

ON THE TOPIC OF PORTFOLIO OPTIMIZATION

PETER J. MERCURIO

A DISSERTATION SUBMITTED TO
THE FACULTY OF GRADUATE STUDIES
IN PARTIAL FULFILMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

GRADUATE PROGRAM IN MATHEMATICS AND STATISTICS
YORK UNIVERSITY
TORONTO, ONTARIO

August 2020

©Peter J. Mercurio 2020

Abstract

This dissertation explores the use of information entropy as a risk measure for the purpose of investment portfolio optimization and selection.

First, we present an improved method of applying entropy as a risk in portfolio optimization. A new family of portfolio optimization problems called the return-entropy portfolio optimization (REPO) is introduced that simplifies the computation of portfolio entropy using a combinatorial approach. REPO addresses five main practical concerns with traditional mean-variance portfolio optimization (MVPO), by using a mean-entropy objective function instead of the mean-variance objective function used in MVPO. REPO also simplifies the portfolio entropy calculation by utilizing combinatorial generating functions in the optimization objective function.

Next, we extend the REPO approach to the optimization problem for assets with discrete distributed returns, such as those from a Bernoulli distribution like binary options. Under a discrete probability distribution, portfolios of binary options can be viewed as repeated short-term investments with an optimal buy/sell strategy or

general betting strategy. Portfolio selection under this setting can be formulated as a new optimization problem called discrete entropic portfolio optimization (DEPO). DEPO creates optimal portfolios for discrete return assets based on expected growth rate and relative entropy. We show how a portfolio of binary options provides an ideal general setting for this kind of portfolio selection.

Finally, we introduce a further generalized portfolio selection method called generalized entropic portfolio optimization (GEPO). GEPO extends DEPO to include intervals of continuous returns, with direct application to a wide range of options strategies. This lays the groundwork for an adaptable optimization framework that can accommodate a wealth of option portfolios. These options strategies exhibit mixed returns: a combination of discrete and continuous returns with performance best measured by portfolio growth rate, making entropic portfolio optimization an ideal method for option portfolio selection. GEPO provides the mathematical tools to select efficient option portfolios based on their growth rate and relative entropy.

Keywords: return entropy; portfolio optimization; portfolio selection; diversification; relative entropy; Kullback-Leibler divergence; Kelly criterion; binary options; sports betting; option portfolios

Acknowledgements

First and foremost, I would like to deeply thank my supervisor Dr. Yuehua Wu for her unwavering support, guidance, and for the honourable opportunity to study under her supervision. Thanks to her I have surpassed my greatest academic dreams.

Additional special thanks goes to Dr. Hong Xie, my industry professional supervisor from Manulife Financial. It was Dr. Xie's brilliant idea to explore the use of entropy in portfolio optimization. I would also like to extend my special thanks to Dr. Xin Gao for sitting on my doctoral committee, and always supporting my studies and professional career.

I am especially grateful to the management and executive team from CIBC for all their emotional and financial support throughout my entire graduate studies.

Finally, I thank my parents, family and friends, and my girlfriend for all their never-ending support, encouragement, patience, guidance, and for always pushing me to follow my heart and try my best.

Table of Contents

Abstract	ii
Acknowledgements	iv
Table of Contents	v
List of Tables	ix
List of Figures	xii
1 Introduction	1
2 An Entropy-Based Approach to Portfolio Optimization	6
2.1 Introduction	6
2.2 Modern Portfolio Theory	9
2.2.1 Markowitz Mean-Variance Portfolio Optimization (MVPO) . .	9
2.2.2 Practical Difficulties with MVPO	10

2.2.3	Literature Review	12
2.3	Entropy as a Risk Measure	16
2.3.1	Shannon Entropy (Information Theory)	16
2.3.2	Portfolio Optimization Based on Entropy	17
2.3.3	Probability Generating Functions	18
2.3.4	Portfolio Entropy Objective Function	22
2.3.5	Return-Entropy Portfolio Optimization (REPO)	22
2.4	A Portfolio Selection Example with REPO	24
2.4.1	Data	24
2.4.2	Efficient Frontier and Portfolio Selection	26
2.4.3	Comparison to MVPO	27
2.4.4	Solutions to the Five Main Issues with MVPO	32

3 Portfolio Optimization for Binary Options Based on Relative Entropy **36**

3.1	Introduction	36
3.1.1	Literature Review	39
3.2	Maximum Exponential Growth Rate	42
3.2.1	The Kelly Criterion	42
3.2.2	Extension of the Kelly Criterion to Multiple Wagers	44

3.3	Shannon Entropy of Discrete Returns	45
3.3.1	Investments Versus Wagers	45
3.3.2	Joint Entropy of a Portfolio of Discrete Return Assets	46
3.4	Minimum Relative Entropy	48
3.4.1	Kullback-Leibler Divergence	48
3.4.2	Relative Entropy as a Convex Risk Measure	50
3.4.3	Minimum Risk Option Portfolios with Relative Entropy	54
3.4.4	Discrete Entropic Portfolio Optimization (DEPO)	56
3.4.5	Risk-Adjusted Performance	59
3.5	Portfolio Selection Examples with DEPO	60
3.5.1	FOREX Binary Option Portfolio Example	60
3.5.2	NFL Sportsbook Example	71

4 Option Portfolio Selection with Generalized Entropic Portfolio Optimization 82

4.1	Introduction	82
4.1.1	Literature Review	84
4.2	Maximum Exponential Growth Rate	87
4.2.1	The Kelly Criterion for Multiple Wagers	87
4.2.2	Extension of the Kelly Criterion to Option Strategies	87

4.3	Minimum Relative Entropy	102
4.3.1	Shannon Entropy	102
4.3.2	Kullback-Leibler Divergence	103
4.4	Option Portfolio Selection Based on Growth Rate and Relative Entropy	105
4.4.1	Generalized Entropic Portfolio Optimization (GEPO)	105
4.4.2	Risk-Adjusted Performance	108
4.5	An Option Portfolio Selection Example with GEPO	109
4.5.1	Data	109
4.5.2	Efficient Frontier and Portfolio Selection	114
4.5.3	Comparison to the Kelly Criterion Over Time	119
5	Conclusions and Future Work	121
5.1	Conclusions	121
5.2	Future Work	124
	Bibliography	125
A	Appendix	134
A.1	Proposition 1	134
A.2	Proposition 2	135

List of Tables

2.1	The ten randomly selected securities from S&P/TSX 60 and the sample means, variances, and entropies of their mean weekly returns over the ten-year period.	25
2.2	Minimum objective and optimal solutions for mean-variance portfolio optimization (MVPO) and return-entropy portfolio optimization (REPO) methods.	28
2.3	Optimal solutions for MVPO And REPO methods with expected returns of 0.37 bps.	29
2.4	Comparison of REPO vs. MVPO portfolios over 20 weeks in 2011: number of portfolios that achieved greater returns.	31
2.5	Optimal solutions via REPO by various risk tolerances.	34
3.1	Mean, in-the-money rate and estimated relative entropy (in bits) of FOREX binary options from January 2019 to January 2020.	62

3.2	Select FOREX binary options for February 2, 2020 11:00PM, with their respective contract strike price, market consensus edge, and estimated probability in-the-money.	64
3.3	The Kelly criterion portfolio of options for contracts expiring February 2, 2020 11:00PM.	67
3.4	DEPO portfolio of options for contracts expiring February 2, 2020 11:00PM.	68
3.5	Mean, cover rate, estimated relative entropy (bits) of NFL teams 2011-19.	73
3.6	Scheduled NFL games for Sunday September 8, 2019, with their respective Las Vegas point spreads and market consensus estimated probabilities of covering.	75
3.7	The Kelly criterion portfolio of wagers with percent allocation for NFL 2019-20 season week 1.	78
3.8	DEPO portfolio of wagers with percent allocation for NFL 2019-20 season week 1.	79
4.1	Mean outcome, average state probabilities and estimated relative entropy (in trits) of select equity credit spread options from July 2012 to January 2018.	111

4.2	Selected equity credit spreads on January 12, 2018, with their respective spread intervals, deltas and state projections.	113
4.3	The Kelly criterion portfolio of options with percent allocation for expiration week 2, January 12, 2018.	116
4.4	GEPO portfolio of options with percent allocation for expiration week 2, January 12, 2018.	118

List of Figures

1.1	Sample asset return payoff function for class of assets used in return-entropy portfolio optimization (REPO).	3
1.2	Sample asset return payoff function for class of assets used in discrete entropic portfolio optimization (DEPO).	3
1.3	Sample asset return payoff function for class of assets used in generalized entropic portfolio optimization (GEPO).	4
2.1	Mean-entropy efficient frontier.	27
2.2	Risk-return efficient frontier: entropy vs. variance.	28
2.3	Optimal portfolio actual returns: return-entropy portfolio optimization (REPO) vs. mean-variance portfolio optimization (MVPO). . . .	30
2.4	Risk-diversification efficient frontier.	33
3.1	Growth rate by relative entropy efficient frontier for FOREX binary option portfolio.	66

3.2	Comparison of FOREX binary option strategies over February 2020. .	70
3.3	Comparison of FOREX binary option strategies over February and March 2020.	71
3.4	Growth rate by relative entropy efficient frontier for NFL sportsbook portfolio.	77
3.5	Comparison of sportsbook strategies over 17 weeks of 2019-20 season NFL games.	81
4.1	Sample expected payoff function for a covered call strategy.	88
4.2	Sample expected payoff function for a married put strategy.	90
4.3	Sample expected payoff function for a put credit spread strategy. . . .	92
4.4	Sample expected payoff function for a call credit spread strategy. . . .	93
4.5	Sample expected payoff function for a straddle strategy.	95
4.6	Sample expected payoff function for a long strangle strategy.	97
4.7	Sample expected payoff function for a butterfly spread strategy. . . .	99
4.8	Sample expected payoff function for an iron condor strategy.	101
4.9	Growth rate by relative entropy efficient frontier.	115
4.10	Comparison of Credit Spread Strategies from January 2018 to January 2020.	120

1 Introduction

Portfolio optimization is the process of selecting an optimal portfolio of assets and allocation weights according to a set of conditions or constraints. Modern portfolio theory (MPT) introduced a novel mathematical framework for data-driven portfolio selection with the objective of maximizing expected return for a chosen level of risk, with risk defined as the variance of expected returns.

This dissertation explores the use of information entropy as a proxy for risk in portfolio optimization. Information entropy holds several advantages over variance in this context: (i) the use of entropy as a risk lends itself very well to highly diversified portfolios, more-so than the use of variance, (ii) the dependence structure of asset returns is well preserved when selecting assets based on entropy, (iii) adjusting model inputs to entropy optimization yields expected and intuitive changes in optimal solutions, (iv) the non-parametric approach of entropy handles non-normality and asymmetry with ease, and (v) entropy optimization has a wide variety of interesting and exciting applications.

The use of entropy enables portfolio optimization to be sufficiently adapted to accommodate various types of investment portfolios, including equity portfolios, binary or digital option portfolios, sportsbook portfolios, and most effectively, derivative or option portfolios. This dissertation introduces three novel portfolio optimization methods to facilitate these goals, which are summarized as follows:

(1) Return-entropy portfolio optimization (REPO): for the selection of portfolios comprising continuous return assets, such as equities, indices, mutual funds, or other non-derivative financial products,

(2) Discrete entropic portfolio optimization (DEPO): for the selection of portfolios comprising strictly discrete return assets, such as binary options, digital options, fixed-return options (FRO), sports bets, or other betting wagers, and

(3) Generalized entropic portfolio optimization (GEPO): for the selection of portfolios comprising assets with mixed-returns (returns that can exhibit both continuous and discrete attributes), like option strategies such as covered puts, married calls, credit spreads, straddles, strangles, butterfly spreads, iron condors, and more.

The applications of these new methods can be illustrated by the following investor return payoff function diagrams:




Figure 1.1: Sample asset return payoff function for class of assets used in return-entropy portfolio optimization (REPO).




Figure 1.2: Sample asset return payoff function for class of assets used in discrete entropic portfolio optimization (DEPO).




Figure 1.3: Sample asset return payoff function for class of assets used in generalized entropic portfolio optimization (GEPO).

Figure 1.1 represents the payoff function for an investor with a regular long position in a stock. For every dollar increase or decrease in the underlying stock price, a proportional increase or decrease is generated in investor return. This payoff function is continuous, since the underlying stock price and thus investor returns fall on the (non-negative) real line.

The second Figure 1.2 illustrates the payoff function for a binary option. This type of option generates a fixed positive return W if the underlying asset price lands above a certain threshold, in this case a strike price of \$60, and generates a fixed negative return L otherwise. This payoff function is strictly discrete, since the only possible return states are $+W$ and $-L$.

The latter Figure 1.3 is an example of a payoff function for a more complex option strategy called a bull put credit spread. This strategy involves buying and selling put options (an option to sell assets at an agreed upon price and date) in such a manner to limit one’s potential risk. The resulting payoff function covers a “spread” outside of which the option behaves just like a binary option—it generates a fixed positive return W if the underlying asset price lands above the upper bound, and generates a negative fixed return L if it lands below the lower bound. In between, if the price lands between the boundaries of the spread, the investor receives a partial return proportional to the distance between the midpoint and the underlying price, as seen in the payoff function. These returns can be both discrete and continuous: $+W$, $-L$, and anything in between. These kind of option strategies are the main motivation behind this research.

This dissertation is organized as follows. Chapter 2 introduces return-entropy portfolio optimization (REPO) with an application to equity portfolio selection. Chapter 3 presents discrete entropic portfolio optimization (DEPO) with applications to FOREX binary options and NFL sportbetting portfolio selection. Chapter 4 presents generalized entropic portfolio optimization (GEPO) with applications to option portfolio selection. Lastly, conclusions and future work on this topic are summarized in Chapter 5.

2 An Entropy-Based Approach to Portfolio Optimization

2.1 Introduction

Markowitz [65] introduced the world's first fundamentally sound quantitative approach to portfolio selection in 1952. He proposed an algorithm that finds the optimal capital allocation across a set of assets based on user-controlled risk parameters. Investors were suddenly given the mathematical tools needed to construct data-driven optimal portfolios according to their preferred risk tolerances. Based on the volatility of random returns, Markowitz's mean-variance portfolio optimization (MVPO) measures the risk of an asset by its second central moment, the variance or squared deviation of returns from the mean. In the multiple-asset case, the risk of a portfolio is measured by the covariance of returns of its comprised assets weighted by their respective capital allocations. The result is a personally-tailored investment

portfolio with the optimal balance between risk and return. Further work on the topic was contributed by Tobin (1958) [98]. Over the years that followed, MVPO quickly became the de facto standard for portfolio selection and capital asset pricing among institutional equity firms, mutual funds, and hedge funds. Its massive influence led to the term *variance* becoming ubiquitous when evaluating risk in the world of finance. Markowitz’s variance-based approach to risk mitigation formed the foundation for modern portfolio theory and investment analysis, and inspired the basis for the capital asset pricing model (CAPM) introduced independently by Sharpe (1964) [89], Lintner (1965) [58; 59], and Mossin (1966) [69].

More recently, topical literature has explored some common difficulties encountered when employing MVPO in the real world. Notably, there are five main issues that complicate the use of MVPO in practice: (i) optimal solutions assigning large allocation weights to high risk assets, (ii) disturbance of the assets’ dependence structure, (iii) drastic variations in optimal solutions when adjusting inputs, (iv) accommodation of non-normal or asymmetric returns, and (v) difficulty in estimating the covariance matrix and expected returns. Researchers have suggested various solutions for addressing these main issues. One popular method is the Black-Litterman asset allocation model (1991,1992) [12; 13] which allows investors to input their own “investor views” without causing unexpected results. A particularly promising ap-

proach was to use entropy as a risk instead of variance, first proposed by McGill (1954) [66] and Garner (1956) [33], and then extended to the portfolio selection problem by Philippatos (1972) [74]. Philippatos' use of joint entropy resulted in a complex computation that proved to be a road-block for practical applications. A main focus of contemporary literature on this topic explored entropy of the portfolio weights as a maximization objective to encourage diversification levels, as found in work by Cheng (2006) [16], Usta (2007) [100] Huang (2007,2008,2012) [38; 39; 40], and Pola (2016) [75]. A Kullback-Leibler view of maximum entropy is demonstrated by Abbas (2017) [1], and the use of cross-entropy minimization was explored by Post (2017) [96]. Further details on the use of entropy for portfolio selection are discussed in Section 2.2. Nevertheless, we believe that entropy-based risk is an ideal approach to addressing the five main difficulties with MVPO. The key is developing a simple yet effective method for calculating the entropy of a portfolio. More detailed discussion is presented in Section 2.3. The chapter is thus organized as follows. Section 2.2 provides a brief review of relevant portfolio optimization methods to date. Section 2.3 introduces the concept of entropy as a risk measure and its favourability as an approach to portfolio optimization, and then details the featured method of this chapter, return-entropy portfolio optimization (REPO). A real-life portfolio selection example using REPO is demonstrated in Section 2.4, and conclusions are discussed

in Chapter 5.

Throughout this chapter, if the size of the distribution support of a discrete random variable is m , we say that the distribution or the random variable has m states.

2.2 Modern Portfolio Theory

2.2.1 Markowitz Mean-Variance Portfolio Optimization (MVPO)

The portfolio selection problem can be stated as such. Given a set of n assets and their respective expected future returns $E(R_1), \dots, E(R_n)$, the goal is to construct the optimal portfolio R_P by allocating weights $w_1, \dots, w_n$ representing the percentages of capital to invest into each asset. The objective function of this optimization problem is designed to minimize the risk and maximize the expected returns of the portfolio. Markowitz (1952) [65] defined risk as the variance of the portfolio returns. Markowitz's MVPO minimizes variance and maximizes expected returns via the following multi-objective function and constraint set,

$$\text{minimize } \text{Var}(R_P) = w_1^2 \sigma_1^2 + \dots + w_n^2 \sigma_n^2 + \sum_i \sum_{j \neq i} w_i w_j \sigma_i \sigma_j \rho_{ij}$$

$$\text{maximize } E(R_P) = w_1 E(R_1) + \dots + w_n E(R_n)$$

$$\text{subject to } w_1 + \dots + w_n = 1, \quad w_i \geq 0 \quad \forall i,$$

where σ_k^2 is the variance of R_k , and $R_P = w_1 R_1 + \cdots + w_n R_n$.

2.2.2 Practical Difficulties with MVPO

There are five main practical difficulties that are often encountered when utilizing MVPO in the real world. These are:

(1) Large weights assigned to high risk assets (sparse solution). In practice, the mean-variance optimization tends to concentrate large-percentage allocations on few assets, often ones with high risk. This is especially common when adjusting the risk parameter to achieve greater returns. This creates a sparse solution with little diversification, which is a consequence opposed to the original intention. This challenge has been studied by various authors who tried to improve the mean-variance portfolio diversification. See Black (1992) [13], Green (1992) [35], Corvalán (2005) [20], and Koumou (2019) [51]. Shannon entropy became a popular method in the sense of diversifying the portfolio weights, and further details on this are found in Section 2.3. Diversification using different entropy measures were explored by Yu (2014) [105]. Lastly, approaches using Rao’s quadratic entropy and diversity measures (Rao, 1982,1985,2004,2010) [77; 78; 79; 80; 81] are discussed in detail by Carmichael (2015) [15].

(2) Disturbing the dependence structure equilibrium. An investor using MVPO

typically calculates a covariance matrix of historical returns for the risk function, but may often wish to input his/her own views (estimates/opinions) about future expected returns. Using investor views instead of historical returns can disturb the dependence structure equilibrium and cause unexpected optimization results, such as (3). See Black (1992) [13] and Babaei (2015) [5].

(3) Drastic variations in optimal solutions when adjusting inputs. An important consequence of (2) is that there are drastic variations in optimal solutions when taking investors' views into consideration. Small changes in the expected return inputs can cause major changes in optimal solutions, which is counterintuitive and unpredictable. See Michaud (1989) [67], Best (1991) [11], Jorion (1992) [46], and Chopra (1993) [17].

(4) Dealing with returns that are non-normal or asymmetric. The Markowitz model relies on symmetry and normality assumptions, and departure from these assumptions can lead to unexpected results. See Jondeau (2005) [44] and Karandikar (2012) [47]. In the real world, asset returns are typically not normally distributed or even symmetric, which makes variance a poor measure for risk. This is not ideal for any investment strategies because an upside volatility is actually welcomed or even desired. Solutions to non-normality and asymmetry in the literature fall into two main categories:

1. Post-modern portfolio theory (PMPT). (i) It only minimizes the downside volatility; (ii) it considers asset distributions as log-normal instead of normal; (iii) it optimizes higher moments than variance (skewness and kurtosis). See Rom (1993) [82] and Sortino (1994) [93].

2. The portfolio entropy minimization method (entropy as a risk). (i) It does not require the normality assumption; (ii) it can accommodate asymmetric distributions; (iii) it is fully non-parametric. See Philippatos (1972) [74], Jiang (2018) [43], and Lassance (2019) [55]—further discussion in Section 2.2.3.

(5) Difficulty in estimating covariance matrix and expected returns. Since portfolio optimization is a forward-looking exercise, historical returns may not be very useful, as past returns are not always indicative of future returns. Forecasts for expected returns are often used instead (i.e., investors' views of expected returns). The covariance matrix can also be difficult to estimate. See Wong (2012) [103] and Sun (2019) [95].

2.2.3 Literature Review

Various authors have managed to address one, two, or even three of these issues, but not all five at once. The following is a brief review of previous research on these topics.

Addressing (1) and (4), Philippatos (1972) [74] aimed to find the optimal portfolio by minimizing the portfolio entropy. Philippatos described three methods by which one can construct mean-entropy diversified portfolios: (i) by calculating the individual and conditional entropies and using them in conjunction with the expected returns; (ii) by computing the individual entropies for each security and their conditional entropies with respect to the level of some acceptable market index (diagonal-index model); and (iii) by computing the security and portfolio entropies directly from the respective variances when it can be assumed that the probability distribution of returns is known. For the first method, it is worth noting that a significant number of conditional entropies must be computed to obtain the joint entropy of the portfolio, since the joint entropy is comprised of the sum of the conditional entropies. For example, calculating the joint entropy of a four-security portfolio with only three states requires one individual entropy plus $3 + 9 + 27 = 39$ conditional entropies. In general, $\sum_{i=0}^{n-1} m^i$ individual and conditional entropy calculations are required, where n is the number of securities and m is the number of probability states, as shown in the appendix. As a result of this complexity, the diagonal-index model was proposed by Philippatos. But this approach provides a poor approximation that does little to account for the true dependence structure of the assets. Lastly, the third method assumes some probability distributions, but is not applicable for unknown distri-

butions or non-parametric approaches. Other measures of entropy, such as Rényi entropy (1960) [86], have been explored by Lassance (2019) [55], but involve similar reliance on a non-parametric estimator of the exponential Rényi entropy function.

Issues (2) and (3) gained the most attention from Black and Litterman at the Goldman Sachs Fixed Income team, who presented methods to help to preserve the dependence structure and stabilize solutions. The Black-Litterman model (Black, 1991,1992) [12; 13] allows users to provide their views that represent their opinions on expected returns and confidence levels. The post-modern portfolio theory was coined by software entrepreneurs Rom and Ferguson (1993) [82] with their work involving the downside risk that targets (4). Additionally, Sortino (1994) [93] introduced the Sortino ratio, which measures the downside-risk-adjusted returns of an asset or portfolio. Cheng (2006) [16], Usta (2007) [100], Huang [38; 39; 40], and Bera (2008) [10] all proposed maximum entropy diversification (MED) methods that maximize the entropy of the portfolio weight vector from MVPO, addressing (1). Pola (2016) [75] argues that portfolio risk and diversification should be managed distinctly, and empirically shows that entropy is a useful means to alleviate the lack of diversification of portfolios on the efficient frontier using a maximum entropy method like MED. A proposed solution called the risk parity given by Asness (2012) [4] was to extend the risk-free tangent line by borrowing capital and

leveraging the portfolio. This is a simple and intuitive solution to (1) but does not address the other issues, and such leverage may not be feasible or available to all investors. Usta (2011) [101] extended the maximum entropy (diversification) approach with the mean-variance-skewness-entropy (MVSE) optimization by adding a multi-objective function to maximize portfolio skewness, which also targets (4). Fono (2011) [32] attempted to obtain an optimal portfolio by introducing semi-kurtosis into the objective function to minimize the low-side tail risk via the mean-semivariance-skewness-semikurtosis (MSSS) optimization. Urbanowicz (2014) [99] took the same approach to the diversification as Cheng and Huang but used Tsallis entropy of portfolio weights instead of Shannon entropy. These methods do not use the joint entropy, as they deal only with the one-dimensional entropy of the portfolio weights vector. A maximum entropy method was proposed by Xu (2014) [104] that aimed to maximize the worst-case portfolio returns. In recent years, entropy was used to evaluate tail risks by Geman (2015) [34]. Zhou (2015) [109] assumed independence between assets to approximate the portfolio entropy by the sum of the individual entropies. Zhou then accommodated the missing dependence structure by also minimizing the portfolio variance, via a multi-objective function. Most recently, Rotela (2017) [84] used the entropic data envelopment analysis (DEA) to improve the diversification of optimized portfolios. Zhou (2017) [108] evaluated six entropy-based

risk measures and declared the mean fuzzy entropy optimization as the best performing method. Dai (2018) [22] used the concept of quadratic entropy to minimize the risk of a portfolio via the mean-quadratic entropy (MQE) optimization by using a multi-objective function that maximizes the entropy of the portfolio weights and minimizes the (approximate) quadratic entropy of the portfolio.

2.3 Entropy as a Risk Measure

2.3.1 Shannon Entropy (Information Theory)

In 1948, Shannon [87; 88] introduced the concept of information entropy. Applied to a probability vector, the information entropy represents the amount of randomness or uncertainty inherent to that probability distribution. It is a measure of how many “choices” are involved in the selection of an event or of how certain we are to its outcome.

For a discrete random variable X with probability mass function $P(\cdot)$ that can take on possible values $x_1, \dots, x_n$, the Shannon entropy H is the average amount of information produced by X , defined as

$$H(X) = E(-\log P(X)) = - \sum_{i=1}^n P(x_i) \log P(x_i).$$

For two discrete random variables X and Y respectively having n and m states,

the joint entropy of X and Y is given by

$$H(X, Y) = - \sum_{i=1}^n \sum_{j=1}^m P_{XY}(x_i, y_j) \log P_{XY}(x_i, y_j).$$

Furthermore, the conditional entropy of Y given X , representing the average amount of information in Y given X , is defined by

$$H(Y | X) = - \sum_{i=1}^n \sum_{j=1}^m P_{XY}(x_i, y_j) \log P_{Y|X}(y_j | x_i) = H(X, Y) - H(X).$$

While variance and entropy are both non-negative quantities, it is important to note a main differences between the two measures: variance takes a value in $[0, \infty)$, whereas the entropy of a random variable with m states is bounded by the “maximal entropy”; i.e., the entropy of the uniform distribution with m states, as shown by using Jensen’s inequality (Jensen, 1906) [42],

$$H(X) = E \left[\log \left(\frac{1}{P(X)} \right) \right] \leq \log \left(E \left[\frac{1}{P(X)} \right] \right) = \log(m).$$

2.3.2 Portfolio Optimization Based on Entropy

By using entropy in place of variance for the portfolio selection problem, all five main difficulties with MVPO can be solved, as (1) low risk portfolios selected by entropy-based methods provide greater diversification levels than those selected by variance-based methods; (2) the dependence structure is maintained, since the entropy is not based on the mean; (3) the optimal solution is more stable under the

adjustments to investors' views; (4) entropy is a non-parametric function designed to accommodate non-normality and asymmetry; and (5) no pre-calculations of any covariance matrices are necessary, as the joint entropy dependence structure can be automatically captured in the objective function.

The aim of this chapter is to introduce a single optimization problem that solves all five issues with the MVPO method. Presented here is an approach to the portfolio optimization using a minimum entropy method, called return-entropy portfolio optimization (REPO), which has not yet been proposed elsewhere. As other methods encountered the difficulty in calculating the portfolio entropy due to the complexity of the joint entropy expression, REPO calculates the resulting portfolio entropy directly in the objective function using combinatorial generating functions, eliminating the need for any intermediary joint entropy calculations. It provides the following practical benefits over MVPO: better stability under changes to inputs, robustness against non-normality and asymmetry, and improved portfolio performance, as shown in Section 2.4.

2.3.3 Probability Generating Functions

Consider an asset's return, a random variable R , and its historical observations $\mathbf{r} = (r_1, \dots, r_T)$. The range of historical observations can be divided into m dis-

distinct probability state partitions, $A_1, \dots, A_m$ with endpoints $[a_0, a_1], \dots, [a_{m-1}, a_m]$ respectively, such that each r_j belongs to only one partition; i.e., $a_{k-1} < r_j \leq a_k$, for some integer k in $1, \dots, m$. Without loss of generality we can assume the partitions to be equally sized. The probability of each event $\{R \in A_k\}$ can be estimated empirically over T time periods as

$$\widehat{f}_T(\mathbf{r}; k) = \frac{1}{T} \sum_{j=1}^T I(a_{k-1} < r_j \leq a_k) \approx P(R \in A_k).$$

Consider now a portfolio with two assets. Denote their returns by R_1, R_2 , respectively. Let their actual returns over T time units be $\mathbf{r}_1 = (r_{11}, \dots, r_{1T})$ and $\mathbf{r}_2 = (r_{21}, \dots, r_{2T})$, with portfolio weights w_1, w_2 such that $w_1 + w_2 = 1$. The portfolio entropy of $R_P = w_1 R_1 + w_2 R_2$, $H(R_P)$, can be estimated by using the empirical probability frequency tables directly via combinatorial generating functions as follows. Take the empirical probability generating function g ,

$$g(x; w_1 \mathbf{r}_1 + w_2 \mathbf{r}_2) = \frac{1}{T} \sum_{j=1}^T x^{\{k : a_{k-1} < w_1 r_{1j} + w_2 r_{2j} \leq a_k\}},$$

for k such that $\{a_{k-1} < w_1 r_{1j} + w_2 r_{2j} \leq a_k\}$. Notice that to sum over all the powers of x one needs to count all the occurrences of the event $\{R_P \in A_k\}$ enumerated in the coefficients similarly to how a histogram or frequency table counts occurrences. The coefficient of x^k estimates the empirical probability of event $\{R_P \in A_k\}$, $\widehat{f}_T(\mathbf{r}_1, \mathbf{r}_2; k)$. These coefficients can be extracted for each polynomial term by taking the k th-derivative of g at $x = 0$, $g^{(k)}(0)$, divided by $k!$. Now with the estimated probabilities

of events, the empirical entropy can be calculated directly. Each coefficient of x^k represents the estimated empirical probability of event $\{R_P \in A_k\}$, which is given by

$$\frac{g^{(k)}(0)}{k!} = \widehat{f}_T(\mathbf{r}_1, \mathbf{r}_2; k) = \frac{1}{T} \sum_{j=1}^T I(a_{k-1} < w_1 r_{1j} + w_2 r_{2j} \leq a_k) \approx P(R_P \in A_k).$$

These estimated probabilities are then substituted into the formula for Shannon entropy (1948) [87; 88] for m total probability states,

$$H(R_P) = - \sum_{k=1}^m \widehat{f}_T(\mathbf{r}_1, \mathbf{r}_2; k) \log \widehat{f}_T(\mathbf{r}_1, \mathbf{r}_2; k).$$

The selection of intervals here, and thus the choice of m , is at the discretion of the user. It should be noted how this selection affects the outcome of the entropy calculation. In the extreme case, arbitrarily small interval sizes would allocate each observation to its own individual interval, with at most one occurrence in each interval. This results in a case of maximal entropy—the uniform distribution—which renders the exercise useless, since every portfolio would equally exhibit maximal entropy, $H(R_P) = \log T$, for T unique states, one for each time period ($T \leq m$). On the other extreme, arbitrarily large interval sizes would create one giant interval that encompasses every single observation. This results in a case of minimal entropy, $H(R_P) = 0$, with no randomness at all. This case is equally ineffective, as every portfolio would equally exhibit zero entropy. The user should explore reasonably

sized intervals that yield the intended level of risk mitigation across portfolios. For more information on density estimation, please see Silverman (1998) [91], the spacing estimates method by Beirlant (1997) [8], and Learned-Miller (2003) [57] or the kernel density estimation method credited to Rosenblatt (1956) [83] and Parzen (1962) [73].

The above method can be extended to the case that there is a portfolio with n assets $R_1, \dots, R_n$, and their actual returns over T time units, $\mathbf{r}_i = (r_{i1}, \dots, r_{iT})$, $i = 1, \dots, n$, with portfolio weights $w_1, \dots, w_n$ such that $w_1 + \dots + w_n = 1$.

It should also be noted that this empirical estimator of entropy is biased. In fact, it has been shown by Paninski (2003) [72] that there does not exist an unbiased estimator of entropy. Corrections to these estimators can be made, but they may not always be satisfactory, as shown by Miller (1955) [68]. In the case of this probability generating function, the severity of the bias depends on the interval selection and choice of m , with larger intervals leading to a stronger bias towards maximal entropy and smaller intervals leading to a stronger bias towards zero entropy, as described above.

2.3.4 Portfolio Entropy Objective Function

Extracting these coefficients, the exact portfolio entropy is given by the following portfolio entropy objective function,

$$H(R_P) = -\frac{g'(0)}{1!} \log \left(\frac{g'(0)}{1!} \right) - \frac{g''(0)}{2!} \log \left(\frac{g''(0)}{2!} \right) - \dots - \frac{g^{(m)}(0)}{m!} \log \left(\frac{g^{(m)}(0)}{m!} \right),$$

for m unique probability states, A_k . Each term in the objective function represents the respective term in the entropy function. This objective function is then to be minimized in the optimization problem.

The reason that this computation is so much easier than methods suggested by other authors is that the entropy is calculated directly on the end-state of the portfolio, *after* the allocation weights have been assigned. Other authors constructed the portfolio entropy by using convoluted combinations of the individual and conditional entropies—drastically increasing the complexity of the calculation as the number of assets n increases.

2.3.5 Return-Entropy Portfolio Optimization (REPO)

The new return-entropy portfolio optimization (REPO) problem uses a multi-objective function that minimizes entropy and maximizes expected returns, formu-

lated as follows,

$$\begin{aligned}
& \text{minimize} & H(R_P) &= - \sum_{k=1}^m \frac{g^{(k)}(0)}{k!} \log \left(\frac{g^{(k)}(0)}{k!} \right) \\
& \text{maximize} & E(R_P) &= w_1 E(R_1) + \cdots + w_n E(R_n) \\
& \text{subject to} & w_1 + \cdots + w_n &= 1, \quad w_i \geq 0 \quad \forall i,
\end{aligned}$$

for $R_P = w_1 R_1 + \cdots + w_n R_n$ and the k th-derivative at $x = 0$ of the probability generating function,

$$g(x; w_1 \mathbf{r}_1 + \cdots + w_n \mathbf{r}_n) = \frac{1}{T} \sum_{j=1}^T x^{\{k : a_{k-1} < \sum_{i=1}^n w_i r_{ij} \leq a_k\}}.$$

The reader should note that REPO evaluates the portfolio entropy as the individual entropy of allocation-weighted portfolio returns $H(aX + bY)$ (a one-to-one dimensional function), whereas Philippatos [74] technically evaluated the portfolio entropy as the joint entropy $H(aX, bY) = H(aX) + H(bY) - I(aX; bY)$ (an n -to-one dimensional function, for mutual information $I(aX; bY)$). However, the key point to note is that the probability generating function method used in REPO works perfectly fine for both $H(aX, bY)$ and $H(aX + bY)$.

Shown in the appendix, as a direct consequence of the well-known data processing inequality (Cover, 1991 [21], and Beaudry, 2012 [7]), $H(aX, bY)$ is always greater than or equal to $H(aX + bY)$ —which means more uniformity—due to the higher dimensionality of the former. To this end, it is our contention that more uniformity

is worse for entropic portfolio optimization because with high enough dimensionality the distributions can quickly all resemble the uniform distribution (maximum entropy), and then no differentiation between portfolios via entropy can be done. Therefore, we decided to use $H(aX + bY)$ as the portfolio entropy measure for the objective function in REPO.

2.4 A Portfolio Selection Example with REPO

2.4.1 Data

In the example provided here are actual market data for ten randomly selected securities gathered from the S&P/TSX 60 stock market index over the ten-year period from January 1, 2001, to December 31, 2010, totaling 520 data points each. Weekly closing prices are recorded and adjusted for stock splits, and relative weekly returns computed as follows,

$$r_{ij} = \frac{P_{ij}}{P_{i,j-1}} - 1,$$

where r_{ij} is the percent return on security i in period j , and P_{ij} represents the price of security i in period j . For this example, the percent returns are discretized simply by using interval sizes of one basis point, with the minimum and maximum across

all returns used as support boundaries;

$$[a_{min}, a_{min} + 1], \dots, [-2, -1], [-1, 0], [0, 1], [1, 2], \dots, [a_{max} - 1, a_{max}],$$

all in units of basis points. According to this dataset, the minimum return across securities is -44 basis points and the maximum return is 42 basis points, for a total of $m = 86$ possible probability states. The ten randomly selected securities and the sample means, variances, and entropies of their respective mean weekly returns over the ten-year period are presented below in Table 1, in which bps and nats are respective abbreviations of basis points and natural units.

Table 2.1: The ten randomly selected securities from S&P/TSX 60 and the sample means, variances, and entropies of their mean weekly returns over the ten-year period.

Company Name	Ticker Symbol	Mean (bps)	Variance (bps ²)	Entropy (nats)
Loblaw Companies Ltd.	L	0.006391	8.078711	2.381352
First Quantum Minerals Ltd.	FM	1.003592	61.97863	3.277249
Thomson Reuters Corp	TRI	-0.019931	11.65211	2.53417
Alimentation Couche-Tard Inc.	ATD.B	0.495919	17.89425	2.798943
Bank of Nova Scotia	BNS	0.242633	11.32819	2.466258
Teck Resources Ltd.	TECK.B	0.729174	60.7617	3.259236
Canadian Tire Corp Ltd.	CTC.A	0.284006	12.18994	2.60514
Inter Pipeline Ltd.	IPL	0.211462	7.847551	2.339923
Manulife Financial Corp	MFC	0.095557	24.68777	2.746475
Suncor Energy Inc.	SU	0.424803	27.367	2.907254

The individual sample entropies are displayed here for demonstration purposes only—they are not actually used in REPO (as the portfolio entropy is calculated directly from the weighted portfolio data points). Notice how the trend validates the assumption that higher (absolute-value) return implies higher entropy. The sample correlation coefficient between absolute values of the sample means and the sample entropies in this sample is 0.880746. Interestingly, the rankings in terms of the sample variance or the sample entropy are almost identical here, except that Alimentation Couch-Tard Inc. and Manulife Financial Corp have swapped positions. The variance measure pegs Manulife as having significantly higher risk security (by almost 7 bps²), but according to the entropy Manulife has lower risk (by 0.05 nats).

2.4.2 Efficient Frontier and Portfolio Selection

In the portfolio selection problem, the efficient frontier refers to the set of the optimal portfolios that yield the greatest expected return for a defined level of risk, or alternatively the least risk for a defined level of expected return (the dual problem). The efficient frontier illustrates the risk-return trade-off for a given set of optimal portfolios.

Here we show the REPO algorithm applied to the data given in Section 2.4.1. Plotted below in Figure 2.1 is the mean-entropy efficient frontier among all possible

optimal portfolio solutions. It is evident that greater risk—higher entropy—must be taken in order to achieve higher returns.




Figure 2.1: Mean-entropy efficient frontier.

2.4.3 Comparison to MVPO

In comparison we show MVPO applied to the same data given in Section 2.4.1. Below, Figure 2.2 displays a translated plot of the mean-variance efficient frontier superimposed onto the mean-entropy efficient frontier. Notice the differences in the shape of the frontier. The variance curve is strictly convex, whereas the entropy curve is concave in the outer portion of the curve. This key difference is what enables REPO to find advantageous portfolio allocations that MVPO misses when

it comes to balancing risk and reward.




Figure 2.2: Risk-return efficient frontier: entropy vs. variance.

Another interesting observation is the fact that REPO and MVPO achieve their minimum risk portfolios at different optimal solutions and expected returns, as shown below.

Table 2.2: Minimum objective and optimal solutions for mean-variance portfolio optimization (MVPO) and return-entropy portfolio optimization (REPO) methods.

Method	Minimum Objective	Expected Return	Optimal Solution
MVPO	3.3993 bps ²	0.1394 bps	(0.3,0.0,0.2,0.1,0.0,0.0,0.1,0.3,0.0,0.0)
REPO	1.9355 nats	0.1630 bps	(0.2,0.0,0.2,0.1,0.1,0.0,0.1,0.3,0.0,0.0)

Shown below are the differences in future, actual returns when using REPO versus using MVPO. Two portfolios are constructed using an expected return constraint equal to 0.37 bps and historical prices from 2001 to 2010: one minimum entropy portfolio using REPO and one minimum variance portfolio using MVPO. Optimal solutions to each strategy are shown in the following table.

Table 2.3: Optimal solutions for MVPO And REPO methods with expected returns of 0.37 bps.

Method	Expected Return	Optimal Solution
MVPO	0.37 bps	(0.0,0.1,0.0,0.4,0.0,0.4,0.0,0.0,0.1,0.0)
REPO	0.37 bps	(0.0,0.4,0.3,0.0,0.0,0.0,0.0,0.2,0.1,0.0)

Forward-looking actual daily prices are collected from the chosen ten securities starting January 1, 2011, and over the ensuing 13 weeks, actual portfolio performances are tracked and compared. The results are shown in Figure 2.3 below.




Figure 2.3: Optimal portfolio actual returns: return-entropy portfolio optimization (REPO) vs. mean-variance portfolio optimization (MVPO).

REPO outperforms MVPO here, with actual returns of 0.12 bps gain compared to -0.05 bps loss. Admittedly, this is only one case—which begs the question of how the methods would compare over the course of repeated trials. To that end, here we emulate 7094 different portfolios, each with unique expected return constraints ranging from -0.0199 bps to 1.0036 bps. Two sets of portfolios are constructed here in the same fashion as the previous example. One set of portfolios by minimizing the portfolio entropy using REPO, the other set by minimizing the portfolio variance using MVPO. The actual portfolio returns are calculated for each one of the 7094

portfolios and compared over the following 20 weeks. In 2925 of these emulations, both methods produce identical portfolios so that their returns are equivalent, but the remaining 4169 emulations each reveal a winning method. Table 4 below shows the results of the analysis, counting how many times each method wins.

Table 2.4: Comparison of REPO vs. MVPO portfolios over 20 weeks in 2011: number of portfolios that achieved greater returns.

	REPO	MVPO	Total	% REPO>MVPO
After 2 weeks	2377	1792	4169	57%
After 4 weeks	3115	1054	4169	75%
After 8 weeks	2537	1632	4169	61%
After 13 weeks	2345	1824	4169	56%
After 20 weeks	1699	2470	4169	41%

The emulation demonstrates that REPO outperforms MVPO handsomely in the near-to-medium term. After roughly four months, the REPO portfolios begin to trail the MVPO portfolios, but in practice it is recommended that portfolios are balanced more frequently than once every four months. Therefore, REPO performs better than MVPO in the short-term time horizon.

2.4.4 Solutions to the Five Main Issues with MVPO

This section explores how the REPO proposed solution handles the five main issues with MVPO mentioned in Section 2.2.

(1) Large weights assigned to high risk assets. REPO lends itself very well to highly diversified portfolios, more-so than MVPO. Figure 2.4 below shows a plot of the most risk-averse portfolios according to each method and their corresponding diversification levels, measured by the Shannon entropy of the portfolio weight vector, $H(\mathbf{w})$. The optimal minimum risk portfolio from MVPO is highlighted in green with a diversification level of only 1.5048 nats. The optimal minimum risk portfolio from REPO is highlighted in red with a diversification level of 1.6957 nats, almost 13% more diversified than the MVPO portfolio.




Figure 2.4: Risk-diversification efficient frontier.

Additionally, the wider selection of lower risk portfolios from REPO tend to provide greater diversification levels than those of MVPO—as confirmed by the stronger inverse relationship. The correlation between REPO’s portfolio risk and diversification level is -0.5609 , whereas the correlation between MVPO’s portfolio risk and corresponding diversification level is only -0.2772 .

(2) Disturbing the dependence structure equilibrium. The dependence structure is well preserved when employing REPO in practice. REPO calculates the portfolio entropy directly in the objective function, as contributed by the securities and their respective weightings. This ensures the dependence structure equilibrium is

maintained true to history (or assumptions). The dependence will be unaffected by any changes made to inputs. This is because the entropy of a random vector is not dependent on its mean in the same way that variance is. Consider two identical random vectors with equal mean, variance and entropy. Changing the value of a uniquely occurring element in one vector will change its mean and variance, but the entropy will be unchanged.

(3) Drastic variations in optimal solutions when adjusting inputs. Adjusting inputs to REPO yields expected and intuitive variations in optimal solutions. The multi-objective function of REPO can be formulated by combining the two objectives with a risk tolerance tuning parameter α ,

$$\text{minimize } H(R_P) - \alpha E(R_P).$$

The sensitivity to changes in risk tolerance parameter α is low for the REPO method. The optimization was run on the same data using three different risk tolerance values: $\alpha = 1.0$, 1.4 , and 1.7 . The optimal solutions were as follows,

Table 2.5: Optimal solutions via REPO by various risk tolerances.

Risk Tolerance	Portfolio Entropy	Expected Return	Optimal Solution
$\alpha = 1.0$	1.9551 nats	0.2311 bps	(0.1,0.0,0.1,0.1,0.2,0.0,0.1,0.3,0.0,0.1)
$\alpha = 1.4$	2.1317 nats	0.3588 bps	(0.1,0.1,0.0,0.2,0.1,0.0,0.1,0.3,0.0,0.1)
$\alpha = 1.7$	2.1419 nats	0.3660 bps	(0.1,0.1,0.0,0.2,0.1,0.0,0.2,0.2,0.0,0.1)

(4) Dealing with returns that are non-normal or asymmetric. REPO handles non-normality and asymmetry with ease. There is no assumption or requirement for the data to be normally distributed or symmetric, and since there are no distribution assumptions at all, the optimization problem is fully non-parametric.

(5) Difficulty estimating covariance matrix and expected returns. REPO eliminates issues with estimating the covariance matrix and expected returns. It does all the risk calculations directly in the objective function, thus eliminating the need for any pre-calculations of covariance matrices or individual security variances. Referring back to (3), since the optimal solutions do not have drastic variations in adjustments to inputs, the pressure of making a sensible and accurate expected return estimate is reduced.

3 Portfolio Optimization for Binary Options Based on Relative Entropy

3.1 Introduction

In the previous Chapter 2, a new class of portfolio optimization problems was introduced called return-entropy portfolio optimization (REPO). REPO uses Shannon entropy (Shannon, 1948) [87; 88] as the discriminatory risk measure for portfolio selection for assets with continuous returns, as opposed to variance used by Markowitz [65] in mean-variance portfolio optimization (MVPO). REPO holds several advantages over the traditional MVPO method: REPO is robust, non-parametric, and indifferent to non-normality and asymmetry, making it an ideal approach to the traditional portfolio selection problem.

The focus of this chapter is on use of the entropic portfolio optimization to select optimal portfolios of assets with discrete returns, where traditional risk management

methods are not applicable. In particular, we concentrate on portfolios of binary options. Binary options have fixed discrete distributed returns of $+100\%$ or -100% as opposed to traditional continuous returns on the real line. Other instruments that behave similarly to binary options are digital options, fixed-return options (FROs), or even sports bets. These types of discrete return instruments can be assumed to follow a Bernoulli distribution with an expected probability of success and fixed profit and loss amounts. Both kinds of investment portfolios can have risk-return choices. Under a discrete probability distribution, a portfolio of binary options can be viewed as repeated short-term investments. Upon the outcome of each event, the portfolio incurs a profit or loss from the success or failure of that event. While the return for an equity portfolio is measured by the expected return of the portfolio, for a binary option portfolio we are most interested in the expected growth rate of the portfolio, as if the bet were to be placed repeatedly ad infinitum, as shown by Kelly (1956) [48].

For the risk management of binary option portfolios, we look to the concept of relative entropy, also known as Kullback-Leibler divergence (1951,1959) [52; 54]. Relative entropy measures the distance between discrete distributions, and the target distribution for a portfolio of binary wagers is the uniform distribution. A binary option portfolio with only one asset (maximum relative entropy) has the highest risk

since the portfolio value will swing 100% in either direction. Allocating that same capital to an increasingly large number of binary events reduces the portfolio risk since the expected net return will approach zero as n increases (minimum relative entropy). For this reason, the uniform distribution is the minimum risk portfolio for binary assets. Therefore the risk of a binary option portfolio is quantified by its relative entropy with respect to the uniform distribution. DEPO finds the optimal allocation of capital across a series of potential binary investments in order to maximize the expected portfolio growth rate and minimize the portfolio relative entropy.

Future implications of this work could include news approaches to binary option pricing in practice, based on historical risk levels of relative entropy. Additionally, the joint dependence structure captured by relative entropy could shine new light on the co-movements and relative behaviours of binary and fixed-return options. As the size of the derivatives market grows, new and alternative risk measures like relative entropy will be increasingly sought after for managing the risk of option portfolios. Beyond the derivatives market, relative entropy can have many applications for any type of investment or structured product with discrete, fixed returns, and creates an opportunity to bring novel risk mitigation tools to a variety of industries.

Most literature on the topic of measuring such risk has focused primarily on the

management of bet sizes, and this is usually driven solely by the expected probability of success or failure. The industry standard method for this type of capital allocation is the Kelly criterion (Kelly, 1956) [48]. This theory is reviewed in Section 3.2. Alternative methods for evaluating the risk of discrete portfolios are discussed in Section 3.3. The rest of this chapter is arranged as follows. Section 3.4 demonstrates how relative entropy is the ideal convex risk measure for making quantitative portfolio allocation decisions for gambling wagers, and introduces a new family of entropic portfolio optimization problems called discrete entropic portfolio optimization (DEPO). An example of DEPO applied to foreign exchange binary options is shown in Section 3.5 with results compared to the leading Kelly criterion methods. To illustrate how DEPO is applied to a portfolio of gambling wagers, Section 3.5 further shows a sportsbook portfolio for the NFL 2019-20 season using DEPO for portfolio selection. Chapter 5 discusses the main conclusions derived from this work.

3.1.1 Literature Review

Research on the topic of this type of portfolio optimization can be classified into two separate but related categories: gambling portfolios and investment portfolios. Some authors tailored their papers into one topic or the other, while other authors suggested their strategies were equally applicable in both of them. Most of the

research to date has focused solely on mathematics of the reward, i.e., maximizing the wealth, but contributed little work to the evaluation of risk. Kelly (1956) [48] discovered that a gambler's exponential growth rate of their capital is maximized at the rate of transmission of information over that channel (the Shannon entropy using base two logarithms), and provided the Kelly criterion defining the optimal bet size to achieve such maximal growth. Further work on the Kelly criterion has been carried out and today it is widely used in investment theory as a standard bet size methodology by gamblers and investors alike, including even Warren Buffet (Benello, 2016) [9].

Applications to the investment securities became quite popular around in the 90s, with early work from Rotundo and Thorp (1992) [85] that applied the Kelly criterion to the U.S. stock market. Browne (1996) [14] derived an optimal gambling and investment policy for general stochastic processes using a continuous-time analog involving Brownian motion. After applying to blackjack and other gambling games, MacLean and Thorp (2010) [63] extended the Kelly criterion and its main variants such as fractional Kelly to applications in the securities market. Das (2016) [23] linked the Kelly criterion to portfolio optimization in the review of Browne (1996) [14], and Lavinio (2000) [56] applied it to a day-trading portfolio by using the d -ratio, or gain-loss ratio. O'Shaughnessy (2012) [71] suggested using correlated events to gain an

edge over bookmakers by combining “for” and “against” bets in win-draw-loss markets. Taking uncertainty into account, Baker (2013) [6] shrank the Kelly bet sizes to compensate for the prediction uncertainty (a modified Kelly approach), which showed an improvement over the “raw” Kelly criterion. The Kelly criterion was analyzed further and demonstrated to be incredibly effective over time by MacLean (2013) [64]. Sinclair (2014) [92] devised a confidence interval for the Kelly criterion by calculating variance of the estimated Kelly criterion ratios. Applied to securities, Davari-Ardakani (2016) [24] developed a multistage optimization method that utilizes options to dynamically mitigate the market risk of an investment portfolio. Faias (2017) [29] optimized European option portfolios by proposing a myopic objective function to overcome limitations due to non-normality and small sample sizes encountered by traditional portfolio selection methods. Chu (2018) [18] recently introduced another fractional Kelly method based on the uncertainty of success probabilities by exploring various loss functions. Most recently, Hubacek (2019) [41] exploited sport-betting markets using a betting strategy that maximizes model prediction accuracy and minimizes model correlation with published bookmaker predictions. In most very recent research in the form of a working paper, Vecer (2020) [102] explores much similarly the use of Kullback-Leibler divergence as the optimal utility for the likelihood ratio of the densities corresponding to market takers and market mak-

ers. In this case it is for the purposes of determining optimal payoff functions and equilibrium for Arrow-Debreu securities (contracts that agree to pay one unit of a currency or commodity if a particular event occurs and zero otherwise, much like the concept of a binary option). Here Kullback-Leibler divergence is measured between the distribution of the market agent and the market equilibrium distribution, but not directly applied as a risk measure for purposes of portfolio optimization.

To date, there has not yet been any suggestion for measuring or managing the risk of these option portfolios in literature. We will tackle it in this chapter. The proposed DEPO mitigates the risk of a binary option portfolio by assessing the relative entropy of the portfolio returns. Additionally, DEPO maximizes the exponential growth rate of the portfolio by extending the Kelly criterion to multiple assets.

3.2 Maximum Exponential Growth Rate

3.2.1 The Kelly Criterion

In probability theory, the Kelly criterion (Kelly, 1956) [48] gives the bet size conditions required for gambling wagers to almost surely achieve the maximum exponential growth rate of wealth (or “bankroll”) based on assumed probability of success if the wager were to be placed repeatedly ad infinitum. For the purposes of this chapter, we will just be concerned with the case that wagers that are paying

rewards equal to the bet size, known as fair wagers, but intended further work on the topic that includes extending this to the generalized case of any size payout and odds. A short summary of the Kelly criterion is provided here, courtesy of Khanna (2016) [50]. Consider a wager with expected probability of success p , and expected probability of failure $(1 - p)$. After N trials, we denote the number of successes by S and number of failures by F (such that $S + F = N$). Let w represent the percentage of the portfolio balance to be wagered on each trial. Then for a starting portfolio balance P_0 and resulting portfolio balance P_N after N trials, we have the following expression,

$$P_N = P_0(1 + w)^S(1 - w)^F.$$

It follows that $e^{N \log(P_N/P_0)^{(1/N)}} = \frac{P_N}{P_0}$, which implies that $\log \left(\frac{P_N}{P_0} \right)^{(1/N)}$ measures the exponential growth rate of wealth per trial. The Kelly criterion achieves the maximum expectation of this function via the growth rate coefficient G , defined as the expectation of the exponential growth rate per trial,

$$G(w) = E \left[\log \left(\frac{P_N}{P_0} \right)^{(1/N)} \right] = E \left[\frac{S}{N} \log(1 + w) + \frac{F}{N} \log(1 - w) \right]$$

As the experiment is repeated N times, S becomes a binomial random variable with parameters (N, p) and mean $E(S) = np$. Analogously, F also becomes a binomial random variable with parameters $(N, 1 - p)$ and mean $E(F) = n(1 - p)$. Therefore

by the additive property of expectations, $G(w)$ can be expressed as

$$G(w) = p \log(1 + w) + (1 - p) \log(1 - w).$$

In order to maximize $G(w)$, we take the first derivative with respect to w and set equal to zero, which yields the optimal bet size $w^* = 2p - 1$. It can be verified that w^* is in fact a maximum by observing that the second derivative is negative at w^* , i.e., $\partial^2/\partial w^2 G(w^*) < 0$, and thus w^* is a local maximum. Additionally, since w^* is the only critical point and $G(0) = 0$ while $G(w) \rightarrow -\infty$ as $w \rightarrow 1^-$, we can confirm that w^* is a global maximum. Therefore, the bet size that maximizes the growth rate of wealth is $w^* = 2p - 1$, as a percentage of portfolio balance.

3.2.2 Extension of the Kelly Criterion to Multiple Wagers

Extensions to the Kelly criterion for multiple wagers have not really been extensively explored, beyond a brief expression for two independent wagers as shown by MacLean and Thorp (2010) [63]. As can be seen in Section 3.4, the joint entropy of multiple wagers is only logically defined for equally-sized wagers, so for the scope of this chapter we can restrict analysis to portfolios that contain equally-sized bet allocations. With this constraint as a foundation, we can now generalize the Kelly criterion for two independent wagers and subsequently extend it to n wagers. Consider a collection of two independent events Y_1 and Y_2 , with probabilities of success p_1

and p_2 respectively. Let w represent the total percentage of bankroll to be wagered.

For a portfolio that allocates the wager equally across amongst events, the growth rate coefficient G would be

$$\begin{aligned} G &= \frac{1}{2} [p_1 \log(1 + w) + (1 - p_1) \log(1 - w)] + \frac{1}{2} [p_2 \log(1 + w) + (1 - p_2) \log(1 - w)] \\ &= \left(\frac{p_1 + p_2}{2} \right) \log(1 + w) + \left(1 - \frac{p_1 + p_2}{2} \right) \log(1 - w). \end{aligned}$$

Therefore, by denoting $\bar{p} = \frac{p_1 + p_2}{2}$ (this can be thought of as a blended probability of success), the Kelly criterion can be used here to identify the optimal size wager for maximum growth rate as $w^* = 2\bar{p} - 1 = p_1 + p_2 - 1$.

This expression can easily be extended to a portfolio of n wagers. For events $Y_1, \dots, Y_n$ with success probabilities $p_1, \dots, p_n$, it follows that

$$G = \left(\frac{1}{n} \sum_{i=1}^n p_i \right) \log(1 + w) + \left(1 - \frac{1}{n} \sum_{i=1}^n p_i \right) \log(1 - w),$$

with growth rate maximized at $w^* = 2\bar{p} - 1$ for $\bar{p} = \frac{1}{n} \sum_{i=1}^n p_i$.

3.3 Shannon Entropy of Discrete Returns

3.3.1 Investments Versus Wagers

As a measure of dispersion, variance is a better tool for measuring risk in situations where the magnitude of observed returns affects the severity of an investor's gain or loss. It is not so much a question of whether you win or lose, but rather how

much you win or lose. The dispersive nature of continuous returns make variance an excellent measure of risk for this purpose. Major losses are the least desired result, so they are effectively penalized on a squared scale.

As a measure of uncertainty, both entropy and relative entropy are more suited for measuring risk for investment strategies where the magnitude of observed random variables do not affect the severity of an investor's gain or loss, for example the discrete outcomes of gambling wagers. Consider a sports wager on the winner of a football game that pays 2 to 1 for a given bet in the sense for \$1 wagered, a loss forfeits the \$1 wager and a win returns \$2 (\$1 winnings plus the original \$1). In the event of a loss, the outcome is -100% , i.e., the severity of the loss is invariant to whether the team loses by 1 point or 20 points. Equivalently for a win, the outcome is $+100\%$, i.e., the severity of the gain is unchanged. In these situations, the information entropy of historical outcomes is more informative than variance because entropy measures the level of randomness in these returns, and is not skewed by the underlying events that have no affect on the magnitude of the returns.

3.3.2 Joint Entropy of a Portfolio of Discrete Return Assets

It is first necessary to establish some theory required to calculate the risk of a portfolio of discrete return assets. We start here with measuring the entropy and

joint entropy of a discrete returns. A portfolio containing just one binary asset (for example, a wager) exhibits a special case of the joint entropy calculation, the univariate Shannon entropy (1948) [87; 88] H . For a discrete random variable X with probability mass function $P(\cdot)$ of taking on values $x_1, \dots, x_n$, the Shannon entropy H is the average amount of information produced by X , defined as

$$H(X) = E(-\log P(X)) = -\sum_{k=1}^m P(x_k) \log P(x_k).$$

For a portfolio with more than one asset, the average amount of information is represented by the joint entropy. For n discrete random variables $X_1, \dots, X_n$ with $m_1, \dots, m_n$ probability respective states, the joint entropy $H(X_1, \dots, X_n)$ is given by

$$H(X_1, \dots, X_n) = -\sum_{k_1=1}^{m_1} \cdots \sum_{k_n=1}^{m_n} P_n(x_{1k_1}, \dots, x_{nk_n}) \log P_n(x_{1k_1}, \dots, x_{nk_n}),$$

for joint probability distribution P_n . The joint entropy of event outcomes is the main discriminator of risk in the DEPO method for portfolio selection presented in this chapter. Greater joint entropy (more fair randomness) represents lesser relative risk.

In the case of a discrete returns portfolio, each asset R_i in the portfolio has a corresponding vector of historical outcomes r_{ij} , for $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, T\}$. For a binary option portfolio, these asset returns r_{ij} can be either $+1$ or -1 (for $+100\%$ or -100% return outcomes corresponding to success and failure respectively).

Then the portfolio outcome at each point in time can be represented by the cross-sectional vector $\mathbf{r}_{\cdot j} = (r_{1j}, \dots, r_{nj})$. These are the event outcomes to be used in the portfolio calculation for joint entropy.

3.4 Minimum Relative Entropy

For the purposes of DEPO, the risk of an option portfolio is defined here as the relative entropy of portfolio returns, with respect to the uniform distribution. We start by providing some background theory on the topic here, and attest to its value as a portfolio risk measure.

3.4.1 Kullback-Leibler Divergence

The Kullback-Leibler divergence [52; 54] measures the distance (or more specifically, the directed divergence) between two probability distributions. If P and Q are two discrete distributions with the support $\mathcal{X}$, the Kullback-Leibler divergence between them, also known as the relative entropy of P with respect to Q , is given by

$$\begin{aligned} D_{KL}(P \parallel Q) &= - \sum_{x \in \mathcal{X}} P(x) \log \left(\frac{Q(x)}{P(x)} \right) \\ &= \sum_{x \in \mathcal{X}} P(x) \log \left(\frac{P(x)}{Q(x)} \right). \end{aligned}$$

Gibbs' inequality (see Mackay, 2003) [62] shows that the relative entropy is always non-negative. That is,

$$D_{KL}(P \parallel Q) \geq 0,$$

with equality if and only if $P = Q$ almost everywhere.

The use of relative entropy requires the existence of a target distribution—a distribution from which the observed distribution is measured. Minimizing the relative entropy ensures the observed distribution resembles the target distribution as closely as possible. In the case of a fair betting wager, we argue that the desired target distribution is the uniform distribution.

Let X be a discrete random variable with probability mass function P , whose Shannon entropy relates to the relative entropy. It can be shown that the Shannon entropy of X is equal to the entropy of the m -state discrete uniform distribution U_m (maximum entropy) less the relative entropy a discrete distribution P with respect to U_m ,

$$H(X) = \log(m) - D_{KL}(P \parallel U_m) \implies D_{KL}(P \parallel U_m) = \log(m) - H(X).$$

Using the same combinatorial technique employed for REPO in Section 2.3.3, the Shannon entropy of portfolio returns can be estimated empirically via probability generating functions. For a collection of n discrete return assets over time period $j = 1, \dots, T$, let $\mathbf{r}_j = (r_{1j}, \dots, r_{nj})$ denote the cross-sectional n -dimensional vector of

outcomes across one observational row of data, and let them be uniquely represented by the collection of $\mathbf{u}_k$'s such that $\mathbf{u}_k = \{\mathbf{r}_j \mid \mathbf{r}_j \neq \mathbf{u}_l, \text{ for some } j, \text{ and any } l \neq k\}$. For example, if the $\mathbf{r}_j$'s were $\{(1, 0, 1), (1, 0, 1), (1, 0, 1), (0, 1, 1), (0, 1, 1)\}$, then the $\mathbf{u}_k$'s would be $\{(1, 0, 1), (0, 1, 1)\}$. The empirical Shannon entropy of option portfolio returns R_Q can then be expressed as

$$H(R_Q) \approx - \sum_{k=1}^m \left(\frac{g^{(k)}(0)}{k!} \right) \log \left(\frac{g^{(k)}(0)}{k!} \right),$$

for k th-derivative at $x = 0$ of generating function

$$g(x; w_1, \dots, w_n) = \frac{1}{T} \sum_{j=1}^T x^{\{k: r_j = \mathbf{u}_k\}}.$$

Therefore the risk of a binary option portfolio can be measured by the relative entropy of portfolio returns R_Q , estimated empirically as

$$D_{KL}(R_Q \parallel U_m) \approx \log(m) + \sum_{k=1}^m \left(\frac{g^{(k)}(0)}{k!} \right) \log \left(\frac{g^{(k)}(0)}{k!} \right).$$

3.4.2 Relative Entropy as a Convex Risk Measure

As outlined by the Education and Examination Committee of the Society of Actuaries (Hardy, 2006) [36], a risk measure is a function $\rho : \mathbb{X} \rightarrow \mathbb{R}$, for a linear space $\mathbb{X}$ and random variable $X \in \mathbb{X}$, that maps a given loss distribution to the real number line. Risk measures are very popular in financial mathematics and actuarial science to quantify risk exposure, and are based on the so-called premium principles,

the purpose of which is to establish an appropriate premium to charge for a given risk. Several common early premium principles include the equivalence principle, expected value principle, variance principle, and standard deviation principle. For example, the variance principle is given by $\rho(X) = E(X) + k\text{Var}(X)$, for a fixed constant $k \geq 0$

Connections between utility maximization and entropy in the context of risk measures have been previously explored in the literature, and a comprehensive review is presented in the text by Follmer and Schied (2011) [31]. Of particular note is the entropic risk measure ρ_β given by $\rho_\beta(X) = \sup_P \left(E_P[-X] - \frac{1}{\beta} D_{KL}(P \parallel Q) \right)$, highlighting the connection between utility-based shortfall risk and divergence risk measures. To that end, here we would like to introduce the *relative entropy principle*, similar to that developed by Ahmadi-Javid (2016) [2], to evaluate the risk of portfolios based on (uniform) relative entropy. For $X \in \mathbb{X}$ according to a discrete distribution P , the relative entropy principle is given by

$$\rho(X) = E(X) + kD_{KL}(P \parallel U),$$

for the relative entropy $D_{KL}(P \parallel U)$ of P with respect to a discrete uniform distribution U that has the same support as P .

Further to the concept of risk measures, a convex risk measure (Artzner, 1999) [3] is a risk measure $\rho : \mathbb{X} \rightarrow \mathbb{R}$ that satisfies the following criteria for each $X, Y \in \mathbb{X}$:

- (i) *Monotonicity.* If $X \leq Y$, then $\rho(X) \leq \rho(Y)$,
- (ii) *Translation invariance.* If $c \in \mathbb{R}$, then $\rho(X + c) = \rho(X) + c$, and
- (iii) *Convexity.* $\rho(\lambda X + (1 - \lambda)Y) \leq \lambda\rho(X) + (1 - \lambda)\rho(Y)$, for $0 \leq \lambda \leq 1$.

We present some mathematical properties of relative entropy first, and then show that the relative entropy qualifies as a convex risk measure.

For independent distributions, relative entropy is additive similar to the way Shannon entropy is additive (Kullback, 1996) [53]. Let P_1, P_2 be independent distributions with joint probability mass function $P(x, y) = P_1(x)P_2(y)$, and similarly let Q_1, Q_2 be independent with $Q(x, y) = Q_1(x)Q_2(y)$, where P_i and Q_i have the same support for $i = 1, 2$. Then it follows that

$$D_{KL}(P \parallel Q) = D_{KL}(P_1 \parallel Q_1) + D_{KL}(P_2 \parallel Q_2).$$

Proposition 3.1 *Based on the entropy principle, relative entropy satisfies the three conditions required to be considered a valid convex risk measure: (i) monotonicity, (ii) translation invariance, and (iii) convexity.*

Proof 3.1 *Let $\rho(\cdot)$ be a risk measure based on the stated relative entropy principle, such that for constant $k \geq 0$ and $X \in \mathbb{X}$ following discrete distribution P , ρ is of the form $\rho(X) = E(X) + kD_{KL}(P \parallel U)$.*

(i) Monotonicity. For a risk measure $\rho(\cdot)$ to be monotonic it must satisfy: If $X, Y \in \mathbb{X}$ and $X \leq Y$ almost surely then $\rho(X) \leq \rho(Y)$ almost surely. Using $\rho(X)$ as

stated, we have $\rho(X) = E(X) + kD_{KL}(P \parallel U)$ and $\rho(Y) = E(Y) + kD_{KL}(Q \parallel U)$ for discrete distributions P and Q . Since $X \leq Y$ implies $D_{KL}(P \parallel U) \leq D_{KL}(Q \parallel U)$ as a consequence of the data processing inequality (Cover, 1991) [21], it follows that

$$\rho(X) = E(X) + kD_{KL}(P \parallel U) \leq E(Y) + kD_{KL}(Q \parallel U) = \rho(Y),$$

almost surely. Therefore ρ is monotonic.

(ii) *Translation invariance.* For a risk measure $\rho(\cdot)$ to exhibit translation invariance it must satisfy: If $X \in \mathbb{X}$ then $\rho(X+c) = \rho(X)+c$. Recall the risk measure based on the relative entropy principle is of the form $\rho(X) = E(X) + kD_{KL}(P \parallel U)$. Since $H(X+c) = H(X)$, for all c , it follows that $D_{KL}(P(X) \parallel U) = D_{KL}(P(X+c) \parallel U)$ and thus

$$\rho(X+c) = E(X+c) + kD_{KL}(P(X+c) \parallel U) = E(X) + c + kD_{KL}(P(X) \parallel U) = \rho(X) + c.$$

Therefore ρ exhibits translation invariance.

(iii) *Convexity.* A risk measure $\rho(\cdot)$ is convex if: For $Z_1, Z_2 \in \mathbb{X}$ and $\lambda \in [0, 1]$ it follows that: $\rho(\lambda Z_1 + (1-\lambda)Z_2) \leq \lambda\rho(Z_1) + (1-\lambda)\rho(Z_2)$. It is known that $D_{KL}(P \parallel Q)$ is convex in the pair of probability mass functions (P, Q) . If (P_1, Q_1) and (P_2, Q_2) are two pairs of probability mass functions, then

$$D_{KL}(\lambda P_1 + (1-\lambda)P_2 \parallel \lambda Q_1 + (1-\lambda)Q_2) \leq \lambda D_{KL}(P_1 \parallel Q_1) + (1-\lambda)D_{KL}(P_2 \parallel Q_2).$$

Therefore ρ is a convex measure. And thus, based on the relative entropy principle,

relative entropy is a valid convex risk measure since all three necessary conditions are satisfied.

3.4.3 Minimum Risk Option Portfolios with Relative Entropy

When selecting a portfolio of securities, the objective was to minimize the discrete entropy of portfolio returns, as seen in Section 2.3.5 (i.e., to minimize entropy, and maximize expected returns). Low entropy means low risk. In contrast, when dealing with gambling outcomes the opposite is true as maximum entropy proves to be the lowest risk option, due to the uniformity of gambling outcomes (as opposed to normality). This can be demonstrated in the following example. Consider an online casino website that offers a virtual roulette game. This roulette game is a black-box (its programming code or internal structure is unknown), but imagine that the site displays the entropy of the red or black outcomes over the past 100 spins. Take two independent roulette wheels, A and B . Wheel A 's past 100 spins were evenly split between red and black, uniformly 50/50, giving an entropy of 1 bit, the maximum entropy possible (using base 2 entropy, since we only have binary outcome here). Wheel B 's past 100 spins were 75 of one colour and 25 of the other, for a (lesser) entropy of 0.8113 bits. Which wheel is lower risk?

Without knowing the colour advantage of wheel B , clearly the fairness of wheel

A is lower risk between the two. Arguably, even if one did know the 75/25 split was in favour of red, wheel B is still riskier. One may play on the “gambler’s fallacy” (Keren, 1994) [49] and bet on black with the conviction that some black spins must be forthcoming to even out the odds. Or one may assume that the wheel’s randomness is flawed and bet on red to capitalize on the fault. Both strategies carry much greater risk than placing any colour bet on wheel A . Therefore it is evident that for uniform-type distributions, such as gambling wager outcomes, maximal entropy is the desired minimum-risk method.

As discussed in Section 3.4.2, based on the relative entropy principle, the relative entropy is proven to be a convex risk measure. Thinking back to the REPO problem in Section 2.3.5, the goal of the objective function was to minimize the portfolio entropy. Here we wish to minimize the portfolio relative entropy with respect to the uniform density, which means the goal is to obtain as close to a uniform distribution as possible. This would be analogous to maximizing the portfolio entropy in REPO, but this proves to be the minimal risk portfolio for a collection of gambling wagers in DEPO.

One interesting point to note is that DEPO assigns an equally-weighted allocation across chosen assets. This is due to the nature of joint entropy and joint relative entropy. Joint entropy is a measurement determined strictly by the inclusion or

exclusion of random variables. Consequently, a fractional inclusion of a random variable contributes the same amount of marginal entropy as the full inclusion, by the property of joint entropy $H(aX_1, bX_2) = H(X_1, X_2)$, for any $0 < a, b \leq 1$. Therefore, while the total allocation is determined by the growth rate objective function, it is split equally amongst the selected assets in the portfolio. As seen in recent works such as Low et al. (2016) [61], the equally-weighted strategy still remains difficult to outperform in the portfolio selection problem.

3.4.4 Discrete Entropic Portfolio Optimization (DEPO)

The new discrete entropic portfolio optimization (DEPO) problem uses a multi-objective function that minimizes the empirical relative entropy and maximizes expected growth rate. Using this optimization, users can make portfolio selections based on a chosen risk tolerance. The highest risk portfolio solely maximizes the expected portfolio growth rate, which is equivalent to the Kelly criterion method. The lowest risk portfolio minimizes the portfolio relative entropy, which is the most diversified portfolio allocating capital to all n options equally. Somewhere in between lies a user's optimal portfolio of choice. For the two-state case of binary options we would have a series of events with possible outcome states L (loss) and W (win), which can be denoted as outcomes $u \in \{0, 1\}$. This leads to the sim-

plified DEPO problem. Consider n events with potential investment opportunities. Let p_i represent the probabilities of success for event $i \in \{1, \dots, n\}$, and let w_i represent the percentage of portfolio to be allocated to event i , with the total portfolio allocation summing to $\omega = w_1 + \dots + w_n$. Let $\mathbf{r}_j = (r_{1j}, \dots, r_{nj})$ denote the cross-sectional n -dimensional vector of outcomes across one observational row of data. Over T data points this leads to m historical unique vectors of 0's and 1's, $\mathbf{u}_k = \{\mathbf{r}_j \mid \mathbf{r}_j \neq \mathbf{u}_l, \text{ for some } j, \text{ and any } l \neq k\}$ for $k = 1, \dots, m$ such that each $\mathbf{u}_k$ is unique, with m bounded by either T or the maximum number of possible combinations λ^n , so $m = \min(T, \lambda^n)$. Basically the collection of $\mathbf{u}_k$'s is a unique representation of the $\mathbf{r}_j$'s with no duplicates. Let us also denote $\eta = \sum_{i=1}^n I(w_i) \leq n$ as the number of chosen options in the portfolio, where $I(w_i)$ is the indicator function for the event $\{w_i > 0\}$. Then the DEPO problem is defined as the following optimization program (using logarithm base 2 since we are dealing with binary event

outcomes u),

$$\begin{aligned}
& \text{minimize} \quad D_{KL}(R_Q \parallel U_m) = \log_2(m) + \sum_{k=1}^m \left(\frac{g^{(k)}(0)}{k!} \right) \log_2 \left(\frac{g^{(k)}(0)}{k!} \right) \\
& \text{maximize} \quad G(\omega) = \left(\frac{1}{\eta} \sum_{i=1}^n I(w_i) p_i \right) \log_2(1 + \omega) + \left(1 - \frac{1}{\eta} \sum_{i=1}^n I(w_i) p_i \right) \log_2(1 - \omega), \\
& \text{subject to} \quad \omega = w_1 + \dots + w_n \leq 1,
\end{aligned}$$

$$w_i \geq 0 \quad \forall i,$$

$$w_i = w_j = \eta^{-1} \omega \quad \forall \{(i, j) : I(w_i) = I(w_j) = 1\},$$

for the m -state uniform distribution U_m and k th-derivative at $x = 0$ of probability generating function

$$g(x; w_1, \dots, w_n) = \frac{1}{T} \sum_{j=1}^T x^{\{k : (I(w_1)r_{1j}, \dots, I(w_n)r_{nj}) = \mathbf{u}_k\}}.$$

The last constraint in the optimization problem stems from the fact that joint entropy measures randomness strictly based on the inclusion or exclusion of a random variable. The joint entropy calculation is completely indifferent to the size of allocation, so any non-zero weight w_i contributes the corresponding marginal entropy from asset i , regardless of the magnitude of w_i . For this reason every asset included in the portfolio is assigned an equal weighting of $\eta^{-1} \omega$.

3.4.5 Risk-Adjusted Performance

Now that a risk-reward framework for finding optimized portfolios has been established, we can calculate a risk-adjusted return ratio for comparing portfolios growth rates of different risk profiles. This is analogous to the “reward-to-variability ratio” by Sharpe (1966) [90], better known as the Sharpe ratio. The Sharpe ratio of a portfolio is defined as

$$S_n = \frac{E(R_a - R_b)}{\sigma_a} = \frac{E(R_a - R_b)}{\sqrt{\text{Var}(R_a - R_b)}},$$

where R_a is the portfolio return, R_b is the risk-free rate of return, and σ_a is the standard deviation of the portfolio excess return.

As shown by Eling (2008) [26] and Rad et al. (2016) [76], for returns that are not normally distributed it is well known that the Sharpe ratio has the potential to underestimate risk, thereby overestimates the risk-adjusted performance. The authors employed methods such as lower partial moment measures and drawdown measures to circumvent this bias. To that end, here we introduce the risk-adjusted ratio for comparing growth rates of discrete return portfolios, called the *Growth Rate Over UNiform Divergence (GROUND) ratio*. This ratio measures the expected growth rate of a portfolio, adjusted by its risk level—relative entropy with respect to the uniform distribution. Let U_m be the m -state discrete uniform distribution. For chosen portfolio R_a and minimum risk portfolio R_b existing in the m -state event

space, the GROUND ratio Γ_m is defined as

$$\begin{aligned}\Gamma_m &= \frac{E(G_a(\omega_a) - G_b(\omega_b))}{D_{KL}(R_a \parallel U_m) - D_{KL}(R_b \parallel U_m)} \\ &= \frac{E(G_a(\omega_a) - G_b(\omega_b))}{H(R_b) - H(R_a)},\end{aligned}$$

where $G_a(\omega_a)$ is the growth rate of the chosen portfolio with weighting ω_a , $G_b(\omega_b)$ is the growth rate of the minimum risk portfolio with weighting ω_b , $D_{KL}(R_a \parallel U_m)$ is the relative entropy of the chosen portfolio with respect to U_m , $D_{KL}(R_b \parallel U_m)$ is the relative entropy of the minimum risk portfolio with respect to U_m , and $H(\cdot)$ is the Shannon entropy.

3.5 Portfolio Selection Examples with DEPO

3.5.1 FOREX Binary Option Portfolio Example

3.5.1.1 Data

In this example, actual binary option data is presented from the ten foreign exchange currency pairs available for trading on the North American Derivatives Exchange (NADEX), which can be found at www.NADEX.com. All historical intraday contract prices, strikes and outcomes are from the time period January 2019 to January 2020, totalling 346,760 historical trades. Daily option contracts can expire at intervals of four hours, namely 3:00AM, 7:00AM, 11:00AM, 3:00PM, 7:00PM and

11:00PM. Note that one can bet either “for” or “against” each option (buy or sell). Each currency pair has a series of available option contracts to choose from, with estimated success rates ranging anywhere from 5% to 95%. Note that all probabilities for success rates of these options are only estimates based on market consensus. For the purposes of this chapter, we focus solely on options with an estimated success rate of around 50% in order to ensure the expected win payoff is close to the amount wagered, so we restrict the data to observations with market consensus of 45% to 55% estimated probability of expiring in-the-money. This results in 10,435 remaining observations encompassing 1,332 possible contract expiry dates for each currency pair. Intraday option contract outcomes versus historical strike prices are recorded and computed as follows,

$$I(E_P > S_P) - I(E_P \leq S_P),$$

where E_P is the currency pair expiration price, and S_P is the strike price of the option (the minimum price necessary to become in-the-money). The result is 1 or -1 for respectively in-the-money or out-the-money, and 0 for an unavailable trade. Using this historical data we are able to empirically calculate the estimated relative entropy for each option. The binary options and their respective outcomes against strike prices over January 2019 to January 2020 are presented below in Table 3.1, together with their estimated relative entropies.

Table 3.1: Mean, in-the-money rate and estimated relative entropy (in bits) of FOREX binary options from January 2019 to January 2020.

Currency Pair	Short Name	Mean Outcome	% In-the-Money	Relative Entropy
Australian Dollar/Japanese Yen	AUD/JPY	-0.42591	0.287045	0.135127
Australian Dollar/US Dollar	AUD/USD	-0.868988	0.065506	0.651076
Euro/Pound Sterling	EUR/GBP	0.0174	0.5087	0.000218
Euro/Japanese Yen	EUR/JPY	-0.243882	0.378059	0.04334
Euro/US Dollar	EUR/USD	-0.109232	0.445384	0.008624
Pound Sterling/Japanese Yen	GBP/JPY	-0.069511	0.465244	0.003488
Pound Sterling/US Dollar	GBP/USD	-0.500692	0.249654	0.18927
US Dollar/Canadian Dollar	USD/CAD	0.468873	0.734437	0.164974
US Dollar/Swiss Franc	USD/CHF	-0.889943	0.055028	0.692615
US Dollar/Japanese Yen	USD/JPY	-0.370885	0.314557	0.101636

For the month of February 2020 there are 110 short-term FOREX binary option contract expiration periods, and projections (estimates) for each option are presented from NADEX market consensus probability of expiring above the strike price (in-the-money). These projections measure with what probability market bettors are expecting currency pairs to land in-the-money, as explained in the NADEX lessons page [70]. At each contract expiration period, there are ten available currency pairs to bet on, as shown in Table 3.1, and the historical results summarized there are used to evaluate the estimated relative entropy risk of each option. The emulation

here shows how DEPO performs against leading Kelly criterion methods for picking a portfolio of options at each contract expiration over the course of February 2020.

For illustrative purposes, let us examine this method applied to the first contract expiry period, Sunday February 2, 2020 11:00PM (as FOREX markets are open globally 24 hours a day during the weekdays, this period is the first available contract expiry in February since markets have already opened Monday morning in Asia). Table 3.2 lists the details of select currency pair binary options, with contract strike price and the market consensus estimates for each option.

Table 3.2: Select FOREX binary options for February 2, 2020 11:00PM, with their respective contract strike price, market consensus edge, and estimated probability in-the-money.

Currency Pair	Contract Strike Price	Market Edge	Probability In-the-Money
Australian Dollar/Japanese Yen	AUD/JPY>72.60	1.75%	48.25%
Australian Dollar/US Dollar	AUD/USD>0.6700	2.25%	52.25%
Euro/Pound Sterling	EUR/GBP>0.8420	0%	50%
Euro/Japanese Yen	EUR/JPY>120.20	0%	50%
Euro/US Dollar	EUR/USD>1.1080	1.5%	51.5%
Pound Sterling/Japanese Yen	GBP/JPY>142.80	1%	49%
Pound Sterling/US Dollar	GBP/USD>1.3180	1.5%	51.5%
US Dollar/Canadian Dollar	USD/CAD>1.3240	2.625%	47.375%
US Dollar/Swiss Franc	USD/CHF>0.9640	2.5%	47.5%
US Dollar/Japanese Yen	USD/JPY>108.40	2.75%	47.25%

DEPO determines which collection of options to buy or sell and what percentage of portfolio to allocate to each, in order to build the optimal risk-reward binary option portfolio.

Each potential portfolio has an expected growth rate given the consensus estimates, and an estimated relative entropy with respect to the uniform distribution, calculated empirically over the historical data. The historical data contains $T = 1332$ data points for each currency pair, and the total number of combinations for $n = 10$

binary outcomes is $2^{10} = 1024$, so the maximum joint entropy that can possibly be exhibited is $\log_2(2^n) = 10$, corresponding to the uniform distribution U_m with $m = 1024$ possible probability states. Therefore, a portfolio R_Q with joint entropy $H(R_Q)$ has an estimated relative entropy of

$$\begin{aligned} D_{KL}(R_Q \parallel U_m) &= \log_2(2^n) - H(R_Q) \\ &= 10 - H(R_Q). \end{aligned}$$

3.5.1.2 Efficient Frontier and Portfolio Selection

In the portfolio selection problem, the efficient frontier refers to the set of optimal portfolios that yield the greatest expected return for a defined level of risk, or equivalently the least risk for a defined level of expected return (the dual problem). The efficient frontier illustrates the risk-return trade-off for a given set of optimal portfolios. Here we show the analogous efficient frontier for portfolios with discrete return assets, comparing the expected growth rates and risk levels (estimated relative entropy) of each efficient portfolio. For the same contract expiry period of February 2, 2020 11:00PM, Figure 3.1 below plots the potential portfolios and their respective expected growth rates against their inherent risk profile, the estimated relative entropy with respect to the uniform distribution $\text{Uniform}(2^n)$ of historical joint outcomes of the portfolio.




Figure 3.1: Growth rate by relative entropy efficient frontier for FOREX binary option portfolio.

For the current month emulation, the Kelly criterion strategy chooses the portfolio R_K that maximizes the expected growth rate. For the first contract expiration period, this leads to the top right-most data point $K = (9.0198, 0.0022)$, with estimated relative entropy of 9.0198 and expected growth rate of 0.0022. This portfolio consists of just the last currency pair listed in Table 3.2, selling USD/JPY>108.40 to land out-the-money with 52.75% probability of success. According to the Kelly criterion, the optimal bet size here is $2p - 1 = 2 \times (0.5275) - 1 = 0.055$, and thus the chosen portfolio for this period is using 5.5% of portfolio funds to sell USD/JPY>108.40, shown in Table 3.3. This strategy disregards any concept of risk associated with the expected portfolio growth rate of $G(\omega) = 0.0022$.

Table 3.3: The Kelly criterion portfolio of options for contracts expiring February 2, 2020 11:00PM.

Option Contract	Buy or Sell	Success Probability	Kelly Wager %
USD/JPY>108.40	Sell	52.75%	5.5%

Alternatively, DEPO chooses the optimal portfolio based on the risk-reward trade-off. For each of the 1,332 contract expiration periods, portfolio selection is performed according to the following DEPO problem with $n = 10$ and $T = 1332$,

$$\begin{aligned} \text{maximize } & G(\omega) = \left(\frac{1}{\eta} \sum_{i=1}^n I(w_i) p_i \right) \log_2(1 + \omega) + \left(1 - \frac{1}{\eta} \sum_{i=1}^n I(w_i) p_i \right) \log_2(1 - \omega), \\ \text{subject to } & D_{KL}(R_Q \| U_m) = \log_2(2^n) + \sum_{k=1}^T \left(\frac{g^{(k)}(0)}{k!} \right) \log_2 \left(\frac{g^{(k)}(0)}{k!} \right) \leq 5, \end{aligned}$$

$$\omega = w_1 + \dots + w_n \leq 1,$$

$$w_i \geq 0 \quad \forall i,$$

$$w_i = w_j = \eta^{-1} \omega \quad \forall \{(i, j) : I(w_i) = I(w_j) = 1\},$$

for Uniform(2^n) as the target distribution, and for the k th-derivative at $x = 0$ of probability generating function

$$g(x; w_1, \dots, w_n) = \frac{1}{T} \sum_{j=1}^T x^{\{k : (I(w_1)r_{1j}, \dots, I(w_n)r_{nj}) = \mathbf{u}_k\}}.$$

While the Kelly criterion places the entire allocation on the option (or options) with

the greatest estimated probability of success, DEPO diversifies the portfolio by distributing the percent allocation across multiple options according to the appropriate risk profile. For period 1, DEPO selects data point $D = (4.9264, 0.0013)$, with estimated relative entropy of 4.9264 and expected growth rate of 0.0013. This corresponds to the optimal portfolio of six option contracts listed below in Table 3.4, with a total portfolio allocation of 4.2% (compared to the 5.5% allocated by the Kelly criterion).

Table 3.4: DEPO portfolio of options for contracts expiring February 2, 2020 11:00PM.

Option Contract	Buy or Sell	Success Probability	DEPO Wager %
USD/JPY>108.40	Sell	52.75%	0.7%
USD/CAD>1.3240	Sell	52.625%	0.7%
USD/CHF>0.9640	Sell	52.5%	0.7%
AUD/JPY>72.60	Sell	51.75%	0.7%
EUR/USD>1.1080	Buy	51.5%	0.7%
GBP/USD>1.3180	Buy	51.5%	0.7%

Portfolio efficiency can be measured by using the risk-adjusted GROUND ratio. In regards to the risk-adjusted expected returns, the DEPO portfolio has a GROUND ratio of $\Gamma_m = (0.0013 - 0.0007)/(4.9264 - 3.0546) = 0.0321\%$, over 25% more efficient than the Kelly criterion portfolio at $\Gamma_m = (0.0022 - 0.0007)/(9.0198 - 3.0546) =$

0.0251%.

The actual results that follow this period saw USD/JPY expire at 108.23, below the strike price of 108.40. Therefore the Kelly criterion strategy experiences a gain of 5.5% of portfolio balance in period 1. In regards to the DEPO portfolio, gains from USD/JPY and AUD/JPY sells expiring out-the-money were offset by USD/CAD and USD/CHF expiring in-the-money (losing both sell options) and EUR/USD and GBP/USD expiring out-the-money (losing both buy options), for a total loss of 1.4% in period 1.

3.5.1.3 Comparison to the Kelly Criterion Over Time

We demonstrate here the performance of DEPO versus the Kelly and Kelly-variant strategies over the full month February 2020 of FOREX binary options on NADEX. Methods in the previous Section 3.5.1.2 are repeated multiple times per day at four-hour interval contract expiry times. The Kelly criterion strategy wagers the optimal bet allocation each period on the option (or options) that have the greatest estimated probability of success. The half Kelly is the same strategy but utilizes the fractional Kelly-variant by wagering half the Kelly criterion bet size on the same options. DEPO optimal risk strategy uses the DEPO algorithm each period to select the portfolio with the greatest expected growth rate subject to the main constraint

that the portfolio has estimated relative entropy of no greater than 5. Each strategy began the month with \$1,000 and the total results are shown below in Figure 3.2.




Figure 3.2: Comparison of FOREX binary option strategies over February 2020.

While the Kelly and half Kelly strategies show massive variability with large portfolio highs and lows, DEPO remains consistently stable and ends the month with a modest profit of 9.1%, up \$91. The half Kelly finishes the month at a loss of 8.4% to \$916, and the full Kelly finishes at a significant loss of 36.4% to \$636. As the top market consensus predictions perform negatively through contracts 30-40 as well as 80+, DEPO's diversification generates consistent returns and builds on profits. The main purpose of DEPO is to mitigate risk of inaccurate predictions, and the goal is well accomplished.

The extended emulation is even more telling. We continue emulating DEPO and

the Kelly strategies over the month of March 2020, and while the Kelly criterion strategy begins to deteriorate rather drastically, DEPO strategy holds quite strong throughout the entire period and finishes the month with just a modest loss of 15% at \$850. The half Kelly finishes at a loss of 34.8% to \$652, and the full Kelly down a whopping 72.3% to only \$277. The results are illustrated below in Figure 3.3.




Figure 3.3: Comparison of FOREX binary option strategies over February and March 2020.

3.5.2 NFL Sportsbook Example

3.5.2.1 Data

In the example provided here, actual game data is presented for all 32 teams from the National Football League (NFL) over the past several seasons. Unfortunately,

the NFL season is a prime example of extremely small sample size data. The regular season consists of only 17 weeks and teams play only 16 regular season games. This would normally pose a problem for traditional portfolio selection methods based on normality, but the non-parametric approach by DEPO makes it well capable of handling such small sample data. Using archived data from www.teamrankings.com/nfl and www.archive.org, we were able to gather eight full past seasons of NFL games with historical Las Vegas point spreads from 2011-12 to 2018-19 seasons, totalling 136 weeks of results. Weekly game outcomes versus historical points spreads were recorded and computed as follows

$$I(M + S_{LV} > 0) - I(M + S_{LV} < 0),$$

where M is the team's winning margin (negative in case of a loss), and S_{LV} is the Vegas pre-game point spread (negative in case of a favourite). This results in 1 or -1 for respectively covering the spread or not, and 0 for a tie or week off. Using this historical data we can empirically calculate the estimated relative entropy for each team. The teams and their respective outcomes versus spreads over the 2011-12 to 2018-19 seasons are presented below in Table 3.5.

Table 3.5: Mean, cover rate, estimated relative entropy (bits) of NFL teams 2011-19.

Team Name	Short Name	Mean Outcome	Cover Rate	Relative Entropy
Arizona Cardinals	ARI	0.056	0.528	0.002263
Atlanta Falcons	ATL	-0.079365	0.460317	0.004548
Baltimore Ravens	BAL	-0.081967	0.459016	0.004852
Buffalo Bills	BUF	-0.04	0.48	0.001154
Carolina Panthers	CAR	0.080645	0.540323	0.004697
Chicago Bears	CHI	-0.031746	0.484127	0.000727
Cincinnati Bengals	CIN	0.173554	0.586777	0.021838
Cleveland Browns	CLE	-0.121951	0.439024	0.010755
Dallas Cowboys	DAL	-0.02439	0.487805	0.00043
Denver Broncos	DEN	0	0.5	0
Detroit Lions	DET	-0.064516	0.467742	0.003004
Green Bay Packers	GB	0.072	0.536	0.003742
Houston Texans	HOU	-0.02439	0.487805	0.00043
Indianapolis Colts	IND	0.088	0.544	0.005593
Jacksonville Jaguars	JAC	-0.114754	0.442623	0.00952
Kansas City Chiefs	KC	0.095238	0.547619	0.006553
Los Angeles Chargers	LAC	-0.031746	0.484127	0.000727
Los Angeles Rams	LAR	-0.112903	0.443548	0.009215
Miami Dolphins	MIA	-0.04065	0.479675	0.001192
Minnesota Vikings	MIN	0.193548	0.596774	0.027194
New England Patriots	NE	0.2	0.6	0.02905
New Orleans Saints	NO	0.129032	0.572581	0.015254
New York Giants	NYG	-0.00813	0.495935	0.000047
New York Jets	NYJ	-0.081967	0.459016	0.004852
Oakland Raiders	OAK	-0.064516	0.467742	0.003004
Philadelphia Eagles	PHI	-0.055118	0.472441	0.002192
Pittsburgh Steelers	PIT	0.024	0.512	0.000415
Seattle Seahawks	SEA	0.163934	0.581967	0.019473
San Francisco 49ers	SF	0.008	0.504	0.000046
Tampa Bay Buccaneers	TB	-0.096774	0.451613	0.006767
Tennessee Titans	TEN	-0.196721	0.401639	0.028099
Washington Redskins	WAS	-0.015625	0.492188	0.000176

Over the 17 weeks of the 2019-20 regular season, projections for each game were gathered from pre-game market betting consensus as presented on the www.covers.com/sports/nfl/matchups website. These estimates measure what percentage of market bettors are betting on either side of the Las Vegas spread line. Interestingly, over the course of the 17 week regular season, these market consensus estimates performed about as well as a coin-toss, with a weighted average of -0.73% versus actual outcomes (approximately 50% accuracy). Despite this, the market consensus estimates are unrealistically high. Therefore, in order to use more sensible estimates, as well as demonstrate the flexibility DEPO has to user inputs, we use posterior probabilities assigned as half of each team's market consensus edge. The historical results summarized in Table 3.5 are used to evaluate the estimated relative entropy risk of each game. On each Sunday, there are between 11 and 16 games taking place. The emulation here shows how DEPO performs against the leading Kelly criterion methods for picking a portfolio of wagers on each of the 17 Sundays throughout the regular season.

For illustrative purposes, let us examine this method applied to the first week: Sunday September 8, 2019, which aired 13 NFL games in total. Table 3.6 lists the scheduled games, pre-game Vegas point spreads and the market consensus estimates for each game.

Table 3.6: Scheduled NFL games for Sunday September 8, 2019, with their respective Las Vegas point spreads and market consensus estimated probabilities of covering.

Away Consensus	Away Team	Home Team	Home Consensus
70%	Kansas City Chiefs -3.5	Jacksonville Jaguars	30%
44%	Tennessee Titans	Cleveland Browns -5.5	56%
53%	Atlanta Falcons	Minnesota Vikings -3.5	47%
43%	Washington Redskins	Philadelphia Eagles -10.5	57%
67%	Baltimore Ravens -7.0	Miami Dolphins	33%
63%	Los Angeles Rams -1.5	Carolina Panthers	37%
46%	Buffalo Bills	New York Jets -2.5	54%
44%	Cincinnati Bengals	Seattle Seahawks -9.5	56%
34%	Indianapolis Colts	Los Angeles Chargers -6.0	66%
43%	San Francisco 49ers	Tampa Bay Buccaneers -1.0	57%
61%	Detroit Lions -2.5	Arizona Cardinals	39%
58%	New York Giants	Dallas Cowboys -7.0	42%
51%	Pittsburgh Steelers	New England Patriots -5.5	49%

Each game represents a potential bet opportunity, particularly betting on the team with higher consensus to cover the spread and betting on the team with lower consensus to not cover the spread. DEPO determines which collection of bets to select and what percentage of bankroll to wager on each, in order to build the optimal risk-reward sportsbook portfolio.

Each potential portfolio has an expected growth rate given the consensus pro-

jections, and an estimated relative entropy with respect to the uniform distribution, calculated empirically using the historical data. The historical data contains $T = 136$ data points for each team, so the maximum joint entropy that can possibly be exhibited is $\log_2(T) = 7.0875$, corresponding to a uniform distribution U_m with $m = 136$ possible probability states. Therefore, a portfolio R_Q with joint entropy $H(R_Q)$ has an estimated relative entropy of

$$\begin{aligned} D_{KL}(R_Q \parallel U_T) &= \log_2(T) - H(R_Q) \\ &= 7.0875 - H(R_Q). \end{aligned}$$

3.5.2.2 Efficient Frontier and Portfolio Selection

For the same first week of Sunday September 8, 2019, Figure 3.4 below plots the potential sportsbook portfolios and their respective expected growth rates against their inherent risk profile, the estimated relative entropy with respect to the uniform distribution $\text{Uniform}(T)$ of historical joint outcomes of the portfolio.




Figure 3.4: Growth rate by relative entropy efficient frontier for NFL sportsbook portfolio.

For the current season emulation, the Kelly criterion strategy chooses the portfolio R_K that maximizes the expected growth rate. For week 1 this leads to the top right-most data point $K = (5.089, 0.029)$, with estimated relative entropy of 5.089 and expected growth rate of 0.029. This portfolio consists of just the first game listed in Table 3.6, KC -3.5 @ JAC, taking KC to cover with a 60% probability of success (half the edge of the 70% market consensus). According to the Kelly criterion, the optimal bet size here is $2p - 1 = 2 \times (0.6) - 1 = 0.2$, and thus the chosen portfolio for week 1 is 20% of bankroll on KC -3.5, shown in Table 3.7. This strategy disregards any concept of risk associated with the expected portfolio growth rate of $G(\omega) = 0.029$.

Table 3.7: The Kelly criterion portfolio of wagers with percent allocation for NFL 2019-20 season week 1.

Bet to Cover	Bet to Not Cover	Success Probability	Kelly Wager %
Kansas City Chiefs -3.5	Jacksonville Jaguars	60%	20%

Alternatively, DEPO chooses the optimal portfolio based on the risk-reward trade-off. For each of the weeks 1 to 17, portfolios selection is performed according the following DEPO problem with n being the number of available games that Sunday and $T = 136$,

$$\begin{aligned} \text{maximize } & G(\omega) = \left(\frac{1}{\eta} \sum_{i=1}^n I(w_i) p_i \right) \log_2(1 + \omega) + \left(1 - \frac{1}{\eta} \sum_{i=1}^n I(w_i) p_i \right) \log_2(1 - \omega), \\ \text{subject to } & D_{KL}(R_Q \| U_T) = \log_2(T) + \sum_{k=1}^T \left(\frac{g^{(k)}(0)}{k!} \right) \log_2 \left(\frac{g^{(k)}(0)}{k!} \right) \leq 2, \end{aligned}$$

$$\omega = w_1 + \dots + w_n \leq 1,$$

$$w_i \geq 0 \quad \forall i,$$

$$w_i = w_j = \eta^{-1} \omega \quad \forall \{(i, j) : I(w_i) = I(w_j) = 1\},$$

for $\text{Uniform}(T)$ as the target distribution, and for the k th-derivative at $x = 0$ of probability generating function

$$g(x; w_1, \dots, w_n) = \frac{1}{T} \sum_{j=1}^T x^{\{k : (I(w_1)r_{1j}, \dots, I(w_n)r_{nj}) = \mathbf{u}_k\}}.$$

While the Kelly criterion places the entire wager on the game (or games) with the greatest estimated win probability, DEPO diversifies the portfolio by distributing the percent allocation across multiple wagers according to the appropriate risk profile. For week 1, DEPO selects data point $D = (1.5049, 0.0226)$, with estimated relative entropy of 1.5049 and expected growth rate of 0.0226. This corresponds to the optimal portfolio of three game wagers listed below in Table 3.8.

Table 3.8: DEPO portfolio of wagers with percent allocation for NFL 2019-20 season week 1.

Bet to Cover	Bet to Not Cover	Success Probability	DEPO Wager %
Kansas City Chiefs -3.5	Jacksonville Jaguars	60%	5.89%
Baltimore Ravens -7.0	Miami Dolphins	58.5%	5.89%
Los Angeles Chargers -6.0	Indianapolis Colts	58%	5.89%

Looking at the portfolio efficiency via the risk-adjusted GROUND ratio, the DEPO portfolio has a GROUND ratio of $\Gamma_m = (0.0226 - 0.0135)/1.5049 = 0.6\%$, twice as efficient as the Kelly criterion portfolio at $\Gamma_m = (0.029 - 0.0135)/5.089 = 0.3\%$.

The actual results that follow for week 1 have KC winning by 14 (covered the spread), BAL winning by 49 (covered the spread), and LAC winning by 6 (push

against the spread). Therefore the Kelly criterion strategy experiences a gain of 20% of bankroll in week 1, while the competing DEPO strategy gains 11.78%.

3.5.2.3 Comparison to the Kelly Criterion Over Time

Demonstrated here is the performance of DEPO versus the Kelly and Kelly-variant strategies over the full 2019-20 NFL regular season. Methods in the previous Section 3.5.2.2 are repeated week by week over the course of 17 weeks. The Kelly criterion strategy wagers the optimal bet size each week on the game (or games) that have the greatest estimated probability of success. The half Kelly is the same strategy but utilizes the fractional Kelly-variant by wagering half the Kelly criterion bet size on the same games. DEPO optimal risk strategy uses the DEPO algorithm each week to select the portfolio with the greatest expected growth rate subject to the main constraint that the portfolio has an estimated relative entropy of no greater than 2. Each strategy begins the season with \$1,000 and the total results are shown below in Figure 3.5.




Figure 3.5: Comparison of sportsbook strategies over 17 weeks of 2019-20 season NFL games.

Though off to a slow start, DEPO ultimately outperforms both the Kelly and half Kelly methods over the 17 week period, and is the only strategy to produce a profit at the end. In fact, the DEPO strategy remains profitable throughout the entire experiment. As the top market consensus predictions deteriorated mid-season, the Kelly strategies suffer massive losses, while the diversification strategy of DEPO holds strong. The main purpose of DEPO is to mitigate risk of inaccurate predictions, and once again the goal is well accomplished. In the end, DEPO finishes the season with a profit of \$100, up 10% to \$1,100, while the half Kelly finishes at a loss of 14% and the full Kelly ends the season down 37% to \$630.

4 Option Portfolio Selection with Generalized Entropic Portfolio Optimization

4.1 Introduction

Our work on the topic of portfolio optimization has introduced novel entropy-based optimization problems that facilitate the selection of a variety of efficient portfolios. The return-entropy portfolio optimization (REPO) allocates capital to an equity portfolio based on the expected return and Shannon entropy of portfolio returns. REPO was adapted to form the discrete entropic portfolio optimization (DEPO) for application to portfolios comprised of assets with discrete distributed returns, like exotic instruments such as binary or digital options, or fixed-return options (FROs). In this chapter we further extend DEPO to accommodate mixed returns with *generalized entropic portfolio optimization (GEPO)*. GEPO can handle a combination of discrete and continuous returns, such as those exhibited by option strategies. Option strategies have expected return payoff functions that contain

intervals of discrete returns and intervals of continuous returns.

Traditional portfolio optimization methods, such as those employed by Markowitz mean-variance portfolio optimization (MVPO) [65], are not suited to handle the discretely distributed nature of option returns as these distributions cannot be described by mean and variance alone. So continuing with the use of relative entropy as the proxy for portfolio risk (as used in DEPO), GEPO is an optimization method containing an objective function that simultaneously maximizes the portfolio growth rate and minimizes the relative entropy of the portfolio with respect to the uniform distribution. In the case of option portfolios, GEPO selects a collection of options from a set of possible choices, for example a portfolio of credit spreads, in order to maximize the expected growth rate for a given level of relative entropy risk. By using the combinatorial generating functions to empirically calculate entropy, as done in REPO and DEPO, we are able to calculate the relative entropy risk of an option portfolio in GEPO by extending the scope to include both discrete and continuous returns. GEPO calculates the relative entropy over multiple return probability states, including return states with continuous returns. In terms of the expected return, the Kelly criterion provides valuable insights into maximizing the portfolio growth rate of an option portfolio. To that effect, we extend the Kelly criterion growth rate to include instances of both discrete and continuous returns. GEPO empowers investors

to quantitatively select a portfolio of options based on their risk-reward tolerance.

The remainder of this chapter is organized as follows. The following Section 4.1.1 gives a literature review of research on the topic of option portfolio optimization. Section 4.2 explains the technical details behind maximizing the portfolio growth rate by extending the Kelly criterion to generalized option strategies, with specific examples for several popular option strategies. Section 4.3 provides a brief review of information theory, Shannon entropy and Kullback-Leibler divergence, the foundation of entropic portfolio optimization. As the main feature of this chapter, the GEPO problem is presented in Section 4.4. Finally, Section 4.5 demonstrates an example of GEPO selecting a portfolio of equity credit spreads chosen from the S&P100 composite index, and shows how GEPO outperformed the Kelly criterion and alternative Kelly criterion methods.

4.1.1 Literature Review

Option portfolio optimization is a considerably new topic of research, with the earliest found work beginning only this millennium. For a constant relative risk aversion investor, Liu (2003) [60] modeled the stochastic volatility to optimize a portfolio comprising one equity, a put option, and cash, deriving an analytic solution for the optimal allocation. Unfortunately for this method, a parametric model must be

specified and the risk is mostly concentrated on just one option. Jones (2006) [45] exploited apparent mispricing of put options to derive an optimal portfolio of options using a general nonlinear latent factor model, but this model is overburdened by numerous required parameters. Eraker (2007) [27] modeled stochastic volatility parametrically, and then used the traditional mean-variance framework to optimize allocation between straddles, puts, and calls, yielding a closed-form solution for portfolio weights. Haugh (2007) [37] used duality and approximate dynamic programming (ADP) methods to facilitate high-dimensional American option pricing and portfolio optimization.

Driessen (2013) [25] used generalized method of moments to maximize expected returns for a portfolio comprising a stock, an option strategy (puts and straddles), and cash. Constantinides (2013) [19] leverage-adjusted portfolios of either calls or puts by using the *omega* or *lambda*, known as the options' elasticity. Fadugba (2014) [28] used a binomial model as a performance measure to price American and European options, and exploited mispricings to derive optimal portfolios. Sumarti (2016) [94] used a fuzzy binomial CRR procedure to perform dynamic asset allocation for optimal portfolios. Fatyanova (2017) [30] developed a constrained optimization problem for constructing an option portfolio that maximizes a certain payoff function. Faias (2017) [29] also noted that traditional portfolio optimization

methods like mean-variance optimization are not suitable for option portfolios due to non-normality and difficulty estimating distribution of returns, then introduced a short-term view objective function used to optimize portfolios of European options mainly by exploiting mispricing between options. Zhao (2018) [107] used first- and second-order moments to model options returns and extended the Markowitz mean-variance framework to include option selection with much lower computational time than off-the-shelf solvers. Zeng (2018) [106] introduced a progressive hedging algorithm using reinforcement learning (Q-learning) for option portfolio optimization that was first of its kind to consider the exercise timings of an American option.

Almost all of these alternative methods require some kind of parametric model assumption and are limited to the number of options composing the portfolio. GEPO is non-parametric and delivers well-diversified portfolios of many options with no limit. Additionally, all these methods misguidedly aim to maximize the portfolio expected returns, and there has not yet been any suggestion for measuring or managing the risk of these option portfolios. GEPO not only minimizes the relative entropy risk of the portfolio, but also maximizes the proper metric for measuring the future performance of option returns: exponential growth rate.

4.2 Maximum Exponential Growth Rate

4.2.1 The Kelly Criterion for Multiple Wagers

As presented in Section 3.2.2, we use the extension of the Kelly criterion (Kelly, 1956) [48] to n wagers. For events $Y_1, \dots, Y_n$ with success probabilities $p_1, \dots, p_n$, let w represent the total percentage of bankroll to be wagered. For a portfolio that allocates the wager equally across amongst events, the growth rate coefficient G would be

$$G(w) = \left(\frac{1}{n} \sum_{i=1}^n p_i \right) \log(1 + w) + \left(1 - \frac{1}{n} \sum_{i=1}^n p_i \right) \log(1 - w),$$

Therefore, by denoting $\bar{p} = \frac{1}{n} \sum_{i=1}^n p_i$ (this can be thought of as a blended probability of success), the Kelly criterion can be used here to identify the optimal size wager for maximum growth rate as $w^* = 2\bar{p} - 1$.

4.2.2 Extension of the Kelly Criterion to Option Strategies

This section is a brief review of several most popular option strategies, with a summary of the details and motivation behind each. Additionally, we extend the Kelly criterion to find optimal growth rates of each option strategy. Further details and strategies can be found in the TMX Group Montreal Exchange guides and strategies documentation (2020) [97].

4.2.2.1 Covered Call

A covered call option strategy involves selling a call that is covered by an equivalent long stock position, providing a hedge on the stock. In exchange for upside potential, an investor is able to earn income on the premium. The motivation behind a covered call is to earn premium income. The maximum loss is limited to the initial stock purchase price slightly reduced by the premium income from sale of the call option. The maximum gain is limited, equal to the strike price less the initial stock purchase price, plus the premium income received. A sample expected payoff function for a covered call strategy is illustrated below in Figure 4.1.




Figure 4.1: Sample expected payoff function for a covered call strategy.

For the price of underlying security S and option return R , the growth rate for a

covered call strategy becomes

$$G(w) = \varrho \log(1 + \alpha w) + p \log(1 + w),$$

where ϱ represents the probability $P(S \in I_1) = P(R < 1)$, p_2 represents the probability $P(S \in I_2) = P(R = 1)$ with $p + \varrho = 1$, and α is the expected value of returns conditional on underlying price falling on interval I_1 , $\alpha = E(R \mid S \in I_1)$. Then G is maximized by differentiating with respect to w and setting equal to zero,

$$\begin{aligned} \frac{\partial}{\partial w} G &= \frac{p}{1 + w} + \frac{\alpha \varrho}{1 + \alpha w} \\ \implies 0 &= p(1 + \alpha w^*) + \alpha \varrho(1 + w^*) \\ \implies w^* &= \frac{p(\alpha - 1)}{\alpha} - 1, \end{aligned}$$

which implies a positive w^* exists if and only if $p > \alpha/(\alpha - 1)$, for $\alpha \neq 0$.

4.2.2.2 Married Put

A married put, or protective put option strategy involves adding a long put position to a long stock position, forming a lower bound for the stock value. The investor profits as the stock price keeps rising. The motivation for a married put is to hedge against a temporary decrease in the stock price. The maximum loss is the stock purchase price less the strike price of the put plus the premium paid for the option. There is unlimited potential gain for a married put. A sample expected payoff function for a married put strategy is illustrated below in Figure 4.2.




Figure 4.2: Sample expected payoff function for a married put strategy.

For the price of underlying security S and option return R , the growth rate for a married put strategy becomes

$$G(w) = q \log(1 - w) + \varrho \log(1 + \alpha w),$$

where q represents the probability $P(S \in I_1) = P(R = -1)$, ϱ represents the probability $P(S \in I_2) = P(R > -1)$ with $q + \varrho = 1$, and α is the expected value of returns conditional underlying price landing in interval I_2 , $\alpha = E(R \mid S \in I_2)$. Then G is maximized by differentiating with respect to w and setting equal to zero,

$$\begin{aligned} \frac{\partial}{\partial w} G &= \frac{\alpha \varrho}{1 + \alpha w} - \frac{q}{1 - w} \\ \implies 0 &= \alpha \varrho (1 - w^*) - q (1 + \alpha w^*) \\ \implies w^* &= 1 - \frac{q(\alpha + 1)}{\alpha}, \end{aligned}$$

which implies a positive w^* exists if and only if $q < \alpha/(\alpha + 1)$, for $\alpha \neq 0$.

4.2.2.3 Credit Spread

A put credit spread, or bull put spread option strategy consists of a short put option at a certain strike price and a long put option at lower strike price. The investor profits with a rise in the underlying stock price. The motivations for a credit spread include to earn income with limited risk and to moderately profit from a rise in the stock price. The maximum loss is limited, equal to the net difference between the higher and lower strike prices less the net premium received. The maximum gain too is limited to the net premium received when putting on the position. A sample expected payoff function for a put credit spread strategy is illustrated below in Figure 4.3.




Figure 4.3: Sample expected payoff function for a put credit spread strategy.

Alternatively, a call credit spread, or bear call spread option strategy consists of a short call option at a certain strike price and a long call option at a higher strike price. This way the investor profits with a decrease in the underlying stock price. The maximum loss is limited, equal to the net difference between the higher and lower strike prices less the net premium received. The maximum gain too is limited to the net premium received when calling on the position. A sample expected payoff function for a put credit spread strategy is illustrated below in Figure 4.4.




Figure 4.4: Sample expected payoff function for a call credit spread strategy.

For the price of underlying security S and option return R , the growth rate for a credit spread strategy becomes

$$G(w) = p \log(1 + w) + \varrho \log(1 + \alpha w) + q \log(1 - w),$$

where p is the probability $P(R = 1)$, ϱ represents the probability $P(S \in I_2) = P(-1 < R < 1)$, q represents the probability $P(R = -1)$, with $p + q + \varrho = 1$, and α is the expected value of returns conditional on underlying price landing in interval I_2 , $\alpha = E(R | S \in I_2)$. Then G is maximized by differentiating with respect to w and

setting equal to zero,

$$\begin{aligned}\frac{\partial}{\partial w}G &= \frac{p}{1+w} + \frac{\alpha \varrho}{1+\alpha w} - \frac{q}{1-w} \\ \implies 0 &= p(1-w^*)(1+\alpha w^*) + \alpha \varrho(1+w^*)(1-w^*) - q(1+w^*)(1+\alpha w^*) \\ \implies w^* &= \frac{(\alpha p - \alpha q - p - q) + \sqrt{(\alpha p - \alpha q - p - q)^2 + 4\alpha(p - q + \alpha - \alpha p - \alpha q)}}{2\alpha},\end{aligned}$$

by the quadratic formula, for $\alpha \neq 0$.

4.2.2.4 Straddle

A straddle option strategy involves buying a call and buying a put with equal strike price and expiration date. The investor profits when the underlying stock price experiences a big move up or down. The motivation behind a straddle is to capitalize on correctly predicting a big price move or high volatility in the near future. The maximum loss for a straddle is limited to the premium paid for the call and put options. The potential gain is unlimited. A sample expected payoff function for a straddle strategy is illustrated below in Figure 4.5.




Figure 4.5: Sample expected payoff function for a straddle strategy.

For the price of underlying security S and option return R , the growth rate for a straddle strategy becomes

$$G(w) = \varsigma \log(1 - \beta w) + \varrho \log(1 + \alpha w),$$

where ς represents the probability $P(S \in I_1)$, ϱ represents the probability $P(S \in I_2)$, and expected values $\beta = -E(R | S \in I_1)$ and $\alpha = E(R | S \in I_2)$, while $\varrho + \varsigma = 1$.

Then G is maximized by differentiating with respect to w and setting equal to zero,

$$\begin{aligned} \frac{\partial}{\partial w} G &= \frac{\alpha \varrho}{1 + \alpha w} - \frac{\beta \varsigma}{1 - \beta w} \\ \implies 0 &= \alpha \varrho (1 - \beta w^*) - \beta \varsigma (1 + \alpha w^*) \\ \implies w^* &= \frac{\varrho}{\beta} + \frac{\varrho - 1}{\alpha}, \end{aligned}$$

which implies a positive w^* exists if and only if the odds ratio $\varrho/(1 - \varrho) > \beta/\alpha$, for $\alpha, \beta \neq 0$.

4.2.2.5 Long Strangle

A long strangle option strategy involves buying an out-of-the-money call option and an out-of-the-money put option with the same expiration date. A strangle is similar to a straddle except a straddle has equal strike price whereas a strangle has a call option with higher strike price than the put option. The investor profits when there is a very big move up or down in the stock price. The motivation behind a long strangle is to capture a big move in the stock price over the term of the option. The maximum loss for a straddle is limited to the net premium paid for the call and put options. A sample expected payoff function for a long strangle strategy is illustrated below in Figure 4.6.




Figure 4.6: Sample expected payoff function for a long strangle strategy.

For the price of underlying security S and option return R , the growth rate for a long strangle strategy becomes

$$G(w) = \varsigma \log(1 - \beta w) + q \log(1 - w) + \varrho \log(1 + \alpha w),$$

where ς represents the probability $P(S \in I_1)$, q represents the probability $P(S \in I_2) = P(R = -1)$, and ϱ represents the probability $P(S \in I_3)$, and expected values $\beta = -E(R|S \in I_1)$ and $\alpha = E(R|S \in I_3)$, while $q + \varrho + \varsigma = 1$. Then G is maximized by differentiating with respect to w and setting equal to zero,

$$\begin{aligned} \frac{\partial}{\partial w} G &= \frac{\alpha \varrho}{1 + \alpha w} - \frac{q}{1 - w} - \frac{\beta \varsigma}{1 - \beta w} \\ \implies 0 &= \alpha \varrho (1 - w^*) (1 - \beta w^*) - q (1 + \alpha w^*) (1 - \beta w^*) - \beta \varsigma (1 + \alpha w^*) (1 - w^*) \\ \implies w^* &= \frac{-(\alpha \beta q - \alpha \beta - \beta \varrho + \beta - \alpha q - \alpha \varrho) - \sqrt{(\alpha \beta q - \alpha \beta - \beta \varrho + \beta - \alpha q - \alpha \varrho)^2 - 4 \alpha \beta (\beta q + \beta \varrho - \beta + \alpha \varrho - q)}}{2 \alpha \beta}, \end{aligned}$$

by the quadratic formula, for $\alpha, \beta \neq 0$.

4.2.2.6 Butterfly Spread

A butterfly spread, or long call butterfly option strategy consists of two short calls at a middle strike price and two long calls, one at the lower strike and one at the higher strike price, all with the same expiration date. The investor profits by correctly predicting the underlying stock price at expiration. The motivation behind a butterfly spread is to capitalize from predicting a target stock price at the options expiry date. The maximum loss for a butterfly spread is the short call strike price less the lower long call strike price less the net premium paid. The potential gain is unlimited. A sample expected payoff function for a butterfly spread strategy is illustrated below in Figure 4.7.




Figure 4.7: Sample expected payoff function for a butterfly spread strategy.

For the price of underlying security S and option return R , the growth rate for a butterfly spread strategy becomes

$$G(w) = \frac{1}{2}q \log(1 - w) + \varrho \log(1 + \alpha w) + \varsigma \log(1 - \beta w) + \frac{1}{2}q \log(1 - w),$$

where q represents the probability $P(S \in I_1 \cup I_4) = P(R = -1)$, ϱ represents the probability $P(S \in I_2)$, ς represents the probability $P(S \in I_3)$, and expected values $\alpha = E(R|S \in I_2)$ and $\beta = -E(R|S \in I_3)$, while $q + \varrho + \varsigma = 1$. Then G is maximized by differentiating with respect to w and setting equal to zero,

$$\begin{aligned} \frac{\partial}{\partial w} G &= \frac{\alpha \varrho}{1 + \alpha w} - \frac{q}{1 - w} - \frac{\beta \varsigma}{1 - \beta w} \\ \implies 0 &= \alpha \varrho (1 - w^*) (1 - \beta w^*) - q (1 + \alpha w^*) (1 - \beta w^*) - \beta \varsigma (1 + \alpha w^*) (1 - w^*) \\ \implies w^* &= \frac{-(\alpha \beta q - \alpha \beta - \beta \varrho + \beta - \alpha q - \alpha \varrho) - \sqrt{(\alpha \beta q - \alpha \beta - \beta \varrho + \beta - \alpha q - \alpha \varrho)^2 - 4 \alpha \beta (\beta q + \beta \varrho - \beta + \alpha \varrho - q)}}{2 \alpha \beta}, \end{aligned}$$

by the quadratic formula, for $\alpha, \beta \neq 0$.

4.2.2.7 Iron Condor

An iron condor, or short condor option strategy involves selling one call and buying another call with a higher strike price, plus selling one put and buying another put with a lower strike price, with the current underlying price falling between the call and put strikes. The investor profits if the underlying stock price is between the call and put strikes at option expiration. The motivation behind an iron condor is when the investor foresees the stock trading in a narrow range over the life of the options. The maximum loss of an iron condor is the greater of the difference between high and low call strikes and high and low put strikes, less the net premium received. A sample expected payoff function for an iron condor strategy is illustrated below in Figure 4.8.




Figure 4.8: Sample expected payoff function for an iron condor strategy.

For the price of underlying security S and option return R , the growth rate for an iron condor strategy becomes

$$G(w) = \frac{1}{2}q \log(1 - w) + \varrho \log(1 + \alpha w) + p \log(1 + w) + \varsigma \log(1 - \beta w) + \frac{1}{2}q \log(1 - w),$$

where q represents the probability $P(S \in I_1 \cup I_5) = P(R = -1)$, ϱ represents the probability $P(S \in I_2)$, p represents the probability $P(S \in I_3) = P(R = 1)$, ς represents the probability $P(S \in I_4)$, and expected values $\alpha = E(R | S \in I_2)$ and $\beta = -E(R | S \in I_4)$, while $p + q + \varrho + \varsigma = 1$. Then G is maximized by differentiating with respect to w and setting equal to zero,

$$\begin{aligned} \frac{\partial}{\partial w} G &= \frac{p}{1 + w} - \frac{q}{1 - w} + \frac{\alpha \varrho}{1 + \alpha w} - \frac{\beta \varsigma}{1 - \beta w} \\ \implies w^* &= \frac{-(\alpha q - \alpha p + p + q) + \sqrt{(\alpha q - \alpha p + p + q)^2 - 4\alpha(\alpha p + \alpha q - \alpha - p + q)}}{2\alpha}, \end{aligned}$$

if and only if $\alpha = -\beta \neq 0$ (reflective symmetry for expected values), by the quadratic formula.

4.3 Minimum Relative Entropy

For the purposes of GEPO, the risk of an option portfolio is defined here as the relative entropy of portfolio returns, with respect to the uniform distribution. In order to calculate the relative entropy, we first must calculate the Shannon entropy.

4.3.1 Shannon Entropy

As a quick review of information theory (Shannon, 1948) [87; 88], the Shannon entropy of a random variable represents the amount of randomness inherent to that variable. For a discrete random variable X with probability mass function $P(\cdot)$ that can take on possible values $x_1, \dots, x_n$, the Shannon entropy H is the average amount of information produced by X , defined as

$$H(X) = E(-\log P(X)) = - \sum_{k=1}^m P(x_k) \log P(x_k).$$

For n discrete random variables $X_1, \dots, X_n$ respectively having $m_1, \dots, m_n$ states, the joint entropy of $\mathbf{X} = (X_1, \dots, X_n)$ is given by

$$H(X_1, \dots, X_n) = - \sum_{k_1=1}^{m_1} \cdots \sum_{k_n=1}^{m_n} P_{\mathbf{X}}(x_{1k_1}, \dots, x_{nk_n}) \log P_{\mathbf{X}}(x_{1k_1}, \dots, x_{nk_n}).$$

This can be calculated empirically given a set of historical data by using the method introduced in REPO (Section 2.3.3). For the case of options returns, many strategies have more than just two (binary) states. Generally there can be up to four total states: positive discrete returns (+1), negative discrete returns (−1), continuous returns positively correlated to underlying price, and continuous returns negatively correlated to underlying price.

4.3.2 Kullback-Leibler Divergence

Kullback and Leibler (1951) [52; 54] introduced the Kullback-Leibler divergence which measures the directed divergence between two probability distributions. For discrete probability distributions P and Q , the Kullback-Leibler divergence between P and Q , also known as the relative entropy of P with respect to Q , is given by

$$\begin{aligned} D_{KL}(P \parallel Q) &= - \sum_{x \in \chi} P(x) \log \left(\frac{Q(x)}{P(x)} \right) \\ &= \sum_{x \in \chi} P(x) \log \left(\frac{P(x)}{Q(x)} \right). \end{aligned}$$

As shown in the previous Chapter 3, relative entropy qualifies as a convex risk measure based on the relative entropy principle. We once again use this quantity as the discriminatory risk measure for option portfolio optimization. For m total possible states, we will use the m -state discrete uniform distribution U_m as the reference distribution and measure from there the distance to the distribution of portfolio

returns. Thus, for Shannon entropy $H(\cdot)$, the risk of an option strategy portfolio R_Q is measured by the relative entropy of R_Q with respect to the uniform distribution U_m ,

$$D_{KL}(R_Q \parallel U_m) = \log(m) - H(R_Q).$$

Using the same combinatorial technique employed for REPO and DEPO in Chapters 2 and 3, the Shannon entropy of option portfolio returns R_Q can be estimated empirically via probability generating functions. For a collection of n discrete return assets over time period $j = 1, \dots, T$, let $\mathbf{r}_j = (r_{1j}, \dots, r_{nj})$ denote the cross-sectional n -dimensional vector of outcomes across one observational row of data, and let them be uniquely represented by the collection of $\mathbf{u}_k$'s such that $\mathbf{u}_k = \{\mathbf{r}_j \mid \mathbf{r}_j \neq \mathbf{u}_l, \text{ for some } j, \text{ and any } l \neq k\}$. For example, if the $\mathbf{r}_j$'s were

$$\{(1, -1, 1), (1, -1, 1), (1, 0, 0), (1, 0, 0), (-1, 1, 1)\},$$

then the $\mathbf{u}_k$'s would be $\{(1, -1, 1), (1, 0, 0), (-1, 1, 1)\}$. The empirical Shannon entropy of option portfolio returns R_Q can then be expressed as

$$H(R_Q) \approx - \sum_{k=1}^m \left(\frac{g^{(k)}(0)}{k!} \right) \log \left(\frac{g^{(k)}(0)}{k!} \right),$$

for k th-derivative at $x = 0$ of generating function

$$g(x; w_1, \dots, w_n) = \frac{1}{T} \sum_{j=1}^T x^{\{k: \mathbf{r}_j = \mathbf{u}_k\}}.$$

Therefore the risk of an option portfolio is given by the relative entropy of portfolio returns R_Q , estimated empirically as

$$D_{KL}(R_Q \parallel U_m) \approx \log(m) + \sum_{k=1}^m \left(\frac{g^{(k)}(0)}{k!} \right) \log \left(\frac{g^{(k)}(0)}{k!} \right),$$

for m -state discrete uniform distribution U_m .

4.4 Option Portfolio Selection Based on Growth Rate and Relative Entropy

4.4.1 Generalized Entropic Portfolio Optimization (GEPO)

The new generalized entropic portfolio optimization (GEPO) problem uses a multi-objective function that minimizes estimated relative entropy and maximizes expected growth rate. Using this optimization, investors can make portfolio selections based on a chosen risk tolerance. The highest risk portfolio solely maximizes the expected portfolio growth rate, equivalent to the Kelly criterion method. The lowest risk portfolio minimizes the portfolio relative entropy, the most diversified portfolio allocating capital to all n options equally. Somewhere in between lies a user's optimal portfolio of choice. For the case of option strategies the returns can be both discrete and continuous in nature. For example, credit spread returns can be either +100% for a success, -100% for a failure, or somewhere in between -100% and +100% on a

continuous scale for partial returns which we will denote by 0. Thus we would have discrete return outcomes u such that $u \in \{-1, 0, +1\}$. This leads to the generalized entropic portfolio optimization problem. Consider n potential option strategy contracts. Let λ be the number of probability states in the single-asset option strategy. For the generalized option strategy, returns can exhibit at most four general unique probability states: $+100\%$, -100% , some continuous return on a positively sloped leg (with mean α), and some continuous return on a negatively sloped leg (with mean β), therefore $1 < \lambda \leq 4$. Let the up-to-four probability states be represented respectively by $p_i, q_i, \varrho_i, \varsigma_i$ for event $i \in \{1, \dots, n\}$, and let w_i represent the percentage of portfolio funds to be allocated on option i , with the total allocation summing to $\omega = w_1 + \dots + w_n$. Let $\mathbf{r}_j = (r_{1j}, \dots, r_{nj})$ denote the cross-sectional n -dimensional vector of outcomes across one observational row of data. Over T data points this leads to m historical unique vectors $\mathbf{u}_k = \{\mathbf{r}_j \mid \mathbf{r}_j \neq \mathbf{u}_l, \text{ for some } j, \text{ and any } l \neq k\}$ for $k = 1, \dots, m$ such that each $\mathbf{u}_k$ is unique, with m bounded by either T or the maximum number of possible combinations λ^n , so $m = \min(T, \lambda^n)$. Basically the collection of $\mathbf{u}_k$'s is a unique representation of the $\mathbf{r}_j$'s with no duplicates. Let us also denote $\eta = \sum_{i=1}^n I(w_i) \leq n$ as the number of chosen options in the portfolio, where $I(w_i)$ is the indicator function for the event $\{w_i > 0\}$. Then the GEPO problem is defined as the following optimization program, generalized for different kinds of

option strategies,

$$\begin{aligned}
& \text{minimize} \quad D_{KL}(R_Q \parallel U_m) = \log_\lambda(m) + \sum_{k=1}^m \left(\frac{g^{(k)}(0)}{k!} \right) \log_\lambda \left(\frac{g^{(k)}(0)}{k!} \right) \\
& \text{maximize} \quad G(\omega) = \frac{1}{\eta} \sum_{i=1}^n I(w_i) p_i \log_\lambda(1 + \omega) + \frac{1}{\eta} \sum_{i=1}^n I(w_i) q_i \log_\lambda(1 - \omega) + \\
& \quad \frac{1}{\eta} \sum_{i=1}^n I(w_i) \varrho_i \log_\lambda(1 + \alpha\omega) + \frac{1}{\eta} \sum_{i=1}^n I(w_i) \varsigma_i \log_\lambda(1 - \beta\omega), \\
& \text{subject to} \quad \omega = w_1 + \dots + w_n \leq 1,
\end{aligned}$$

$$w_i \geq 0 \quad \forall i,$$

$$w_i = w_j = \eta^{-1}\omega \quad \forall \{(i, j) : I(w_i) = I(w_j) = 1\},$$

for the m -state uniform distribution U_m and k th-derivative at $x = 0$ of probability generating function

$$g(x; w_1, \dots, w_n) = \frac{1}{T} \sum_{j=1}^T x^{\{k : (I(w_1)r_{1j}, \dots, I(w_n)r_{nj}) = \mathbf{u}_k\}}.$$

The last constraint in the optimization problem stems from the fact that joint entropy measures randomness strictly based on the inclusion or exclusion of a random variable. The joint entropy value does not change upon changes to non-zero percentages of asset allocation, so any non-zero weight w_i contributes the corresponding marginal entropy from asset i , regardless of the magnitude of w_i . For this reason every asset included in the portfolio is assigned an equal weighting of $\eta^{-1}\omega$.

4.4.2 Risk-Adjusted Performance

Here we will utilize the risk-adjusted ratio for comparing growth rates of gambling portfolios introduced in Section 3.4.5, called the Growth Rate Over UNiform Divergence (GROUND) ratio. This ratio measures the expected growth rate of a portfolio, adjusted by its risk level—relative entropy with respect to the uniform distribution. Let U_m be the m -state discrete uniform distribution. Then for chosen portfolio R_a and minimum risk portfolio R_b existing in the m -state event space, the GROUND ratio Γ_m is defined as

$$\begin{aligned}\Gamma_m &= \frac{E(G_a(\omega_a) - G_b(\omega_b))}{D_{KL}(R_a \parallel U_m) - D_{KL}(R_b \parallel U_m)} \\ &= \frac{E(G_a(\omega_a) - G_b(\omega_b))}{H(R_b) - H(R_a)},\end{aligned}$$

where $G_a(\omega_a)$ is the growth rate of the chosen portfolio with weighting ω_a , $G_b(\omega_b)$ is the growth rate of the minimum risk portfolio with weighting ω_b , $D_{KL}(R_a \parallel U_m)$ is the relative entropy of the chosen portfolio with respect to U_m , $D_{KL}(R_b \parallel U_m)$ is the relative entropy of the minimum risk portfolio with respect to U_m , and $H(\cdot)$ is the Shannon entropy.

4.5 An Option Portfolio Selection Example with GEPO

4.5.1 Data

In this example, actual put and call option data is presented for 20 randomly selected equities from the S&P100 composite index using the Wharton Research Data Services (WRDS) from the Wharton School of the University of Pennsylvania, found at `wrds-www.wharton.upenn.edu`. Prices, volumes, expiration dates, and other essential option data is compiled from June 2012 to January 2018 (June 2012 is the earliest month available that contains data for all listed securities). Included in this data are the Greek parameters for options, which are described in great detail in the Equity Options Reference Manual from TMX Group Montreal Exchange guides and strategies documentation (2020) [97]. The *delta* of each option, from the Greek parameters for options, effectively represents an estimated probability of landing in-the-money at expiry, and this parameter serves as the estimated success rates for the portfolio optimization here. For the purposes of this chapter, we are only concerned with *deltas* closest to, but not less than 50%, to ensure every possible option yields a slight edge. Using these data archives, we are able to build historical weekly bull put spread and bear call spread options by selecting the buy-sell pairs that most closely resemble a 1:1 odds wager for each expiration date, creating 287 unique data

points for each equity. Weekly credit spread outcomes versus historical strike prices are recorded and computed as follows,

$$I(E_P > S_P) - I(E_P \leq B_P) + (E_P - M_P)/(S_P - M_P), \text{ for put spreads, and}$$

$$I(E_P > S_P) - I(E_P \leq B_P) + (E_P - M_P)/(B_P - M_P), \text{ for call spreads,}$$

where E_P is the expiration price for the underlying equity in question, S_P is the strike price for the sold option, B_P is the strike price for the bought option, and M_P is the midpoint of the two strike prices. The result is a value between +1 and -1 inclusive, on the continuous line, where positive 1 is awarded for an expiration price above the upper strike, -1 for an expiration price below the lower strike, and partial continuous returns between -1 and +1 for spreads expiring in the middle interval in between strike prices from the credit spread payoff function diagrams in Figures 4.3 and 4.4. Using this historical data we are able to empirically calculate the estimated relative entropy of each option. The selected credit spread options and their respective outcomes against strikes over June 2012 to January 2018 are presented below in Table 4.1.

Table 4.1: Mean outcome, average state probabilities and estimated relative entropy (in trits) of select equity credit spread options from July 2012 to January 2018.

Company Name	Symbol	Mean			
		Outcome	p -prob	q -prob	D_{KL}
Apple Inc.	AAPL	0.10124	52.3%	42.9%	0.226675
Accenture	ACN	0.146462	54.3%	38.1%	0.183754
American Intl. Group	AIG	0.121123	49.8%	38.2%	0.118573
Bank of America Corp	BAC	0.139333	49.6%	32.6%	0.071421
Biogen	BIIB	-0.048327	51.4%	46.1%	0.281129
Caterpillar Inc.	CAT	0.151175	53.7%	38.5%	0.180714
Capital One Fin. Corp	COF	0.013458	46.6%	45.3%	0.16507
Costco Wholesale Corp	COST	0.177373	51.8%	37.3%	0.135782
Cisco Systems	CSCO	0.340769	59.6%	25%	0.141729
Facebook Inc.	FB	0.128127	53.5%	42.3%	0.242447
Intl. Business Machines	IBM	0.148669	53.2%	38.7%	0.173439
Intel Corp	INTC	0.143774	51.7%	35.5%	0.115093
Johnson & Johnson	JNJ	0.290017	61.9%	32.6%	0.251597
JPMorgan Chase & Co.	JPM	0.223986	58%	36.4%	0.230925
MasterCard Inc.	MA	0.125993	53.3%	40.8%	0.209407
McDonald's Corp	MCD	0.160243	53%	37.7%	0.157863
3M Company	MMM	0.18277	55.7%	35.7%	0.17661
Merck & Co.	MRK	0.177165	54%	37.9%	0.17795
Microsoft	MSFT	0.170696	55.3%	38.5%	0.209972
Oracle Corp	ORCL	0.286192	60.1%	30.6%	0.191343

Over the following calendar year 2018, there are 52 weekly equity option expiration periods, and estimated success probabilities for each option are defined as

follows. For put options, the negative *delta* represents the probability of landing in-the-money for bought puts, and $(1 + \textit{delta})$ represents landing out-the-money for sold puts. For call options, the *delta* represents landing in-the-money for bought calls, and $(1 - \textit{delta})$ represents landing out-the-money for sold calls. The historical results summarized in Table 4.1 are used to evaluate the estimated relative entropy risk of each option, as well as the combined estimated relative entropy of the composite portfolio. The emulation here shows how GEPO performs against leading Kelly criterion methods for picking a portfolio of options for each week throughout 2018.

For illustrative purposes, let us examine this method applied to the second week, with option expiration dates January 12, 2018. Table 4.2 lists the details of the selected group of equity credit spreads built of the appropriate option pairs.

Table 4.2: Selected equity credit spreads on January 12, 2018, with their respective spread intervals, deltas and state projections.

Symbol	Spread Type	Spread Interval	Sell Delta	Buy Delta	p-proj	q-proj
AAPL	Put	[167.5, 170]	-0.496757	-0.405426	50.3%	40.5%
ACN	Call	[149, 150]	0.492771	0.447763	50.7%	44.8%
AIG	Call	[60, 61]	0.489573	0.345599	51%	34.6%
BAC	Put	[28.5, 29]	-0.497381	-0.401975	50.3%	40.2%
BIIB	Put	[317.5, 320]	-0.495617	-0.448496	50.4%	44.8%
CAT	Put	[149, 150]	-0.497975	-0.43738	50.2%	43.7%
COF	Call	[92.5, 93]	0.498336	0.465034	50.2%	46.5%
COST	Put	[182.5, 185]	-0.497485	-0.417452	50.3%	41.7%
CSCO	Put	[36.5, 37]	-0.496195	-0.374409	50.4%	37.4%
FB	Put	[170, 172.5]	-0.487722	-0.392952	51.2%	39.3%
IBM	Put	[150, 152.5]	-0.494561	-0.311505	50.5%	31.2%
INTC	Put	[44, 44.5]	-0.496887	-0.403415	50.3%	40.3%
JNJ	Call	[142, 143]	0.499681	0.408484	50%	40.8%
JPM	Call	[105, 106]	0.499166	0.435065	50.1%	43.5%
MA	Call	[146, 147]	0.494829	0.433672	50.5%	43.4%
MCD	Put	[170, 172.5]	-0.499101	-0.335054	50.1%	33.5%
MMM	Call	[242.5, 245]	0.495376	0.398127	50.5%	39.8%
MRK	Call	[55, 55.5]	0.48691	0.405147	51.3%	40.5%
MSFT	Call	[84.5, 85]	0.498834	0.451614	50.1%	45.2%
ORCL	Call	[50, 51]	0.490397	0.374001	51%	37.4%

GEPO determines which collection of credit spreads to select and what percentage of portfolio funds to allocate to each, in order to build the optimal risk-reward

credit spread portfolio.

Each potential portfolio has an expected growth rate, given the *delta* projections, and an estimated relative entropy with respect to the uniform distribution. The historical data contains $T = 287$ data points for each option, so the maximum joint entropy that can possibly be exhibited is $\log_3(T) = 5.1515$, i.e., uniform distribution with $m = 287$ possible probability states. Therefore, a portfolio R_Q has an estimated relative entropy of

$$\begin{aligned} D_{KL}(R_Q \parallel U_T) &= \log_3(T) - H(R_Q) \\ &= 5.1515 - H(R_Q), \end{aligned}$$

for joint entropy $H(R_Q)$.

4.5.2 Efficient Frontier and Portfolio Selection

In the portfolio selection problem, the efficient frontier refers to the set of the optimal portfolios that yield the greatest expected return for a defined level of risk, or equivalently the least risk for a defined level of expected return (the dual problem). The efficient frontier illustrates the risk-return trade-off for a given set of optimal portfolios. Here we show the analogous efficient frontier for portfolios with discrete returns, comparing the expected growth rates and risk levels (estimated relative entropy) of each efficient portfolio. For the same week of January 12, 2018, Figure

4.9 below plots the potential portfolios and their respective expected growth rates against their inherent risk profile, the estimated relative entropy with respect to the uniform distribution $\text{Uniform}(T)$ to historical joint outcomes of the portfolio.




Figure 4.9: Growth rate by relative entropy efficient frontier.

For the current season emulation, the Kelly criterion strategy chooses the portfolio R_K that maximizes the expected growth rate. For week 2 this leads to the top right-most data point $K = (4.5206, 0.0056)$, with estimated relative entropy of 4.5206 and expected growth rate of 0.0056. This portfolio consists of just one equity credit spread from Table 4.2: bull put spread on IBM, consisting of buying a put with strike 150 and selling a put with strike 152.5, betting on IBM to expire above 152.5 with a 50.5% probability of success, 31.2% probability of loss, and 18.3% of a partial return. According to our extended Kelly criterion conditions from Section 4.2.2, the optimal

bet size here is 12%, and thus the chosen portfolio for week 2 is 12% of portfolio funds on $[150, 152.5] < \text{IBM}$, shown below in Table 4.3. This strategy disregards any concept of risk associated with the expected portfolio growth rate of $G(\omega) = 0.0056$.

Table 4.3: The Kelly criterion portfolio of options with percent allocation for expiration week 2, January 12, 2018.

Symbol	Spread Type	Spread Interval	p-proj	q-proj	ϱ-proj	Kelly Allocation %
IBM	Put	[150, 152.5]	50.5%	31.2%	18.3%	12%

Alternatively, GEPO chooses the optimal portfolio based on the risk-reward trade-off. For each of the weeks 1 to 52, portfolio selection is performed according to the following GEPO problem with $n = 20$, $\alpha = -0.5$ and $T = 287$ (using

logarithm base 3 here since we are dealing with three outcome states: 1, 0, and -1),

$$\begin{aligned}
& \text{maximize } G(\omega) = \frac{1}{\eta} \sum_{i=1}^n I(w_i) p_i \log_3(1 + \omega) + \frac{1}{\eta} \sum_{i=1}^n I(w_i) q_i \log_3(1 - \omega) + \\
& \quad \frac{1}{\eta} \sum_{i=1}^n I(w_i) \varrho_i \log_3(1 + \alpha\omega), \\
& \text{subject to } D_{KL}(R_Q \| U_m) = \log_3(m) + \sum_{k=1}^m \left(\frac{g^{(k)}(0)}{k!} \right) \log_3 \left(\frac{g^{(k)}(0)}{k!} \right) \leq 2, \\
& \quad w_1 + \dots + w_n \leq 1, \\
& \quad w_i \geq 0 \ \forall i, \\
& \quad w_i = w_j = \eta^{-1}\omega \ \forall \{(i, j) : I(w_i) = I(w_j) = 1\},
\end{aligned}$$

for $\text{Uniform}(T)$ distribution as the target distribution, and for the k th-derivative at $x = 0$ of probability generating function with T data points,

$$g(x; w_1, \dots, w_n) = \frac{1}{T} \sum_{j=1}^T x^{\{k : (I(w_1)r_{1j}, \dots, I(w_n)r_{nj}) = \mathbf{u}_k\}}.$$

While the Kelly criterion places the entire wager on the option (or options) with the greatest expected growth rate, GEPO diversifies the portfolio by distributing the percent allocation across multiple options according to the appropriate risk profile. For week 2, GEPO selects data point $D = (1.5845, 0.0035)$, with estimated relative entropy of 1.5845 and expected growth rate of 0.0035. This corresponds to the optimal portfolio of the six credit spreads listed below in Table 4.4, with a total portfolio allocation of 9% (compared to the 12% allocated by the Kelly criterion).

Table 4.4: GEPO portfolio of options with percent allocation for expiration week 2, January 12, 2018.

Symbol	Spread Type	Spread Interval	p-proj	q-proj	ϱ-proj	GEPO Allocation %
IBM	Put	[150, 152.5]	50.5%	31.2%	18.3%	1.5%
AIG	Call	[60, 61]	51%	34.6%	14.1%	1.5%
MCD	Put	[170, 172.5]	50.1%	33.5%	16.4%	1.5%
ORCL	Call	[50, 51]	51%	37.4%	11.6%	1.5%
FB	Put	[170, 172.5]	51.2%	39.3%	9.5%	1.5%
MRK	Call	[55, 55.5]	51.3%	40.5%	8.2%	1.5%

Looking at the portfolio efficiency via the risk-adjusted GROUND ratio, the GEPO portfolio has a GROUND ratio of $\Gamma_m = (0.0035 - 0.0021)/(1.5845 - 0.05) = 0.213\%$, more than twice as efficient as the Kelly criterion portfolio at $\Gamma_m = (0.0056 - 0.0021)/(4.5206 - 0.05) = 0.078\%$.

The actual expiration prices that follow for week 2 are IBM at 163.14 (win), AIG at 60.97 (94% partial loss), MCD at 173.57 (win), ORCL at 49.51 (win), FB at 179.37 (win), and MRK at 58.66 (loss), for a total return of 3.2% in week 2. Therefore the Kelly criterion strategy experiences a gain of 12% of portfolio balance in week 2, while the competing GEPO strategy gains 3.2%.

4.5.3 Comparison to the Kelly Criterion Over Time

We demonstrate here the performance of GEPO versus the Kelly and Kelly variant strategies over the entire 2018 calendar year, executing option strategies at weekly intervals. Methods in the previous Section 4.5.2 are repeated week by week over the course of 52 weeks. The Kelly criterion strategy allocates the optimal investment size each week on the option (or options) that yield the greatest expected growth rate. Half Kelly is the same strategy but utilizes the fractional Kelly variant by allocating only half the Kelly criterion weighting on the same options. GEPO optimal risk strategy employs the GEPO algorithm each week to select the portfolio with the greatest growth rate subject to the main constraint that the portfolio has an estimated relative entropy of no greater than 2. Each strategy begins the year with \$10,000 and the total results are shown below in Figure 4.10.




Figure 4.10: Comparison of Credit Spread Strategies from January 2018 to January 2020.

With consistent, sustainable returns, GEPO ultimately outperforms both the Kelly and half Kelly methods over the 52 week period, and more than doubles the initial investment by the end of the year. As the Kelly strategies experiences gross variability with the see-saw pattern returns, the diversification strategy of GEPO holds strong and consistently returns profits month after month. The main purpose of GEPO is to mitigate risk of inaccurate predictions, and goal is well accomplished. In the end, GEPO finishes the year at a profit of \$10,062, more than 100% ROI, while the Kelly criterion gives up most gains and only retains \$2,645 (26.45%) profit, with half Kelly finishing up \$1,770 (17.7%).

5 Conclusions and Future Work

5.1 Conclusions

Presented here is a new entropy-based combinatorial approach to portfolio selection called return-entropy portfolio optimization (REPO) that addresses the five main practical concerns with MVPO: (i) optimal solutions assigning large allocation weights to high risk assets, (ii) disturbance of the assets' dependence structure, (iii) drastic variations in optimal solutions when adjusting inputs, (iv) accommodating non-normal or asymmetric returns, and (v) difficulty estimating a covariance matrix and expected returns. By using combinatorial generating functions, REPO greatly simplifies the portfolio entropy computation. REPO is robust, non-parametric, and indifferent to non-normality and asymmetry, making it an ideal approach to the portfolio selection problem. In addition to these practical improvements over MVPO, REPO significantly outperforms the mean-variance method with greater future portfolio returns, especially in the short-term.

Also presented here is a new entropy-based combinatorial approach to binary option portfolio selection called discrete entropic portfolio optimization (DEPO). DEPO introduces a robust method for evaluating risk of binary option portfolios and gambling portfolios alike, and gives the mathematical tools to make data-driven portfolio selection decisions to mitigate risk. DEPO is robust, non-parametric, and indifferent to non-normality, asymmetry and small sample size data, making it an ideal approach to the binary option portfolio selection problem. Compared to previous research in this space, DEPO is first to introduce the concept of managing the risk of binary options as an additional dimension to the optimization of binary option portfolios. We show how relative entropy qualifies as a convex risk measure and is therefore an ideal minimization objective for the discrete return portfolio selection problem. DEPO also adapts the Kelly criterion to a collection of binary options by extending the results to multiple wagers. By minimizing the relative entropy of portfolio returns, DEPO is able to balance risk and reward to obtain the optimal portfolio growth rate according to investor risk criteria. DEPO consistently outperforms leading Kelly criterion strategies choosing optimal portfolios of FOREX binary options. Applied to an NFL sportsbook portfolio, DEPO ultimately outperforms the industry standard quantitative methods for bet size allocation. Other possible applications of DEPO include optimizing portfolios of digital options and fixed-return

options, as well as other more alternative portfolios like sportsbooks with parlays, or fantasy sports teams. Even further, any contracts with deterministic outcomes, such as Arrow-Debreu securities or related contracts, could be evaluated by their relative entropy risk and potential expected growth rate. DEPO can usher in a new range of portfolio optimization applications that were previously unavailable with the traditional mean-variance optimization or Kelly criterion alone.

Lastly presented here is a new entropy-based combinatorial approach to option strategy portfolio selection called generalized entropic portfolio optimization (GEPO). GEPO is the most general method of the entropic portfolio optimizations introduced in our research series. We extend the notorious Kelly criterion to accommodate multiple assets and mixed returns, with direct application to option strategies. Using the convex risk measure relative entropy, GEPO presents a robust method for evaluating risk of option strategy portfolios and gives the mathematical tools to make data-driven portfolio selection decisions to mitigate risk. GEPO is robust, non-parametric, and indifferent to non-normality, asymmetry and small sample size data, making it an ideal approach to the option strategy portfolio selection problem. We show how GEPO comfortably outperforms leading Kelly criterion strategies in choosing optimal portfolios of equity credit spreads over 2018, both absolutely and in terms of risk-adjusted performance via the GROUND ratio. GEPO has a wide

range of applications including option strategies such as covered calls, married puts, credit spreads, straddles, strangles, butterfly spreads, iron condors, and more.

5.2 Future Work

Future work on this topic of research includes:

(1) Expansion of DEPO and GEPO to accommodate non-even odds wagers by the use of a uniformity adjustment, and instead of measuring the entropy of returns, we measure the entropy of expected payout.

(2) Develop a macro portfolio selection model that exists as a layer on top of REPO, DEPO and GEPO, that allocates weights for each portfolio to form a macro mixed-asset optimal portfolio.

Bibliography

- [1] Abbas, A., Cadenbach, A., and Salimi, E. A Kullback–Leibler view of maximum entropy and maximum log-probability methods. *Entropy*, 19:232–245, 2017.
- [2] Ahmadi-Javid, A. Entropic value-at-risk: A new coherent risk measure. *Journal of Optimization Theory and Applications*, 155:1105–1123, 2012.
- [3] Artzner, P., Delbaen, F., Eber, J., and Heath, D. Coherent measures of risk. *Mathematical Finance*, 9:203–228, 1999.
- [4] Asness, C., Frazzini, A., and Pedersen, L. Leverage aversion and risk parity. *Financial Analysts Journal*, 68(1):47–59, 2012.
- [5] Babaei, S., Sepehri, M., and Babaei, E. Multi-objective portfolio optimization considering the dependence structure of asset returns. *European Journal of Operational Research*, 244:525–539, 2015.
- [6] Baker, R. and McHale, I. Optimal betting under parameter uncertainty: Improving the Kelly criterion. *Decision Analysis*, 10:189–199, 2013.
- [7] Beaudry, N. and Renner, R. An intuitive proof of the data processing inequality. *Quantum Information and Computation*, 12:432–441, 2012.
- [8] Beirlant, J., Dudewicz, E., Györfi, L., and van der Meulen, E. Nonparametric entropy estimation: An overview. *International Journal of Mathematical and Statistical Sciences*, 6:17–39, 1997.
- [9] Benello, A., van Biema, M., and Carlisle, T. Concentrated investing: Strategies of the world’s greatest concentrated value investors. John Wiley and Sons: New York, USA, 2016.
- [10] Bera, A. and Park, S. Optimal portfolio diversification using maximum entropy principle. *Econometric Reviews*, 27:484–512, 2008.

- [11] Best, M. and Grauer, R. On the sensitivity of mean-variance-efficient portfolios to changes in asset means: Some analytical and computational results. *The Review of Financial Studies*, 4(2):315–342, 1991.
- [12] Black, F. and Litterman, R. Global asset allocation with equities, bonds, and currencies. *Goldman Sachs Fixed Income Research*, pages 1–40, 1991.
- [13] Black, F. and Litterman, R. Global portfolio optimization. *Financial Analysts Journal*, 48(5):28–43, 1992.
- [14] Browne, S. and Whitt, W. Portfolio choice and the Bayesian Kelly criterion. *Advances in Applied Probability*, 28(4):1145–1176, 1996.
- [15] Carmichael, B., Koumou, G., and Moran, K. Unifying portfolio diversification measures using Rao’s quadratic entropy. *CIRANO Technical Report*, 16:1–45, 2015.
- [16] Cheng, C. Improving the Markowitz model using the notion of entropy. <https://www.diva-portal.org/smash/get/diva2:304730/FULLTEXT01.pdf>, 2006.
- [17] Chopra, V. and Ziemba, W. The effect of errors in means, variances, and covariances on optimal portfolio choice. *Journal of Portfolio Management*, 19:6–11, 1993.
- [18] Chu, D., Wu, Y., and Swartz, T. Modified Kelly criterion. *Journal of Quantitative Analysis in Sports*, 14:22, 2018.
- [19] Constantinides, G., Jackwerth, J., and Savov, A. The puzzle of index option returns. *The Review of Asset Pricing Studies*, 3:229–257, 2013.
- [20] Corvalán, A. Well diversified efficient portfolios. *Working Papers Central Bank of Chile*, 336:1–10, 2005.
- [21] Cover, T. and Thomas, J. Elements of information theory. John Wiley and Sons: New York, NY, USA, 1991.
- [22] Dai, W. Mean-entropy models for uncertainty portfolio selection. Springer: Singapore, 2018.
- [23] Das, S. Data science: Theories, models, algorithms, and analytics. https://srdas.github.io/Papers/DSA_Book.pdf, 2016.

- [24] Davari-Ardakani, H., Aminnayeri, M., and Seifi, A. Multistage portfolio optimization with stocks and options. *International Transactions in Operations Research*, 23(3):593–622, 2016.
- [25] Driessen, J. and Maenhout, P. The world price of volatility and jump risk. *Journal of Banking & Finance*, 37:518–536, 2013.
- [26] Eling, M. Does the measure matter in the mutual fund industry? *Financial Analysts Journal*, 64:54–66, 2008.
- [27] Eraker, B. The performance of model based option trading strategies. *Review of Derivatives Research*, 16:1–23, 2007.
- [28] Fadugba, S. Performance measure of binomial model for pricing American and European options. *Applied and Computational Mathematics*, 3:18–30, 2014.
- [29] Faias, J. and Santa-Clara, P. Optimal option portfolio strategies: Deepening the puzzle of index option mispricing. *Journal of Financial and Quantitative Analysis*, 52(1):277–303, 2017.
- [30] Fatyanova, M. and Semenov, M. Model for constructing an options portfolio with a certain payoff function. *Proceedings of 3rd International Conference of Information Technology and Nanotechnology*, 3:254–262, 2017.
- [31] Follmer, H. and Schied, A. Stochastic finance: An introduction in discrete time. De Gruyter: Berlin, Germany, 2011.
- [32] Fono, L., Kandem, J., and Tassak, C. Kurtosis and semi-kurtosis for portfolio selection with fuzzy returns. *International Statistical Institute: Proceedings of 58th World Statistical Congress*, pages 6517–6522, 2011.
- [33] Garner, W. and McGill, W. The relation between information and variance analyses. *Psychometrika*, 21(3):219–228, 1956.
- [34] Geman, D., Geman, H., and Taleb, N. Tail risk constraints and maximum entropy. *Entropy*, 17(6):3724–3737, 2015.
- [35] Green, R. and Hollifield, B. When will mean variance efficient portfolios be well diversified? *The Journal of Finance*, 47:1785–1809, 1992.

- [36] Hardy, M. An introduction to risk measures for actuarial applications. https://www.researchgate.net/profile/Mary_Hardy2/publication/242469445_An_Introduction_to_Risk_Measures_for_Actuarial_Applications/links/0a85e5342bd491d188000000/An-Introduction-to-Risk-Measures-for-Actuarial-Applications.pdf, 2006.
- [37] Haugh, M. and Kogan, L. Duality theory and approximate dynamic programming for pricing American options and portfolio optimization. *Handbooks in Operations Research and Management Science*, 15:925–948, 2007.
- [38] Huang, X. Portfolio selection with fuzzy returns. *Journal of Intelligent and Fuzzy Systems*, 18:383–390, 2007.
- [39] Huang, X. Mean-semivariance models for fuzzy portfolio selection. *Journal of Computational and Applied Mathematics*, 217:1–9, 2008.
- [40] Huang, X. An entropy method for diversified fuzzy portfolio selection. *International Journal of Fuzzy Systems*, 14(1):161–165, 2012.
- [41] Hubacek, O., Sourek, G., and Zelezny, F. Exploiting sports-betting market using machine learning. *International Journal of Forecasting*, 35(2):783–796, 2019.
- [42] Jensen, J. Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta Mathematica*, 30(1):175–193, 1906.
- [43] Jiang, L., Wu, K., and Zhou, G. Asymmetry in stock comovements: An entropy approach. *Journal of Financial and Quantitative Analysis*, 53:1479–1507, 2018.
- [44] Jondeau, E. and Rockinger, M. Conditional asset allocation under non-normality: How costly is the mean-variance criterion? *Institute of Banking and Finance, HEC Lausanne*, pages 1–42, 2005.
- [45] Jones, C. A nonlinear factor analysis of S&P 500 index option returns. *Journal of Finance*, 61:2325–2363, 2006.
- [46] Jorion, P. Portfolio optimization in practice. *Financial Analysts Journal*, 48:68–74, 1992.
- [47] Karandikar, R. and Sinha, T. Modelling in the spirit of Markowitz portfolio theory in a non-Gaussian world. *Current Science*, 103:666–672, 2012.

- [48] Kelly, J. A new interpretation of information rate. *Bell System Technical Journal*, 35(4):917–926, 1956.
- [49] Keren, G. and Charles, L. The two fallacies of gamblers: Type I and type II. *Organizational Behavior and Human Decision Processes*, 60(1):75–89, 1994.
- [50] Khanna, N. The Kelly criterion and the stock market. https://sites.math.washington.edu/~morrow/336_16/2016papers/nikhil.pdf, 2016.
- [51] Koumou, G. and Dionne, G. Coherent diversification measures in portfolio theory. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3351423, 2019.
- [52] Kullback, S. Information theory and statistics. John Wiley and Sons: New York, USA, 1959.
- [53] Kullback, S. Information theory and statistics. Dover Publications: NY, USA, 1996.
- [54] Kullback, S. and Leibler, R. On information and sufficiency. *Annals of Mathematical Statistics*, 22(1):79–86, 1951.
- [55] Lassance, N. and Vrins, F. Minimum Rényi entropy portfolios. *Annals of Operations Research*, pages 1–37, 2019.
- [56] Lavinio, S. The hedge fund handbook. McGraw-Hill, Irwin Library of Investment & Finance: New York, USA, 2000.
- [57] Learned-Miller, E. and Fisher, J. ICA using spacings estimates of entropy. *Journal of Machine Learning Research*, 4:1271–1295, 2003.
- [58] Lintner, J. Securities prices, risk, and maximal gains from diversification. *The Journal of Finance*, 20(4):587–615, 1965.
- [59] Lintner, J. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *The Review of Economics and Statistics*, 47:13–37, 1965.
- [60] Liu, J. and Pan, J. Dynamic derivative strategies. *Journal of Financial Economics*, 69:401–430, 2003.

- [61] Low, R., Faff, R., and Kjersti, A. Enhancing mean–variance portfolio selection by modeling distributional asymmetries. *Journal of Economics and Business*, 85:49–72, 2016.
- [62] MacKay, D. Information theory, inference and learning algorithms. Cambridge University Press: United Kingdom, 2003.
- [63] MacLean, L. The Kelly capital growth investment criterion: Theory and practice. World Scientific: Singapore, 2010.
- [64] Maclean, L. and Ziemba, W. The handbook of the fundamentals of financial decision making. World Scientific: Singapore, 2013.
- [65] Markowitz, H. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952.
- [66] McGill, W. Multivariate information transmission. *Psychometrika*, 19(2):97–116, 1954.
- [67] Michaud, R. The Markowitz optimization enigma: Is optimized optimal? *Financial Analysts Journal*, 45:31–42, 1989.
- [68] Miller, G. Note on the bias of information estimates. *Information Theory in Psychology: Problems and Methods*, pages 95–100, 1955.
- [69] Mossin, J. Equilibrium in a capital asset market. *Econometrica*, 34(4):768–783, 1966.
- [70] NADEX. What are binary options and how do they work? <https://www.nadex.com/learning-center/courses/binary-options/components-binary-option-pricing>, 2020.
- [71] O’Shaughnessy, D. Optimal exchange betting strategy for win-draw-loss markets. *MATHSPORT 2012: 11th Australasian Conference on Mathematics and Computers in Sport*, 1:1, 2012.
- [72] Paninski, L. Estimation of entropy and mutual information. *Neural Computation*, 15:1191–1253, 2003.
- [73] Parzen, E. On estimation of a probability density function and mode. *Annals of Mathematical Statistics*, 33:1065–1076, 1962.
- [74] Philippatos, G. and Wilson, C. Entropy, market risk, and the selection of efficient portfolios. *Applied Economics*, 4:209–220, 1972.

- [75] Pola, G. On entropy and portfolio diversification. *Journal of Asset Management*, 17:218–228, 2016.
- [76] Rad, H., Low, R., and Faff, R. The profitability of pairs trading strategies: Distance, cointegration and Copula methods. *Quantitative Finance*, 16:1541–1558, 2016.
- [77] Rao, C. Diversity and dissimilarity coefficients: A unified approach. *Theoretical Population Biology*, 21:24–43, 1982.
- [78] Rao, C. Diversity: Its measurement, decomposition, apportionment and analysis. *Sankhyā: The Indian Journal of Statistics, Series A*, 44:1–22, 1982.
- [79] Rao, C. Rao’s axiomatization of diversity measures. John Wiley and Sons: New York, USA, 2004.
- [80] Rao, C. Quadratic entropy and analysis of diversity. *Sankhyā: The Indian Journal of Statistics, Series A*, 72:70–80, 2010.
- [81] Rao, C. and Nayak, T. Cross entropy, dissimilarity measures, and characterizations of quadratic entropy. *IEEE Transactions on Information Theory*, 31:589–593, 1985.
- [82] Rom, B. and Ferguson, K. Post-modern portfolio theory comes of age. *The Journal of Investing*, 3(3):11–17, 1993.
- [83] Rosenblatt, M. Remarks on some nonparametric estimates of a density function. *Annals of Mathematical Statistics*, 27:832–837, 1956.
- [84] Rotela Jr., P., Rocha, L., Aquila, G., Balestrassi, P., Peruchi, R., and Lacerda, L. Entropic data envelopment analysis: A diversification approach for portfolio optimization. *Entropy*, 19(9):352, 2017.
- [85] Rotundo, L. and Thorp, E. The Kelly criterion and the stock market. *The American Mathematical Monthly*, 99:922–931, 1992.
- [86] Rényi, A. On measures of information and entropy. *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, 4:547–561, 1960.
- [87] Shannon, C. A mathematical theory of communication: Part 1. *Bell Systems Technical Journal*, 27(3):379–423, 1948.

- [88] Shannon, C. A mathematical theory of communication: Part 2. *Bell Systems Technical Journal*, 27(4):623–656, 1948.
- [89] Sharpe, W. Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance*, 19(3):425–442, 1964.
- [90] Sharpe, W. Mutual fund performance. *The Journal of Business*, 39(1):119–138, 1966.
- [91] Silverman, B. Density estimation for statistics and data analysis. Chapman and Hall: London, UK, 1998.
- [92] Sinclair, E. Confidence intervals for the Kelly criterion. *Journal of Investment Strategies*, 3:65–74, 2014.
- [93] Sortino, F. and Price, L. Performance measurement in a downside risk framework. *The Journal of Investing*, 3(3):59–64, 1994.
- [94] Sumarti, N. and Nadya, P. A dynamic portfolio of American option using fuzzy binomial method. *Journal of Information Technology Education*, 3:85–92, 2016.
- [95] Sun, R., Ma, T., Liu, S., and Sathye, M. Improved covariance matrix estimation for portfolio risk measurement: A review. *Journal of Risk and Financial Management*, 12:48–80, 2019.
- [96] T., P. and Pot, V. Portfolio analysis using stochastic dominance, relative entropy, and empirical likelihood. *Management Science*, 63:153–165, 2017.
- [97] TMX Group, L. Montréal exchange: Guides and strategies. https://m-x.ca/educ_guides_strat_en.php, 2020.
- [98] Tobin, J. Liquidity preference as behavior towards risk. *Review of Economic Studies*, 25(2):65–86, 1958.
- [99] Urbanowicz, K. Entropy and optimization of portfolios. <http://arxiv.org/pdf/1409.7002v1.pdf>, 2014.
- [100] Usta, I. and Kantar, Y. Portfolio optimization with entropy measure. https://www.researchgate.net/publication/261859605_Portfolio_optimization_with_entropy_measure, 2007.
- [101] Usta, I. and Kantar, Y. Mean-variance-skewness-entropy measures: A multi-objective approach to portfolio selection. *Entropy*, 13:117–133, 2011.

- [102] Vecer, J. Optimal distributional trading gain: State price density equilibrium and Bayesian statistics. Working paper. <https://ssrn.com/abstract=3616661>, 2020.
- [103] Wong, W., Leung, P., and Ng, H. An improved estimation to make Markowitz’s portfolio optimization theory users friendly and estimation accurate with application on the US stock market investment. *European Journal of Operational Research*, 222:85–95, 2012.
- [104] Xu, Y., Wu, Z., Jiang, L., and Song, X. A maximum entropy method for a robust portfolio problem. *Entropy*, 16(6):3401–3415, 2014.
- [105] Yu, J., Lee, W., and Chiou, W. Diversified portfolios with different entropy measures. *Applied Mathematics and Computation*, 241:47–63, 2014.
- [106] Zeng, Y. and Klabjan, D. Portfolio optimization for American options. *Journal of Computational Finance*, 22:37–64, 2018.
- [107] Zhao, L. and Palomar, D. A Markowitz portfolio approach to options trading. *IEEE Transactions on Signal Processing*, 66:4223–4238, 2018.
- [108] Zhou, R., Liu, X., Yu, M., and Huang, K. Properties of risk measures of generalized entropy in portfolio selection. *Entropy*, 19(12):657, 2017.
- [109] Zhou, R., Zhan, Y., Cai, R., and Tong, G. A mean-variance hybrid-entropy model for portfolio selection with fuzzy returns. *Entropy*, 17:3319–3331, 2015.

A Appendix

A.1 Proposition 1

Let the number of securities be n and the number of probability states be m . Then $\sum_{i=0}^{n-1} m^i$ individual and conditional entropies are required for computing the joint entropy.

Proof A.1 *The joint entropy of a portfolio with n securities can be expressed as a sum of their conditional entropies, defined as:*

$$H(R_1, \dots, R_n) = \sum_{i=1}^n H(R_i \mid R_{i-1}, \dots, R_1)$$

We have the following claim: For n securities and m probability states $A_1, \dots, A_m$, with $n, m \in \mathbb{N}$, $H(R_1, \dots, R_n)$ has $\sum_{i=0}^{n-1} m^i$ terms. We show it by induction on n . It is straightforward to show that the claim holds true for $n = 2$. Now we assume

that the claim is true for k securities $H(R_1, \dots, R_k)$ that have $\sum_{i=0}^{k-1} m^i$ terms,

$$\begin{aligned}
H(R_1, \dots, R_k, R_{k+1}) &= \sum_{i=1}^{k+1} H(R_i \mid R_{i-1}, \dots, R_1) \\
&= \sum_{i=1}^k H(R_i \mid R_{i-1}, \dots, R_1) + H(R_{k+1} \mid R_k, \dots, R_1) \\
&= H(R_1, \dots, R_k) + H(R_{k+1} \mid R_k, \dots, R_1) \\
&= H(R_1, \dots, R_k) + H(R_{k+1} \mid R_k = A_1, \dots, R_1 = A_1) + \dots \\
&\quad + H(R_{k+1} \mid R_k = A_m, \dots, R_1 = A_m) \\
&\equiv \sum_{i=0}^{k-1} m^i \text{ terms} + m^k \text{ terms} \\
&= \sum_{i=0}^{k+1-1} m^i \text{ terms.}
\end{aligned}$$

Thus, the claim holds true for any n .

A.2 Proposition 2

Let X and Y be two random variables that are not necessarily independent, and let a and b be scalar weights such that $0 \leq a, b \in \mathbb{R}$. Then, the joint entropy $H(aX, bY)$ is always greater than or equal to the individual entropy $H(aX + bY)$.

Proof A.2

$$P(aX + bY = ax + by) \geq P(aX = ax, bY = by)$$

$$\log P(ax + by) \geq \log P(ax, by)$$

$$P(ax + by) \log P(ax + by) \geq P(ax + by) \log P(ax, by)$$

$$P(ax + by) \log P(ax + by) \geq P(ax, by) \log P(ax, by)$$

$$\begin{aligned} \sum_{i=1}^n \sum_{j=1}^m P(ax_i + by_j) \log P(ax_i + by_j) &\geq \sum_{i=1}^n \sum_{j=1}^m P(ax_i, by_j) \log P(ax_i, by_j) \\ - \sum_{i=1}^n \sum_{j=1}^m P(ax_i + by_j) \log P(ax_i + by_j) &\leq - \sum_{i=1}^n \sum_{j=1}^m P(ax_i, by_j) \log P(ax_i, by_j) \\ \implies H(aX + bY) &\leq H(aX, bY). \end{aligned}$$

This property can be easily extended to n variables. While the level of uniformity increases with less mutual information and decreases with greater mutual information, $H(aX, bY)$ will certainly always have an entropy greater than or equal to $H(aX + bY)$. This is basically a variation of the data processing inequality that states that no clever manipulation of the data can improve inference. See Cover (1991) [21] and Beaudry (2012) [7]. Consider a probability model described by the Markov Chain: $X \rightarrow Y \rightarrow Z$, where $X \perp Z | Y$. It follows that $I(X; Y) \geq I(X; Z)$, so no transformation of the received code Y can give more information about the sent code X than Y itself. In reference to the example provided, $aX + bY$ is the so-called clever manipulation of the data, and it cannot exceed the information contained in the pair (X, Y) .