards for market data feeds.

In general, ITCH data are timestamped sequences of messages, measured in nanoseconds after midnight, revealing the nature of various orders arriving at the exchange. Messages related to a specific stock then accumulate to effectively

create the LOB for that stock. Although a detailed description of the messages or the protocol is beyond the scope of this tutorial, the basic kinds of messages include things like *add orders*, which adds either a buy or a sell order to the book, or *modify* messages, which indicate things like when an execution occurs or when an order is partially or fully cancelled. Note that it does not cost a trader anything to place or remove orders on the book; costs are only incurred by transactions.

Since ITCH data are simply delivered as a sequence of messages ordered by timestamp, called *Level II data*, it takes some effort to convert this sequence into the data structure we call the LOB; e.g., see [31] as in Figure 7. Here, consistent with our state-space modelling, we encode each message as a pair of vectors in \mathbb{Z}^n to model the effect of each order type on the LOB, translating the message stream into the input time series $u(k)$, driving the dynamics. Indeed, beginning with the simple *no action* case, we take $u^b(k) = u^a(k) = 0$. Now, for brevity, we provide three more examples how this translation of the message stream works where \mathbb{Z} denotes the integers, $q \in \mathbb{Z}$ is the number of shares ordered, and $*$ represents standard multiplication: Now, we provide three more examples illustrating how this translation of the message stream works:

Add a bid Order for q Shares at Price p

$$\begin{aligned} u_i^b(k) &= \begin{cases} q & i \leq p \\ 0 & i > p \end{cases} & \forall i \in \mathbb{P}, \\ u_j^a(k) &= 0 & \forall j \in \mathbb{P}. \end{aligned}$$

Cancel Ask Order for q Shares at Price p

$$\begin{aligned} u_i^b(k) &= 0 & \forall i \in \mathbb{P}, \\ u_j^a(k) &= \begin{cases} 0 & j < p \\ -q & j \geq p \end{cases} & \forall j \in \mathbb{P}. \end{aligned}$$

Add a Market Order to Sell q Shares at the Best Price

$$\begin{aligned} u_i^b(k) &= -q & \forall i \in \mathbb{P}, \\ u_j^a(k) &= 0 & \forall j \in \mathbb{P}. \end{aligned}$$

D. LOB Example

To illustrate LOB dynamics, consider the LOB for a single stock, and suppose that market rules define the price of this stock at any time step k as the price of the last execution. For this illustration, let $\mathbb{P} = \{\$0.01, \$0.02, \$0.03, \$0.04, \$0.05, \$0.06\}$ and suppose that at time $k = 0$ the price of the stock is $S(0) = \$0.05$, with the initial book given by:

$$z(0) = \begin{array}{cc|l} \text{bid} & \text{ask} & \\ \hline 0 & 0 & \leftarrow \$0.01 \\ 2 & 0 & \leftarrow \$0.02 \\ 1 & 0 & \leftarrow \$0.03 \\ 3 & 0 & \leftarrow \$0.04 \\ 0 & 1 & \leftarrow \$0.05 \\ 0 & 2 & \leftarrow \$0.06 \end{array} \Rightarrow x(0) = \begin{bmatrix} 6 & 0 \\ 6 & 0 \\ 4 & 0 \\ 3 & 0 \\ 0 & 1 \\ 0 & 3 \end{bmatrix},$$

corresponding to bids of size 2, 1, and 3 at prices \$0.02, \$0.03, and \$0.04 respectively, as well as asks of size 1 and 2 at prices \$0.05 and \$0.06. This order book is balanced, since there are no overlapping bids and asks, and the *bid-ask spread*, or the difference between the price of the highest bid and lowest ask, is \$0.01.

At the same time, suppose that a message arrives to add a bid order for 1 share at price \$0.01. This is encoded as a matrix $u(0)$ and, following the calculations in (9) and (8), we find the cumulative state $x(1)$ and the corresponding shape of the LOB, $z(1)$, as:

$$u(0) = \begin{bmatrix} 1 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \Rightarrow x(1) = \begin{bmatrix} 7 & 0 \\ 6 & 0 \\ 4 & 0 \\ 3 & 0 \\ 0 & 1 \\ 0 & 3 \end{bmatrix} \Rightarrow z(1) = \begin{bmatrix} 1 & 0 \\ 2 & 0 \\ 1 & 0 \\ 3 & 0 \\ 0 & 1 \\ 0 & 2 \end{bmatrix}$$

Notice that the add order effectively just added onto the order book, i.e. $z(1) = z(0) + u(0)$, and since no executions were involved, the price (of the last execution) remains unchanged: $S(1) = \$0.05$.

However, suppose that at the next time step a market order to sell 5 shares arrives. These orders must be filled, if possible, at the best possible price, and, since the last reported price (for at least the last two time steps) has been \$0.05, the trader may be expecting to sell these five shares for a total price of \$0.25. Encoding the order, we see that

$$x(2) = f \left(\max \left(\begin{bmatrix} 7 & 0 \\ 6 & 0 \\ 4 & 0 \\ 3 & 0 \\ 0 & 1 \\ 0 & 3 \end{bmatrix} + \begin{bmatrix} -5 & 0 \\ -5 & 0 \\ -5 & 0 \\ -5 & 0 \\ -5 & 0 \\ -5 & 0 \end{bmatrix}, 0 \right) \right)$$

$$= f \left(\begin{bmatrix} 2 & 0 \\ 1 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 1 \\ 0 & 3 \end{bmatrix} \right) = \begin{bmatrix} 2 & 0 \\ 1 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 1 \\ 0 & 3 \end{bmatrix} \Rightarrow z(2) = \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 1 \\ 0 & 2 \end{bmatrix}$$

Here we see that $z(2) \neq z(1) + u(1)$, and using the cumulative state representation for $x(k)$ made the dynamics of *walking the book* easy to execute. What happened was that the first 3 shares of the market order were bought by the highest bid at a price of \$0.04, but then the next best price was one share at \$0.03, and the last share was filled at a price of \$0.02, for a total price of \$0.17 instead of the expected \$0.25. The last execution is now the last share filled, so $S(2) = \$0.02$, jumping from $S(1) = S(0) = \$0.05$.

Seeing the price, which has been steadily at \$0.05 a share suddenly jump down to \$0.02 a share, a trader who decides that this stock is now undervalued may decide to buy a few shares. Of course, in practice traders cannot receive and act on market information this fast, but continuing with this situation for pedagogical reasons, we suppose a market

order is placed to buy two shares. This being the case, a straightforward but lengthy calculation yields

$$x(3) = f \left(\begin{bmatrix} 2 & 0 \\ 1 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 1 \end{bmatrix} \right) = \begin{bmatrix} 2 & 0 \\ 1 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 1 \end{bmatrix} \Rightarrow z(3) = \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 1 \end{bmatrix}$$

and we see that the order is filled at the best available prices, which were \$0.05 for the first share, and \$0.06 for the second, totalling \$0.11 for the order. This is quite different from the \$0.04 the investor may have been expecting when placing the order, and we see that although \$0.01 of the discrepancy is due to the order “walking the book” and \$0.06 of the discrepancy comes from both shares of the order jumping the \$0.03 bid-ask spread.

So far, none of the incoming orders have unbalanced the book, but the re-balancing dynamics can easily be seen if we suppose that $u(3)$ is a limit bid order for 5 shares at \$0.04, and $u(4)$ is a limit ask order for 2 shares at \$0.03. Again, by a straightforward but lengthy calculation, we obtain

$$x(5) = \max \left(\begin{bmatrix} 7 & 0 \\ 6 & 0 \\ 5 & 2 \\ 5 & 2 \\ 0 & 2 \\ 0 & 3 \end{bmatrix} - \begin{bmatrix} 2 & 2 \\ 2 & 2 \\ 2 & 2 \\ 2 & 2 \\ 2 & 2 \\ 2 & 2 \end{bmatrix}, 0 \right) = \begin{bmatrix} 5 & 0 \\ 4 & 0 \\ 3 & 0 \\ 3 & 0 \\ 0 & 0 \\ 0 & 1 \end{bmatrix}$$

which demonstrates how the dynamics automatically re-balance the book. Although the ask was placed at \$0.03, there was already a better offer on the book, so it was filled at \$0.04, showing us that even though $S(4)$ remained unchanged from $S(3)$ at \$0.06, $S(5) = \$0.04$. This example illustrates how the state model from (9) and (8) calculate changes to the LOB as new orders arrive, and how subtle dynamics like walking the book and jumping the bid-ask spread can lead to unexpected outcomes and volatile prices.

VIII. BACKTESTING: AN ILLUSTRATIVE EXAMPLE

Backtesting is a process for evaluating the performance of a trading scheme based on historical data. For instance, for backtesting a daily strategy, sources such as Yahoo Finance, Google Finance and Quandl will suffice. For a tick-by-tick strategy, a large number of universities have licensed access to the Wharton School’s WRDS data. For backtesting of higher frequency trading schemes, costly NASDAQ ITCH protocol data can be required. At the outset, for a given trading strategy, it is important to distinguish between demonstration of “mechanics” of backtesting it versus judging whether it is likely to be an “excellent performer” when implemented in the marketplace rather than with historical data; here our simulation is being used solely for demonstration of mechanics purposes. In order to make a judgment whether the strategies we consider are “ready for prime time,” one would need to perform a large number of backtests with data covering many different scenarios — different time windows



Fig. 8. Prices for Grayscale Bitcoin ETF: Dec.12, 2021 to Jan. 10,2024

including both bull and bear markets, periods when events such as earnings or interest changes take place, times when geopolitical turmoil takes place, to name a few.

A. Performance Metrics for Account Value Plots

For a given account value plot, there are a wide variety of performance metrics which can be used for evaluation purposes. Depending on the underlying risk preferences of a trader, one metric may be preferred to another. For example, in the pursuit of a high return, some traders may be willing to tolerate a large *Maximum Percentage Drawdown*

$$d_{max} = \max_{k < n, i > k} \frac{V_k - V_i}{V_k}$$

along the way and others may not. We illustrate use of this metric in the simulations to follow, along with the classical *Raw Return*:

$$R = \frac{V_n - V_0}{V_0}$$

to evaluate our V -plot performance. Although not considered in our simulations, the reader would also be well advised to learn about the so-called *Sharpe Ratio (SR)* which was introduced in a widely celebrated 1966 paper [48] and has been used extensively ever since to evaluate the performance of money managers.

B. Trading the Grayscale Bitcoin ETF

In this section, per discussion above, to demonstrate the *mechanics of backtesting*, we consider the following simple scenario: For the time period from December 8, 2021 to January 10, 2024, beginning with Yahoo daily data for the Grayscale Bitcoin ETF whose ticker is GBTF, we examine the performance of three of the trading strategies in earlier sections. In all of the price and account value plots to follow, the units assumed are U.S. dollars. The starting point for our simulation is the daily closing prices for GBTC over the 526 days under consideration; see Figure 8. Each of the three cases to follow is a variation on a linear feedback controller — each with its own distinctive features. To make

these backtests more consistent with simulations performed by a practitioner, we assume that “idle cash” held overnight in a trader’s account receives interest at an annual risk-free interest rate of 3.75% which is converted to a daily rate $r = (1 + 0.0375)^{1/252} - 1 \approx 1.461 \cdot 10^{-4}$ for simulation purposes. This averaged rate, obtained by examination of U.S. Treasury data, is used for simplicity, and the simulations to follow are rather insensitive to use of r in lieu of more accurate daily rate r_k . When we implement the trading schemes in Sections III, IV and VI, the updated equation for V_{k+1} has one extra term attributable to this interest. For example, for the trader using the linear feedback $I_k = I_0 + Kg_k$ in Section IV, the account value update equation becomes

$$V_{k+1} = V_k + I_k X_k + (V_k - I_k)r.$$

C. The Limit Order Trader

The scenario in this first simulation is as follows: Imagine a trader who is long-only and has opportunity to buy new shares or sell existing shares at the market close of each day. On Day k , a limit is entered with the goal of bringing the account to a desired investment level specified by a linear feedback control of the form $u_k = KV_k$. With C_k being the closing price, a price improvement of 2% is being sought on both buy and sell orders; i.e., in Section III, we take $S'_k = 0.98C_k$ on buy orders and $S'_k = 1.02C_k$ on sell orders. The results of the simulation, shown for illustrative values $K = 0.05, 0.25, 0.50$ are seen in Figure 9.




Fig. 9. Account Value for Limit Order Trader

For the performance metrics R and d_{max} given above, we obtained:

- for $K = 0.05$: $R \approx 0.06$ and $d_{max} \approx 0.02$;
- for $K = 0.25$: $R \approx 0.09$ and $d_{max} \approx 0.20$;
- for $K = 0.50$: $R \approx 0.14$ and $d_{max} \approx 0.53$.

D. The Linear Feedback Trader

The scenario in this second simulation is as follows: A trader is going long-only again, but this time, at the close of each day k , a market order at price $S_k = C_k$ instead of a limit order is used. Again, a linear feedback control of the

form $u_k = KV_k$ is applied. The results of the simulation, shown for illustrative values $K = 0.30, 0.60, 0.90$ are seen in Figure 10.




Fig. 10. Account Value for Linear Feedback Trader

For the performance metrics R and d_{max} given above, we obtained:

- for $K = 0.3$: $R \approx 0.22$ and $d_{max} \approx 0.58$;
- for $K = 0.6$: $R \approx 0.15$ and $d_{max} \approx 0.49$;
- for $K = 0.9$: $R \approx 0.07$ and $d_{max} \approx 0.42$.

E. The Data-Driven Kelly Betting Trader

The scenario in this third and final simulation is for a trader is going long using the Kelly Betting described in Section VI. Again, at the close of each day k , a market order at price $S_k = C_k$ with trade size is determined by a Kelly-optimal time-varying linear feedback of the form $u_k = K_k^* V_k$ which is obtained by the maximization of expected logarithmic growth using lookback window of variable size s to obtain the probability mass function required in the optimization. The results of the simulation, shown for illustrative values $s = 40, 78, 90$ are seen in Figure 11.




Fig. 11. Account Value for Data-Driven Kelly Betting Trader

For the performance metrics R and d_{max} given above, we obtained:

- for $s = 40$: $R \approx 1.27$ and $d_{max} \approx 0.39$;
- for $s = 78$: $R \approx 1.85$ and $d_{max} \approx 0.26$;
- for $s = 90$: $R \approx 1.09$ and $d_{max} \approx 0.31$.

F. Additional Comments on the Simulations

Looking at the account value plots and the reported values for the raw return R and maximum percentage draw-down d_{max} , in order to make a performance comparison between two such results, it is important to point out that traders may disagree which of the two results is “preferred.” To make this point, imagine a trader comparing the $s = 40$ plot for the Kelly Betting Trader with the $K = 0.05$ plot for the Limit Order Trader. In making this comparison, someone uninitiated in the world of finance might instantly conclude that “Kelly is preferred” because its raw return of 127% far outweighs the “measly” 6% return for the limit order method. However, such reasoning is flawed because other performance attributes also need to be considered. For example, for the comparison under consideration, the Kelly method has a maximum percentage drawdown of 39%, while the limit order method has a maximum percentage drawdown of only 2%. A conservative trader may be unwilling to get drawn down by such a large percentage because it is unknown if the stock price will recover and eventually return to a high level of profitability. Suffice it to say, coming up with a performance rating for $V(k)$ plots involves many considerations beyond the scope of this paper. Here, for illustrative purposes we only used the 2-tuple (R, d_{max}) as performance metrics but there may be many other factors which a practitioner would taken into account in a more general multi-attribute setting; e.g., see [49].

IX. CONCLUSION

The main objective of this tutorial was to facilitate entry into the stock trading research area for control scientists without a substantial background in finance and economics. The so-called “jump start,” included in the title of this paper, was accomplished in a number of ways: We covered just a few topics whose central ideas are easily understood because they do not require significant technical development. In addition, we imposed a number of simplifying assumptions so as to avoid obfuscation of key issues with technical or mathematical detail. Finally, as much as possible, we formulated and described the stock trading problems under consideration in control-theoretic terms, the language of the systems and control community.

Of all of the assumptions we imposed, we point to three which are felt to be worthy of future study because their relaxations are not straightforward and are felt to be amenable to research using control community tools: First and foremost among these is our assumption that the market is *frictionless*. Invocation of this assumption for simplification purposes is foundational in many models and theories — the celebrated Black-Scholes pricing model being a prime example; see [47]

for a literature review. In our view, relaxation of this assumption is an important problem, particularly for a practitioner who emphasizes making money rather than theorem proving. However, the reader should be forewarned there are a number of issues which arise when considering such a relaxation: to be considered: latency associated with the time required for orders to travel from the trader to the server at the stock exchange, degree of liquidity as indicated by bid-ask spread and share volume in the limit order book at various price points and the number of shares in the traders desired transaction, to name a few.

Per discussion above, we also mention two more assumptions, *no-leverage* and *cash settlement at closing*, also important to the practitioner, and seemingly appear “easier” to relax. For the case when leverage is allowed, the details of margin account trading need to be brought into play. This includes but not limited dealing with brokerage maintenance requirements and margin calls margin interests rates. Currently, most large brokerages offer margin loan interests rates which are a function of a trader’s account size and are typically in the range of 7% to 15%. Hence, for many investors, the use of a significant margin loan provides a significant obstacle to profitable trading. As far as relaxation of the cash settlement requirement is concerned, the analysis of a trader’s strategies will need to account for the fact that there may be times when the “available cash” in an account is insufficient to carry out a desired trade. For these situations, to carry out such trades a cash account will not suffice and a margin account is needed.

Our approach has been to formulate trading decisions as control-theoretic problems, followed by a validation exercise calibrating the performance of our resulting trading strategies in a process known as backtesting; see Section VIII. This is accomplished by checking the performance of trading strategies on actual historical data, and, once this is deemed satisfactory, the serious trader will often then check the performance of their strategies on live market data, in real time, using a process called *paper trading*. Paper trading simulates a trading strategy using fictitious money but operates in real time rather than using historical data. This allows traders to gain further confidence that their backtests do not inadvertently “cheat” by using future information in trading decisions or relying on assumptions upon stock prices upon which the strategy is based. Many brokers offer paper trading services and there are a number of third party applications doing the same. This process of thorough backtesting and timely paper trading offers the control researcher an opportunity to integrate theory and practice in the pursuit of stock trading algorithms with real impact.

REFERENCES

- [1] D. G. Luenberger, *Investment Science*, Oxford University Press, 1998.
- [2] J. C. Hull, *Options, Futures, and Other Derivatives*, Prentice Hall, 2012.
- [3] S. Achelis, *Technical Analysis From A to Z*, McGraw-Hill, 2000.
- [4] C. D. Kirkpatrick and J. R. Dahlquist, *Technical Analysis: The Complete Resource for Financial Market Technicians*, Financial Times Press, 2007.
- [5] K. L. Grant, *Quantitative Financial Analytics: The Path to Investment Profits*, Wiley, 2007.
- [6] R. L. Kosowski and S. N. Neftci, *Principles of Financial Engineering*, Academic Press, 2015.
- [7] P. Wilmott, S. Howison and J. Dewynne, *The Mathematics of Financial Derivatives: A Student Introduction*, Cambridge University Press, 1995.
- [8] B. R. Barmish and J. A. Primbs, “On a New Paradigm for Stock Trading via a Model-Free Feedback Controller,” *IEEE Transactions on Automatic Control*, vol. 61, pp. 662–676, 2015.
- [9] R. C. Merton, *Continuous-Time Finance*, Blackwell, 1990.
- [10] B. R. Barmish and J. A. Primbs, “Stock Trading via Feedback Control Methods,” in *Encyclopedia of Systems and Control*, J. Baillieul and T. Samad, pp. 2196–2204, Springer, International Publishing, 2021.
- [11] B. R. Barmish, “On Performance Limits of Feedback Control-Based Stock Trading Strategies,” *Proceedings of the American Control Conference*, pp. 3874–3879, San Francisco, 2011.
- [12] A. Deshpande and B. R. Barmish, “A Generalization of the Robust Positive Expectation Theorem for Stock Trading via Feedback Control,” *Proceedings of the European Control Conference*, pp. 514–520, 2018.
- [13] H. M. Markowitz, “Portfolio Selection,” *Journal of Finance*, vol. 3, pp. 77–91, 1952.
- [14] S. Formentin, F. Previdi, G. Maroni, and C. Cantaro, “Stock Trading via Feedback Control: An Extremum Seeking Approach,” *Proceedings of Mediterranean Conference on Control and Automation*, pp. 1–9, Zadar, Croatia, 2018.
- [15] G. Maroni, S. Formentin, and F. Previdi, “A Robust Design Strategy for Stock Trading via Feedback Control,” *Proceedings of the European Control Conference*, pp. 447–452, Naples, Italy, 2018.
- [16] F. Abbracciavento, S. Formentin, and S. M. Savaresi, “Data-driven Stock Trading in Financial Markets: An Adaptive Control Approach,” *International Journal of Control*, vol. 95, pp. 1032–1041, 2022.
- [17] C. H. Hsieh, “On Solving Robust Log-Optimal Portfolio: A Supporting Hyperplane Approximation Approach,” *European Journal of Operational Research*, vol. 313, pp. 1129–1139, 2024.
- [18] C. H. Hsieh, “On Asymptotic Log-Optimal Portfolio Optimization,” *Automatica*, vol. 151, pp. 110901:1–110901:11, 2023.
- [19] G. Cornuejols, J. Pena, R. Tütüncü, *Optimization Methods in Finance*, Cambridge University Press, 2018.
- [20] F. J. Fabozzi, P. N. Kolm, D. Pachamanova, S. M. Focardi, *Robust Portfolio Optimization and Management*, Wiley, 2007.
- [21] A. Shapiro, D. Dentcheva and A. Ruszcynski, *Lectures on Stochastic Programming: Modeling and Theory*, SIAM, 2014.
- [22] J. L. Kelly, “A New Interpretation of Information Rate,” *Bell System Technical Journal*, vol. 35, pp. 917–926, 1956.
- [23] E. O. Thorp, “The Kelly Criterion in Blackjack Sports Betting, and the Stock Market,” in *Handbook of Assets*, pp. 385–428, 2008.
- [24] F. Black and R. Litterman, “Global Portfolio Optimization,” *Financial Analysts Journal*, vol. 48, no. 5, pp. 28–43, 1992.
- [25] T. M. Cover and T. J. Thomas, *Elements of Information Theory*, Wiley, 2006.
- [26] L. C. MacLean, E. O. Thorp and W. T. Ziemba, *The Kelly Capital Growth Investment Criterion: Theory and Practice*, World Scientific, 2011.
- [27] M.-E. Wu, H.-H. Tsai, W.-H. Chung, C.-M. Chen, “Analysis of Kelly Betting on Finite Repeated Games,” *Applied Mathematics and Computation*, vol. 373, pp. 125028, 2020.
- [28] A. V. Proskurnikov and B. R. Barmish, “On the Benefit of Nonlinear Control for Robust Logarithmic Growth: Coin Flipping Games as a Demonstration Case,” *IEEE Control System Letters*, vol. 7, pp. 2275–2280, 2023.
- [29] A. Despons, L. Peliti and D. Lacoste, “Adaptive Strategies in Kelly’s Races Model,” *Journal of Statistical Mechanics: Theory and Experiment*, pp. 1-17, 2022.
- [30] M. D. Gould, M. A. Porter, S. Williams, D. J. Fenn and S. D. Howison, “Limit Order Books,” *Quantitative Finance*, vol. 13, pp. 1709–1742, 2013.
- [31] R. Huang and T. Polak, “LOBSTER: Limit Order Book Reconstruction System,” Available at SSRN 1977207, 2011.
- [32] F. Abergel, A. Chakraborti, A. Jedidi and I. Toke and M. Anane, *Limit Order Books*, Cambridge University Press, 2016.
- [33] C. A. Parlour and D. J. Seppi, “Limit Order Markets: A Survey,” *Handbook of Financial Intermediation and Banking*, vol. 5, pp. 63–95, 2008.

- [34] I. Zaznov, J. Kunkel, A. Dufour and A. Badii, "Predicting Stock Price Changes Based on the Limit Order Book: A Survey," *Mathematics*, vol. 10, no. 8, pp. 1234, 2022.
- [35] P. Jorion, "International Portfolio Diversification with Estimation Risk," *Journal of Business*, pp. 259–278, 1985.
- [36] V. K. Chopra and W. T. Ziemba, "The Effect of Errors in Means, Variances, and Covariances on Optimal Portfolio Choice," in *Handbook of the Fundamentals of Financial Decision Making: Part I*, pp. 365–373, World Scientific, 2013.
- [37] S. P. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004.
- [38] S. Satchell and A. Scowcroft, "A Demystification of The Black–Litterman Model: Managing Quantitative and Traditional Portfolio Construction," *Journal of Asset Management*, vol. 1, pp. 138–150, 2000.
- [39] C. H. Hsieh, "Necessary and Sufficient Conditions for Frequency-Based Kelly Optimal Portfolio," *IEEE Control Systems Letters*, vol. 5, no. 1, pp. 349–354, 2020.
- [40] S. Diamond, S. Boyd, "CVXPY: A Python-Embedded Modeling Language for Convex Optimization," *Journal of Machine Learning Research*, vol. 17, no. 83, pp. 1–5, 2016.
- [41] B. R. Barmish, J. A. Primbs, S. Malekpour, and S. Warnick, "On the Basics for Simulation of Feedback-Based Stock Trading Strategies: An Invited Tutorial Session," *Proceedings of the IEEE Conference on Decision and Control*, pp. 7181–7186, Florence, Italy, 2013.
- [42] B. R. Barmish, S. Condie, S. Warnick, J. Primbs, and D. Materrasi, "On NASDAQ Order Book Dynamics: New Problems for the Control Field," *Proceedings of the American Control Conference*, pp. 5671–5672, Boston, 2016.
- [43] B. R. Barmish, J. A. Primbs, and S. Warnick, "On Feedforward Stock Trading Control Using a New Transaction Level Price Trend Model," *IEEE Transactions on Automatic Control*, vol. 67, no. 2, pp. 902–909, 2022.
- [44] C. L. Li and C. H. Hsieh, "On Adaptive Portfolio Management with Dynamic Black-Litterman Approach," *arXiv preprint arXiv:2307.03391*, 2023
- [45] R. G. Cestari, F. Barchi, R. Busetto, D. Marazzina and S. Formentin, "Hawkes-Based Cryptocurrency Forecasting via Limit Order Book Data," *arXiv preprint arXiv:2312.16190*, 2023.
- [46] G. A. Hawkes, "Hawkes Processes and Their Applications to Finance: a Review," *Quantitative Finance*, vol. 18, no. 2, pp. 193–198, 2018.
- [47] D. Adler, "The New Field of Liquidity and Financial Frictions," *Research Foundation Literature Reviews*, vol. 7, no. 2, pp. 1–37, 2012.
- [48] W. F. Sharpe, "Mutual Fund Performance," *Journal of Business*, vol. 39, no. 1, pp. 119–138, 1966.
- [49] E. Barucci and C. Fontana, *Financial Markets Theory Equilibrium, Efficiency and Information*, Springer Finance Textbooks, 2017.
- [50] B. A. Francis, W. M. Wonham. The Internal Model Principle of Control Theory, *Automatica*, vol. 12, no. 5, pp. 457-465, 1976.
- [51] P. Samuelson, "Proof that Properly Anticipated Prices Fluctuate Randomly," *Industrial Management Review*, vol. 41, 1965.
- [52] E. Fama, "Efficient Capital Markets: A Review of Theory and Empirical Work," *The Journal of Finance*, vol. 25, pp. 383-417, 1970.
- [53] E. Fama, "Random Walks in Stock Market Prices," *Financial Analysts Journal*, vol. 51, pp. 75-80, 1995.
- [54] A. Lo and A. C. MacKinlay, "Stock Market Prices Do Not Follow Random Walks : Evidence From a Simple Specification Test," *Review of Financial Studies*, vol. 1, pp. 41-66, 1988.
- [55] A. Lo and A. C. MacKinlay, *A Non-Random Walk Down Wall Street*, Princeton University Press, 2011.