

Article

Dynamic Asset Allocation with Expected Shortfall via Quantum Annealing

Hanjing Xu ¹, Samudra Dasgupta ^{2,3,4}, Alex Pothen ¹ and Arnab Banerjee ^{2,3,*} ¹ Department of Computer Science, Purdue University, West Lafayette, IN 47906, USA; xu675@purdue.edu² Department of Physics, Purdue University, West Lafayette, IN 47906, USA³ Oak Ridge National Laboratory, Quantum Computing Institute, Oak Ridge, TN 37831, USA⁴ Bredeben Center, University of Tennessee, Knoxville, TN 37996, USA

* Correspondence: arnabb@purdue.edu

Abstract: Recent advances in quantum hardware offer new approaches to solve various optimization problems that can be computationally expensive when classical algorithms are employed. We propose a hybrid quantum-classical algorithm to solve a dynamic asset allocation problem where a target return and a target risk metric (expected shortfall) are specified. We propose an iterative algorithm that treats the target return as a constraint in a Markowitz portfolio optimization model, and dynamically adjusts the target return to satisfy the targeted expected shortfall. The Markowitz optimization is formulated as a Quadratic Unconstrained Binary Optimization (QUBO) problem. The use of the expected shortfall risk metric enables the modeling of extreme market events. We compare the results from D-Wave's 2000Q and Advantage quantum annealers using real-world financial data. Both quantum annealers are able to generate portfolios with more than 80% of the return of the classical optimal solutions, while satisfying the expected shortfall. We observe that experiments on assets with higher correlations tend to perform better, which may help to design practical quantum applications in the near term.

Keywords: portfolio optimization problem; Quadratic Unconstrained Binary Optimization (QUBO); quantum annealing; hybrid algorithm



Citation: Xu, H.; Dasgupta, S.; Pothen, A.; Banerjee, A. Dynamic Asset Allocation with Expected Shortfall via Quantum Annealing. *Entropy* **2023**, *25*, 541. <https://doi.org/10.3390/e25030541>

Academic Editors: Giuliano Benenti and Brian R. La Cour

Received: 2 February 2023

Revised: 1 March 2023

Accepted: 17 March 2023

Published: 21 March 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

We describe a hybrid quantum-classical algorithm to solve a dynamic asset allocation problem where the targeted return and expected-shortfall (ES)-based risk appetite are specified. Since both the return as well as the shortfall are functions of the chosen asset allocation, we treat the return as a constraint in a modified Markowitz framework, and optimize the allocation strategy to meet the requirements of the expected shortfall using an iterative procedure that solves the Markowitz Optimization problem at each iteration. The latter optimization problem is solved by a Quadratic Unconstrained Binary Optimization (QUBO) formulation on a quantum annealer, while the iterative procedure to compute the shortfall is performed by a classical algorithm.

Quantum annealing offers a highly parallelized approach for solving optimization problems by using quantum tunneling from a manifold of high-energy solutions to the ground state. A common approach to embed the optimization problem into an Ising quantum annealer is to convert it to a QUBO problem [1–4]. Several examples have been explored so far in the literature, including the maximum clique [5], scheduling [6] and graph coloring problems [7], among others [8].

The portfolio optimization problem, introduced by Harry Markowitz [9] in 1952, investigates how investors could use the power of diversification to optimize portfolios by minimizing risk, and serves as a foundation for later models, such as the Black–Litterman model [10]. The original Markowitz Optimization problem used volatility as the measure of the risk. However, it is now known that volatility changes with time [11]; hence, treating

it as a constant is risky and sub-optimal. Furthermore, it often fails to characterize the market during extreme events or “shocks”, for example, the 2008 mortgage crisis which led to an abrupt collapse of the market with the insolvency of Lehman Brothers. As a result, modern finance practitioners prefer to use a time-varying risk metric such as stochastic volatility, Value-at-Risk (VaR) or the expected shortfall. The latter is defined as the average loss that can be expected when the loss has already exceeded a specific threshold [12]. The advantages of expected shortfall over other risk measurements such as volatility or Value-at-Risk are discussed in [11].

It is NP-Hard to solve general quadratic optimization problems [13]. For convex quadratic optimization problems such as the portfolio optimization problem, however, there exist polynomial time algorithms that takes $O(n^{7/2}L)$ time [14], where n is the number of variables and L bounds the number of digits for each integer. Hence it is prohibitive to solve large-scale portfolio optimization problems exactly using classical methods due to the high time complexity. Hence as more versatile and scalable quantum computing devices—currently quantum annealers—enter the market, we explore solving the portfolio optimization problem on two such machines available today using QUBO formulations.

In Grant et al. [15], the authors have benchmarked the performance of a D-Wave 2000Q quantum annealer on solving the Markowitz Optimization Problem with a relatively small size of 20 logical variables and random data. Our study has the following novel contributions:

- We demonstrate how optimization problems with non-polynomial constraints such as the Expected Shortfall could be solved with a hybrid quantum-classical, iterative approach that requires no additional qubits. An alternative approach would encode such constraints directly into a QUBO by converting them first to a multilinear polynomial through Fourier analysis [16], and then to a quadratic polynomial using methods described in [17–19]. However, in this approach, in the worst-case the number of binary variables will grow exponentially due to non-trivial higher-order terms generated from the Fourier expansion, which severely limits the problem sizes that we can solve on the current generation of quantum hardware.
- To the best of our knowledge, quantum computing has not been employed prior to this study for solving Expected-Shortfall based dynamic asset-allocation problems [12]. Previous approaches (e.g., [15]) have employed the classical Mean-variance framework. However, static variance is no longer used in modern finance as it is well known that volatility fluctuates with time and hence it needs to be modeled in a statistical framework that captures non-stationarity. Moreover, industrial practitioners prefer tail-risk measures such as Value at Risk and Expected Shortfall (the latter is considered cutting edge in risk management) since true risk is associated with the fluctuations in the negative return, and is not symmetric with respect to positive and negative returns (i.e., no one minds a surprise positive return).
- Thirdly, this is one of the first papers that uses quantum computing for portfolio optimization using real financial data (using ETF and currency data) on a real quantum computer (i.e., not simulation) in an accurate industry setting. Previous approaches have used random data (e.g., [15]).

We further explored our algorithm’s performance on two generations of quantum annealers offered by D-Wave, with up to 115 logical variables. We provide experimental results on both the Advantage (Pegasus topology) and 2000Q (Chimera topology) D-Wave quantum annealers. The results are generally close to the optimal portfolios obtained by classical optimization methods, in terms of final returns and Sharpe ratios (return/standard deviation of the return in a time period).

The paper is structured as follows. Section 2 defines the expected shortfall based dynamic asset allocation problem and lays out a hybrid algorithm for solving it. Section 3 provides the technical background on D-Wave’s quantum annealer and maps the Mean-Variance Markowitz problem on it. Section 4 discusses the experimental results on both

D-Wave 2000Q and Advantage systems. Section 5 states our conclusions and lists future research directions.

2. The Problem of Dynamic Asset Allocation

The problem of dynamic asset allocation is to allocate/invest an amount of money into N assets, while satisfying an expected return and keeping the risk below a given threshold. To make the problem more specific, we need to describe the input data and variables.

- The historical return matrix R is obtained from Yahoo Finance [20–25] for the assets mentioned in Section 4.2 with N rows and T_{total} columns, where N is the number of assets and T_{total} is the number of days data are collected. We divide the return matrix R into periods of T days and index the data for each time period, for example, R_t represents the return data from t -th time period.
- The vector of asset means μ_t is computed from R_t .
- The co-variance matrix C_t calculated from the matrix R_t as

$$C_{t,i,j} = \frac{(e_i^T R_t - \mu_{t,i} \mathbf{1}) \cdot (e_j^T R_t - \mu_{t,j} \mathbf{1})}{T - 1}, \tag{1}$$

where e_j is the column vector of all zeros except with a one at the j -th position.

Asset allocation is especially interesting to financial practitioners during a time period with unpredictable market turbulence with goal of minimizing risk while achieving a target return. The risk is upper bounded by a consumer-driven risk appetite. A data sheet of the assets' daily returns (profit one can earn if buying an asset the previous day and selling the next) for the previous three months is available. The risk threshold can be set using observed market metrics from a volatile time period, for example, the 2008 market crash. The algorithm uses the assets' historical return data to estimate the trends of the assets' performance and their correlations. Options for risk measurements include:

- Volatility: the standard deviation of the portfolio return.
- Value-at-Risk at level α : the smallest number y such that the probability that a portfolio does not lose more than $y\%$ of total budget is at least $1 - \alpha$.
- Expected Shortfall at level α : the expected return from the worst $\alpha\%$ cases. It is defined as follows:

$$ES_\alpha(w_t, R_t) = \text{mean}(\text{lowest } \alpha\% \text{ from } w_t^T R_t). \tag{2}$$

We focus on the expected shortfall as our risk measurement for the rest of the paper as it is the modern approach preferred by practitioners (as mentioned earlier in Section 1). The problem can be expressed as follows where the weight vector w indicates what fraction of the budget is invested in each asset:

(P1) Minimize the expected shortfall $ES_\alpha(w_t, R_t)$ under the constraints that the expected return is satisfied, the variance of the portfolio is small, and all assets are invested.

It is possible to write the expected shortfall based portfolio optimization as a linear program [26], but it requires adding $N + 1$ variables and $2N$ constraints where N is the number of assets. Since the expected shortfall cannot be expressed by a quadratic formulation natively, we opt to use it as a convergence criterion instead of including it in the optimization problem directly. To justify this approach, assuming that the assets' historical returns follow a Gaussian distribution, we can approximate the expected shortfall of a given portfolio P by:

$$ES_\alpha(P) = \mu + \sigma \frac{\phi(\Phi^{-1}(\alpha))}{1 - \alpha}, \tag{3}$$

where μ is the expected return, σ is the volatility of the portfolio, and $\phi(x)$ and $\Phi(x)$ are the Gaussian probability distribution and cumulative distribution functions, respectively [27].

The expected shortfall has positive correlation with the volatility and in turn, the variance of the portfolio ([11,28]).

Hence we propose a bilevel optimization approach described in Figure 1 to solve Problem (P1).

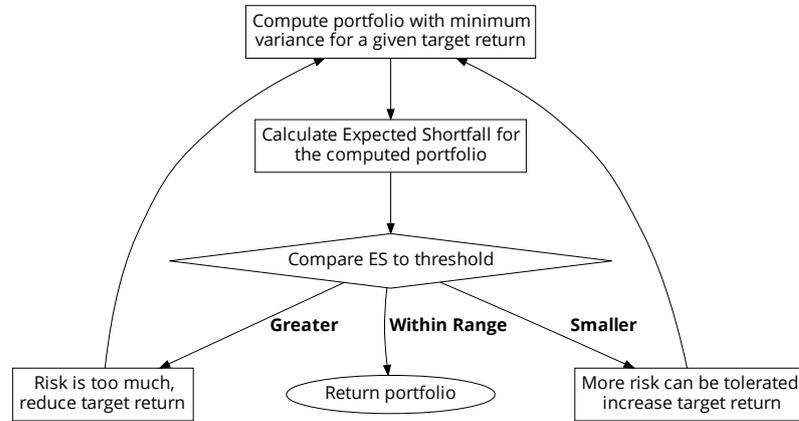


Figure 1. A flowchart for the proposed algorithm for computing optimal portfolio with a threshold on the expected shortfall.

Given a balance sheet of the assets’ returns in the history, we create multiple time periods each with T days. After picking target return p_t for one of the time periods t , we choose a reference asset that is representative of the portfolio, and set a target expected shortfall for that asset computed from its volatility in the year 2008, its volatility in the time period t , and its shortfall in 2008. (The precise expression is included in Algorithm 1). Then we use the Markowitz Optimization problem [9] to allocate assets within the portfolio in order to minimize volatility with the constraint that the target return is met. Next we compute the expected shortfall from the current allocation of assets. If the target expected shortfall is not met by the current allocation, then we adjust the target return value and iteratively solve the Markowitz Optimization problem. We terminate when the target expected shortfall is met, or the target return cannot be met as the maximum return of all assets is smaller than the target return.

Now we describe the Markowitz portfolio optimization procedure. Its QUBO formulation will be provided in the next section. The Markowitz Optimization problem can be expressed by the quadratic optimization problem

$$\begin{aligned}
 \min_w \quad & w_t^T C_t w_t \\
 \text{s.t.} \quad & \mu_t^T w_t = p_t, \\
 & \sum_i w_{t,i} = 1, w_{t,i} \geq 0 \quad \forall i.
 \end{aligned} \tag{4}$$

where p_t is the target portfolio return during t -th time period. The constraint $\mu_t^T w_t = p_t$ ensures that the target return is met, $\sum w_{t,i} = 1$ indicates we want to invest all of the resources, and $w_{t,i} \geq 0$ means short selling is not allowed. With all the constraints satisfied, we minimize $w_t^T C_t w_t$, that is, the variance of the portfolio at $t - th$ time period. However, with our bilevel optimization, we treat the constraints as soft constraints, that is, small violations of their values are permitted. The optimizer can return portfolios with small variance even when the expected return falls short of the target; if the sum of weight is not equal to 1, we can scale the weights of the assets to sum to 1.

Algorithm 1: Expected Shortfall based Dynamic Asset Allocation during t .

```

Input:  $\mu_t, R_t, \sigma_{ref}, \sigma_{ref,t}, ES_{ref}, \alpha$ 
Output:  $w_t$ 
Let  $p_t = \text{mean}(\mu_t)$ ; ▷ Initialize the target return
Let  $EST_t = \frac{\sigma_{ref}}{\sigma_{ref,t}} ES_{ref}$ ; ▷ Initialize the target expected shortfall
Let  $C_t = \text{cov}(R_t)$ ; ▷ Compute the co-variance matrix from the returns
while True do
  if  $p > \max(\mu_t)$  then
    | Return 0; ▷ Return constraint cannot be satisfied
  Solve Equation (4) for  $w_t$ ;
  Let  $ES_t = ES_\alpha(w_t, R_t)$ 
  if  $\frac{|ES_t|}{|EST_t|} > 1 + \epsilon$  then
    | Let  $p_t = p_t(1 - \delta)$ ; ▷ Decrease target return for lower risks
  else if  $\frac{|ES_t|}{|EST_t|} < 1 - \epsilon$  then
    | Let  $p_t = p_t(1 + \delta)$ ; ▷ Increase target return as there is room for more risks
  else
    | Return  $w_t$ ;
end

```

Algorithm 1 provides the pseudo-code for the algorithm for expected shortfall based asset allocation. Here σ_{ref} is the volatility of a reference asset’s returns during the market crash in 2008; the reference asset is chosen from among the assets to be representative of the market trend, for example, SPDR S&P 500 ETF Trust (SPY). The variable $\sigma_{ref,t}$ is the volatility of the reference asset’s returns during the time window t ; ES_{ref} is reference asset’s expected shortfall during the market crash; EST_t is the target expected shortfall at time window t ; α is the risk level parameter; ES_t is the expected shortfall for the computed portfolio during the optimization process at time window t ; ϵ is the error tolerance parameter; and δ is the momentum parameter that is adjusted dynamically. Figure 2 shows the ratio between the variance and expected shortfall in different iterations of Algorithm 1 from ETFs consisting of 6 assets whose returns were obtained from December 2019 to May 2020. The monotonic one-to-one tracking justifies why optimization problems with expected shortfall constraint can be solved iteratively using the Markowitz Mean-Variance framework.

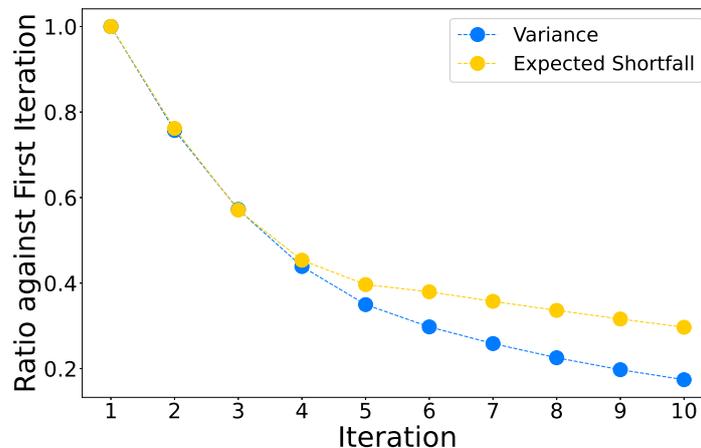


Figure 2. The y-axis tracks the ratio between the variance and expected shortfall with $\alpha = 5$ in later iterations against their respective values in the first iteration of Algorithm 1 running on 6 ETF assets. The expected shortfall decreases at a different rate from the variance but each iteration in the algorithm is guaranteed to make progress towards the target expected shortfall, which ensures convergence.

3. A Hybrid Quantum Classical Algorithm

3.1. Algorithm Overview

We will use a hybrid quantum classical algorithm to solve the quadratic optimization problem given in Equation (4) with a quantum annealer backend.

Quantum annealing (QA) [29–31] is the quantum analog of the classical annealing where the disorder is introduced quantum mechanically instead of thermally via applying the Pauli matrix x on every qubit as in Equation (5).

$$\mathcal{H}_I = - \sum_{i=1}^N \sigma_i^x, \tag{5}$$

This Hamiltonian does not commute with the problem Hamiltonian in Equation (6)

$$\mathcal{H}_P = - \sum_i h_i \sigma_i^z - \sum_{i<j} J_{ij} \sigma_i^z \sigma_j^z, \tag{6}$$

where σ_i^z is the Pauli matrix z acting on qubit i , h_i is the magnetic field on qubit i and J_{ij} defines the coupling strength between qubits i and j [32]. The spin configurations of the ground states of Equation (6) also minimize the Ising model problem:

$$\min_s E(s) = - \sum_i h_i s_i - \sum_{i,j} J_{ij} s_i s_j, \quad s_i \in \{-1, 1\}, \tag{7}$$

where s_i is the spin, h is the external longitudinal magnetic field strength vector and the matrix J represents the coupler interactions. Moreover, the general two-dimensional Ising problem within a magnetic field is NP-hard [33]. And in the case of spin-glass three-dimensional Ising model with lattice size of $N = lmn$, the complexity is $O(2^{mn})$ [34], which is NP-Hard as well.

During the QA process, combining both Hamiltonians in Equations (5) and (6), at time t the system evolves under the following Hamiltonian:

$$\mathcal{H}(t) = A\left(\frac{t}{T}\right)\mathcal{H}_I + B\left(\frac{t}{T}\right)\mathcal{H}_P. \tag{8}$$

Here T is the total annealing time and the system is initialized to the ground state of \mathcal{H}_I , which is a superposition of all qubits in the z basis. Functions $A\left(\frac{t}{T}\right)$ and $B\left(\frac{t}{T}\right)$ describe the change of influences from disorder and problem Hamiltonians on the system. \mathcal{H}_I dominates \mathcal{H}_P initially and slowly (adiabatically) changes to the opposite while the influence of \mathcal{H}_I vanishes at the end of the annealing process, thus removing disorder from the system. The system will then settle into one of the low energy states.

Due to unavoidable experimental compromises [35], QA serves as an intermediate step towards universal adiabatic quantum computation (AQC) [36,37] as the system evolves under a time-dependent Hamiltonian

$$\mathcal{H} = [1 - s(t)]\mathcal{H}_I + s(t)\mathcal{H}_P, \tag{9}$$

where $s(t)$ changes from 0 to 1. When conditions on internal energy gap and time scales are met [38], the system will remain in its ground state at all times, which is different from QA.

A Quadratic Unconstrained Binary Optimization (QUBO) problem of the form

$$\min_x Q(x) = \sum_i h_i x_i + \sum_{i,j} J_{ij} x_i x_j, \quad x_i \in \{0, 1\} \tag{10}$$

aims to minimize a mathematical function with linear and quadratic terms; here any combination of $x_i \in \{0, 1\}, \forall i$ is feasible. It can be converted to the Ising model shown in Equation (7) by a one-to-one mapping of the variables: $x_i = \frac{1+s_i}{2}$. We will use the QUBO

formulation for the rest of the paper but note that quantum annealers from D-Wave require the QUBO problems to be transformed into Ising models before execution.

Consider a standard binary optimization problem with a linear or quadratic objective function $f(x)$ and linear constraints $Ax = b$, where $A \in \mathbb{R}^{m \times n}$, and $b \in \mathbb{R}^{m \times 1}$.

$$\begin{aligned} & \min_x f(x) \\ & \text{s.t. } Ax = b, \\ & \text{where } x \in \{0, 1\}^{n \times 1}. \end{aligned} \tag{11}$$

We can rewrite it as a QUBO

$$Q(s) = f(x) + \lambda(Ax - b)^T(Ax - b) \tag{12}$$

to be minimized by quantum annealers with a large enough $\lambda \in \mathbb{R}^+$ to guarantee that the constraint is satisfied in the optimal solutions.

We will now discuss how to convert the Markowitz Optimization problem with continuous variables in Equation (4) to a QUBO problem.

First we write Equation (4) as an unconstrained optimization problem with penalty coefficients λ_1 and λ_2 (the subscripts t are dropped for better readability):

$$Q = \left(\sum_i^n \sum_j^n C_{i,j} w_i w_j \right) + \lambda_1 \left(\sum_i^n \mu_i w_i - p \right)^2 + \lambda_2 \left(\sum_i^n w_i - 1 \right)^2, \tag{13}$$

where λ_1 and λ_2 scale the constraint penalties. Minimizing Equation (13) is equivalent to

$$\begin{aligned} \min Q = & \left(\sum_i^n \sum_j^n C_{i,j} w_i w_j \right) + \lambda_1 \left[\left(\sum_i^n \mu_i w_i \right)^2 - 2p \sum_i^n \mu_i w_i \right] \\ & + \lambda_2 \left[\left(\sum_i^n w_i \right)^2 - 2 \sum_i^n w_i \right], \end{aligned} \tag{14}$$

after expanding the squared terms and eliminating the constants. When the constraints are satisfied exactly, we have

$$\lambda_1 \left[\left(\sum_i^n \mu_i w_i \right)^2 - 2p \sum_i^n \mu_i w_i \right] = -\lambda_1 p^2, \tag{15}$$

and

$$\lambda_2 \left[\left(\sum_i^n w_i \right)^2 - 2 \sum_i^n w_i \right] = -\lambda_2. \tag{16}$$

We use k binary variables $x_{i,1} \dots x_{i,k} \in \{0, 1\}$ to approximate each continuous variable w_i in Equation (4) with a finite geometric series

$$w_i = \sum_{a=1}^k 2^{-a} x_{i,a}. \tag{17}$$

The larger k is, the more precision w_i has. However, larger k also widens the differences between the coupler strengths— J terms in Equation (10). Although the coupler strengths for D-Wave annealers can be set at any double-precision floating point number between -1 and 1 , precision errors may pose a challenge due to integrated control errors (ICE) [39]. In our experiments, we set $k = 5$ which empirically gives us the best approximations to the optimal solutions for Equation (4). For larger k we risk the errors dominating the

coupler coefficients, rendering those additional qubits unreliable. We set λ_1 to p^{-2} and λ_2 to 1 to bound the penalty terms in Equation (14) to -2 . Additionally, we scale the objective by a factor of λ_3 to around 1 such that penalty terms in the optimal solutions to Equation (14) remain relatively small while not overwhelming the objective. If penalties dominate the objective, it may introduce numerous local minima to the energy landscape and the optimizer will suffer from barren plateaus. Alternatively, if the objective dominates penalties, constraints will be violated significantly. The soft constraints enable us to obtain better portfolios as presented in Section 4.4.

Substituting Equation (17) in to Equation (14), we have the final binary optimization formalism

$$f(x) = \left[\left(\sum_i^n \sum_a^k \mu_i 2^{-a} x_{i,a} \right)^2 - 2p \sum_i^n \sum_a^k \mu_i 2^{-a} x_{i,a} \right] + p^{-2} \left[\left(\sum_i^n \sum_a^k 2^{-a} x_{i,a} \right)^2 - 2 \sum_i^n \sum_a^k 2^{-a} x_{i,a} \right] + \lambda_3 \sum_i^n \sum_j^n \sum_a^k \sum_b^k C_{i,j} 2^{-a-b} x_{i,a} x_{j,b}, \quad (18)$$

which is quantum-annealable as it only has linear and quadratic interactions.

3.2. Previous Work

Rosenberg et al. [40] solve the multi-period portfolio optimization problem using D-Wave’s quantum annealer:

$$\begin{aligned} \max_w \quad & \sum_{t=1}^T (\mu_t^T w_t - \frac{\gamma}{2} w_t^T \Sigma_t w_t - \Delta w_t^T \Lambda_t \Delta w_t \\ & + \Delta w_t^T \Lambda'_t \Delta w_t) \\ \text{s.t.} \quad & \sum_{n=1}^N w_{nt} = K, \forall t, \quad w_{nt} \leq K', \forall t, \forall n. \end{aligned} \quad (19)$$

Here T is the number of time steps, and N is the number of assets. At each time step t , μ_t represents the forecast returns, w_t are holdings for each asset, Σ_t is the forecast covariance matrix, Λ_t and Λ'_t are coefficients for transaction costs related to temporary and permanent market impacts, respectively, which penalize changes in the holdings if the corresponding terms are positive. Additionally, γ is the risk aversion factor.

Equation (19) seeks to maximize returns considering constraints on asset size. Specifically, the sum of asset holdings is constrained by K and the maximum allowed holdings of each asset is K' . For small problems ranging from 12 to 584 variables, D-Wave’s 512 and 1152-qubit systems are able to find optimal solutions with high probability.

Venturelli and Kondratyev [41] focus on the following QUBO problem where the task is to select M assets from a pool of N assets:

$$\min_q \quad \sum_{i=1}^N a_i q_i + \sum_{i=1}^N \sum_{j=i+1}^N b_{ij} q_i q_j + P \left(M - \sum_{i=1}^N q_i \right). \quad (20)$$

The variable q_i is 1 if asset i is selected and 0 otherwise. The coefficient a_i indicates the attractiveness of the i -th asset and b_{ij} is the pairwise diversification penalties (positive) or rewards (negative). The penalty coefficient P scales the constraint on the number of selected assets to make sure it is satisfied in the optimal solution. The authors have explored

the benefits of reverse annealing on D-Wave systems, and report one to three orders of magnitude speed-up in time-to-solution with reverse annealing.

The problem considered by Phillipson and Bhatia [42] is similar to the Markowitz Optimization problem but with binary variables indicating asset selections instead of real weights. The authors report comparable results from D-Wave's hybrid solver to other state of the art classical algorithms and solvers including simulated annealing [43,44], genetic algorithm [45,46], linear optimization problems [47] and local search [48].

Grant et al. [15] benchmark the Markowitz Optimization problem on a D-Wave 2000Q processor with real weight variables and price data generated uniformly at random, and explore how embeddings, spin reversal and reverse annealing affect the success probability. Hegade et al. [49] solve the same problem with added counterdiabatic terms on circuit-based quantum computers and see improvements on success probabilities using digitized-adiabatic quantum computing (DAdQC) and Quantum Approximation Optimization Algorithm (QAOA) [49].

We extend the general QUBO formulation in [15] to solve the asset allocation problem, with the expected shortfall as the risk metric, using the Markowitz Optimization problem as a subroutine at each iterative step of the algorithm. We use our algorithms on real-world ETF and currency data. Additionally, we present the results on the newly-available Advantage processor and experiment on problems with up to 115 logical variables, up from 20 in [15].

4. Experimental Setup and Results

4.1. D-Wave Quantum Annealer

We start by discussing the latest quantum annealing technologies offered by D-Wave as the solvability of the problem is dependent on the architecture. D-Wave quantum annealers are specifically designed to solve Ising problems natively. Currently two types of quantum annealers are offered by D-Wave: 2000Q processor with Chimera topology and the Advantage processor with Pegasus topology. The latter was made publicly available in 2020 and it has more qubits (5760 vs. 2048) and better connectivity than the former. The qubits in the Chimera topology have 5 couplers per qubit while in the Pegasus topology they have 15 couplers per qubit [50]. It is not always possible to formulate an optimization problem to match the Chimera or Pegasus topologies exactly. Therefore minor embeddings are necessary to map the problems to D-Wave processors. Such embeddings usually require the users to map multiple physical qubits to one logical variable with constraints such that every qubit on the 'chain' behaves the same, which significantly reduces the total size of the problems that can be solved on the quantum annealers.

Furthermore, it is advisable to have uniform chain lengths (number of qubits representing a single variable) for more predictable chain dynamics during the anneal [51]. Algorithms in [52] detail such procedures for fully-connected graphs which is the underlying logical graph for the portfolio optimization problem. A full-yield 2000Q processor can map up to 64 logical variables and an Advantage processor can map around 180 logical variables. A comparison between the embedding of the two topologies is shown in Figure 3. In our experiments, we use the *find_clique_embedding* function from *dwave-system* to map fully-connected graphs to either the Chimera or the Pegasus topology.

4.2. Test Input and Annealer Parameters

We pick the top-six ETFs by trading volumes, EEM, QQQ, SPY, SLV, SQQQ and XLF, and six major currencies' USD exchange rates, AUD, EUR, GBP, CNY, INR and JPY, for most of the tests below. The reference assets for ETF and currency tests are SPY and EUR, respectively. For the tests in Section 4.5 we use 12 and 23 assets respectively and pick the top ETFs by trading volumes again. We choose the parameter α in the definition of expected shortfall to be 5%.

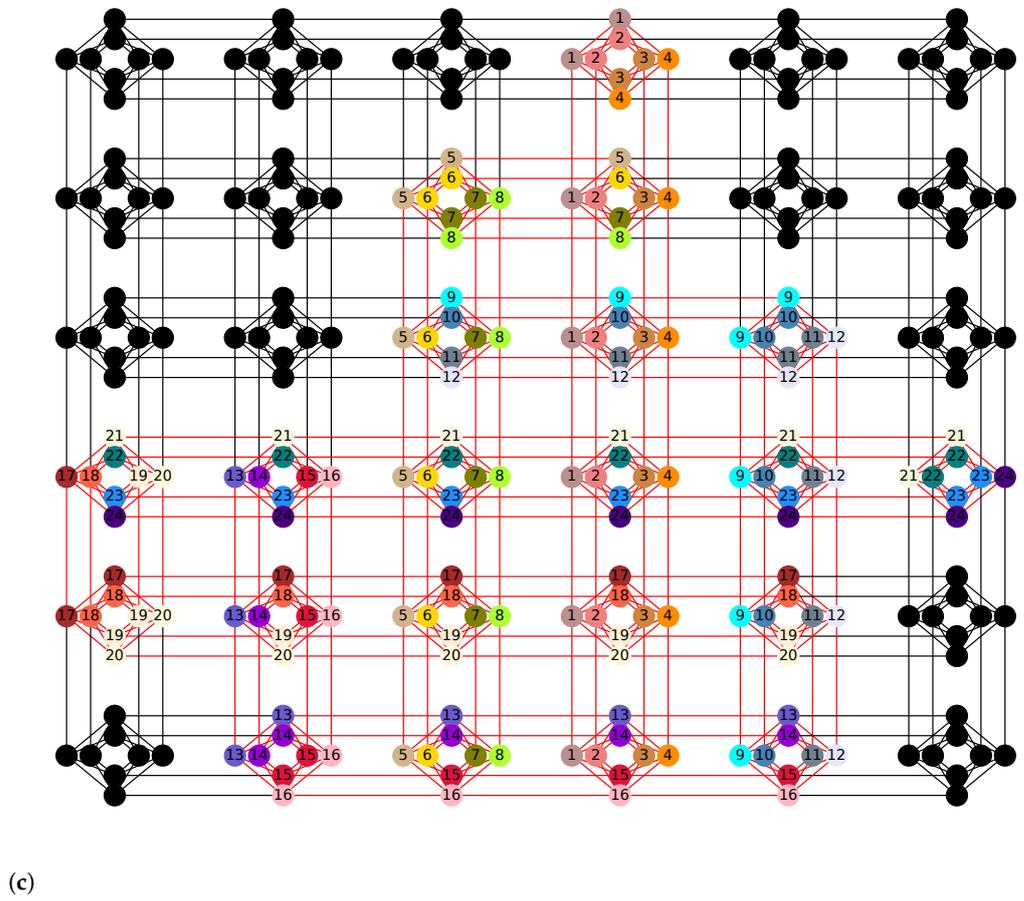
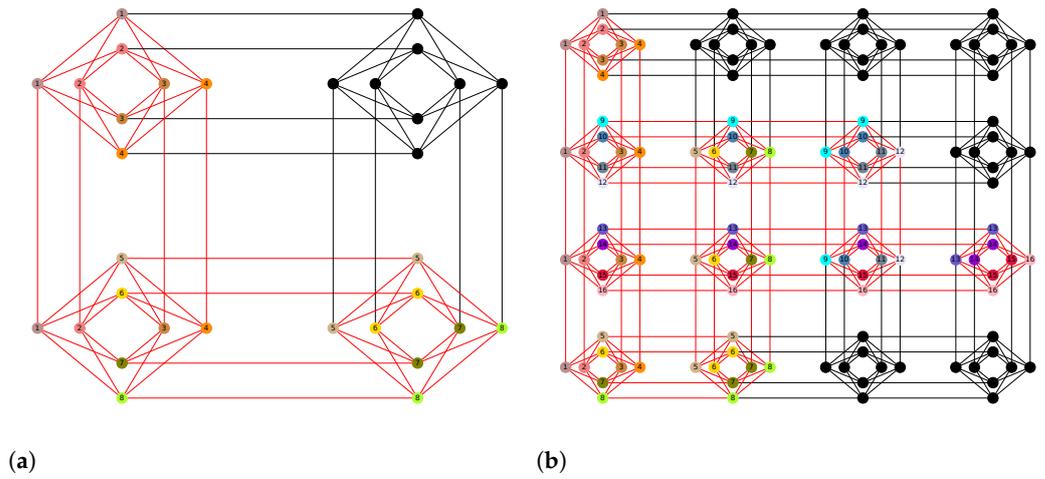


Figure 3. Cont.

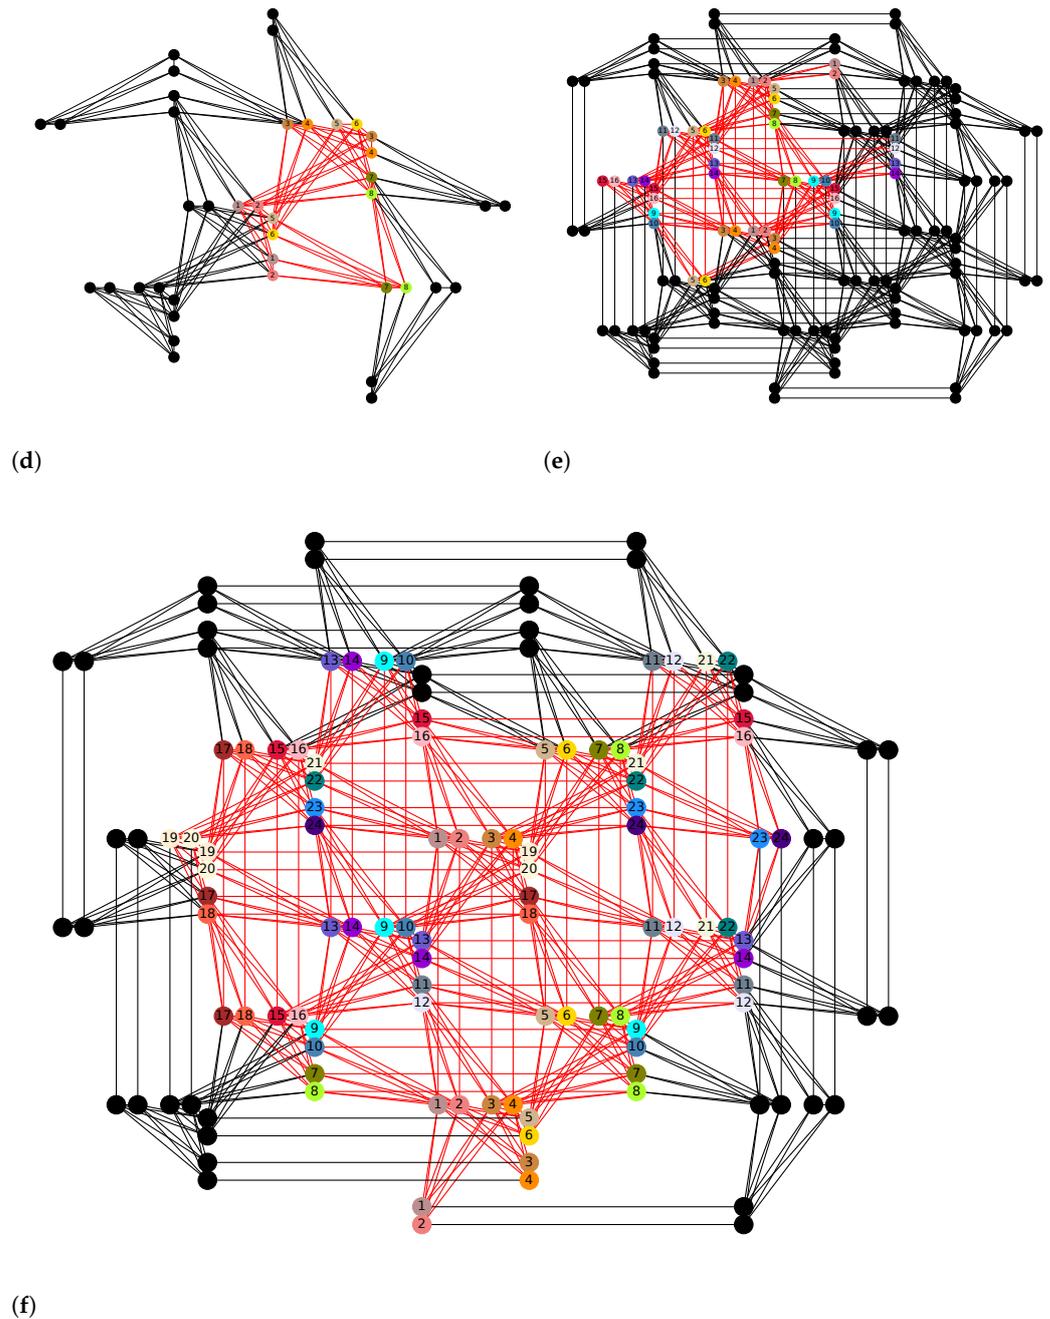


Figure 3. A comparison between minor embeddings on Chimera (2000Q) and Pegasus (Advantage) lattices of D-Wave processors for cliques (fully connected graphs) of different sizes. The vertices with the same color or label represent the same logical variable in Equation (7) and the chain length is defined as the number of qubits used to represent one logical variable. Each Chimera cell is a 4 by 4 complete bipartite graph ($K_{4,4}$) with 4 additional edges connecting neighboring cells. Each Pegasus cell has 24 qubits which include three $K_{4,4}$ graphs as in the Chimera cell and the cells are connected with each other using $K_{2,4}$ edges. To minor embed cliques of 8 vertices ($K = 8$), the chain length on the Chimera lattice is 3 while on the Pegasus lattice it is 2. With $K = 16$, the chain lengths are 5 and 2–3, respectively, and with $K = 24$, they are 7 and 3–4, respectively. This shows that Pegasus processor scales better for larger clique problems, which may lead to better performance. (a) Embedding $K_{8,8}$ on Chimera topology; (b) Embedding $K_{16,16}$ on Chimera topology; (c) Embedding $K_{24,24}$ on Chimera topology; (d) Embedding $K_{8,8}$ on Pegasus topology; (e) Embedding $K_{16,16}$ on Pegasus topology; (f) Embedding $K_{24,24}$ on Pegasus topology.

We can control a range of annealer parameters that may impact the solution quality in various degrees. Specifically, we set the number of spin reversal transforms [53] to 100 and readout thermalization to 100 μs as suggested in [54,55]. The spin reversal transform flips the signs of 100 variables and coefficients of the Ising model, which leaves the ground state invariant. The goal is to average out the system errors thus improving the quality of the solutions [53]. 100 μs readout thermalization allows the system enough time to cool back to the base temperature after each anneal. We set the annealing time at 1 μs as longer annealing time sees no statistically significant improvements to the solutions similarly reported in [15]. Results from 2000Q and Advantage processors are both included in the following sections. Additionally, we report the results from D-Wave’s post processing utility on 2000Q processors, which decomposes the underlying graph induced by the QUBO into several low tree-width subgraphs [56], and then solves them exactly using belief propagation on junction trees [57].

We sample all QUBOs 30,000 times with both D-Wave backends and report the samples with the lowest objective value from Equation (17) each time. Figure 4 shows an example distribution of the samples.

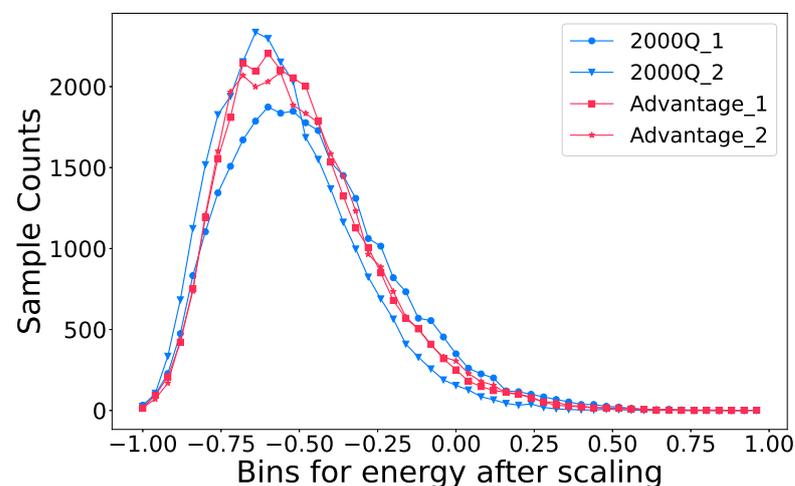


Figure 4. The distribution of the samples for 4 different QUBOs with both D-Wave backends. Each QUBO is sampled 30,000 times and the objective values of the samples is scaled to be between $(-1, 1)$. We divide the objective value range into 50 equally-spaced bins and count the number of samples in each bin. All four samples exhibit the Poisson distribution, and thus we only report the samples with the lowest objective value for the experiments in this section since they can be reproduced reliably.

4.3. Embedding Comparison on D-Wave Annealers

As discussed in Section 4.1, D-Wave quantum annealers require the problems to be minor embedded to the Chimera or Pegasus topology. For small problems this means there may be multiple valid embeddings and in this section we will measure how different embeddings can make an impact on the solution quality.

We compute four different embeddings that use different sets of physical qubits from both 2000Q and Advantage processors. Otherwise, the embedding graphs are the same, and hence they use the same number of qubits and chain lengths. We sample the same QUBO—first iteration of Algorithm 1 on the ETFs from December 19 to May 20—10 times with 10,000 samples each. We then pick the best solutions in terms of QUBO objective value from all 10 sample sets for each embedding and obtain their average and minimum values. Tables 1 and 2 report the results as ratios against the best objective values computed by simulated annealing for better readability. Since the objective values are negative, we compute ratios of the magnitudes instead.

Table 1. Embedding comparison on the 2000Q processor with 30 logical variables or 270 physical qubits after minor embedding. The objective value is calculated from Equation (18) and is normalized against the simulated annealer solving the same QUBO. All energies computed are negative, and their respective magnitudes are used for the comparison. The second embedding out of these four is able to find the solution with the better average and best objective value.

Embedding	Average Objective	Best Objective
1	95.66%	98.78%
2	96.83%	99.66%
3	96.53%	98.37%
4	96.21%	98.49%

Table 2. Embedding comparison on the Advantage processor with 30 logical variables or 134 physical qubits after minor embedding. Different embeddings on the Advantage processor show no statistically significant differences.

Embedding	Average Objective	Best Objective
1	99.25%	99.89%
2	99.52%	99.94%
3	99.22%	99.95%
4	99.48%	99.89%

We can see from Tables 1 and 2 that the impact that different embeddings make is statistically insignificant. However, it is clear that the Advantage processors have higher ratios than the 2000Q processors, which we will address next.

4.4. Annealing Results Comparison

We benchmark our algorithm on both simulated and quantum annealers using, as the baseline algorithm, a classical optimization solver, namely, *cvxpy* [58]. We create five ETF test datasets and four currency test datasets from 100 days of return data with different starting dates from 2010 to 2020.

The results are normalized against the optimal classical solution. The quantum algorithm fails to converge for the first two currency tests on the 2000Q processor, and the corresponding bars are missing in Figures 5 and 6.

In Figures 5 and 6, we used $k = 5$ binary variables to represent each asset weight. The simulated annealing results follow the optimal solutions closely in most tests. We note that in tests 2 and 5 from the ETF tests and tests 1, 2 and 3 from Currency tests, simulated annealing, and in some cases, quantum annealing produce portfolios of higher returns than those of the exact classical quadratic optimization problem solver. This is due to how Markowitz Optimization problems are formulated as QUBOs with discretized variables in Equation (18), which changes the optimization problem slightly, and also the optimal asset allocations. In test 4 from Currency tests, quantum annealers are able to find a portfolio with higher returns than simulated annealing as it returns a portfolio with a slightly increased risk that is still acceptable, but higher returns. This is not optimal in terms of QUBO objective values as the constraint penalty is now higher, yet the solution is still feasible. We also observe that the currency tests generally perform better than ETF tests on both quantum annealer backends. Figure 7 shows how quantum annealers perform with respect to the average of absolute correlation coefficients over all pairs of assets in each test. Higher correlation coefficients seem to lead to higher returns.

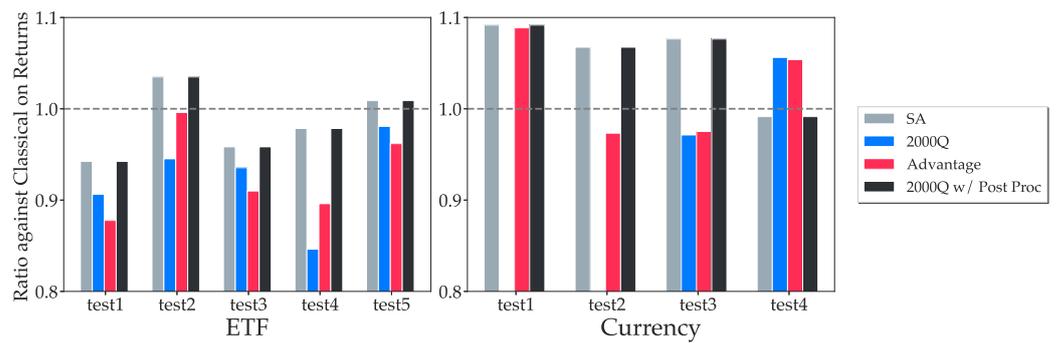


Figure 5. The comparison of the final returns between all four backends. A higher ratio means the backend can return portfolios with higher returns. Each test uses 100 days of return data with different starting dates from 2010 to 2020. The results from 2000Q with post processing yields identical results from simulated annealing. Both 2000Q and Advantage processors are able to compute returns that are consistently more than 80% of the optimal, except the two currency test cases where the algorithm fails to converge on the 2000Q.

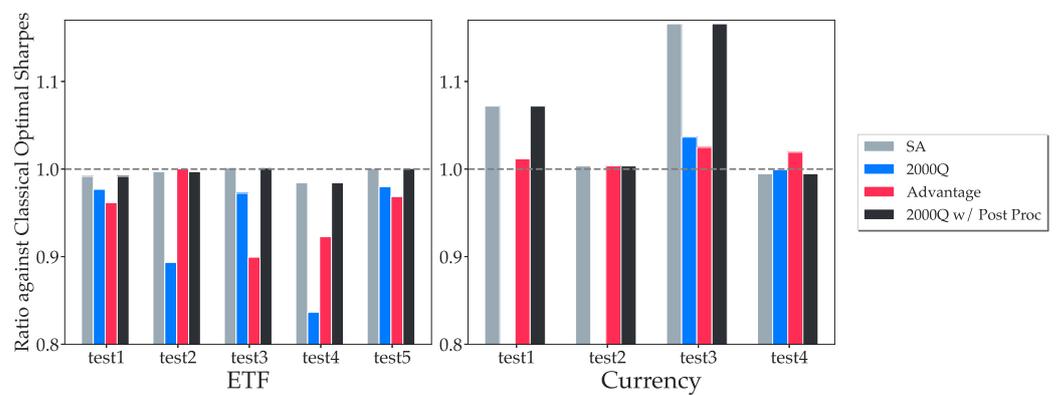


Figure 6. The comparison of the final Sharpe ratios between all four backends. Recall that the Sharpe ratio is the ratio of the return to the standard deviation of an asset for a set time period. Given a portfolio defined by the weight vector w , the Sharpe ratio of this portfolio is calculated as $\frac{\mu^T w}{\sqrt{w^T C w}}$. A higher ratio means the backend can return portfolios with higher Sharpe ratios. The results confirm that the portfolio variances returned by the quantum processors are close to the optimal results obtained from classical optimization methods, and it is effective to solve standard constrained optimization problems as a QUBO.

Although we acknowledge there may be other factors contributing to our observations that currency tests do better than ETF tests on for quantum annealers, Figure 7 implies that more correlated assets tend to perform better. Detailed analysis on which attributes of the assets have an impact on the quantum annealing performance and how much the impacts are requires more research in the future. Ref. [59] used machine learning models such as decision tree and regression to predict the accuracy of D-Wave’s quantum annealer on maximum clique problems.

4.5. State-of-the-Art on D-Wave Annealers

The embeddings of the six asset tests on both 2000Q and Advantage processors leave plenty of unused qubits. D-Wave’s clique embedding algorithm [52] suggests that we can embed fully connected graphs with 64 and 180 vertices to full-yield 2000Q and Advantage processors, respectively. Due to the defective qubits and connectors in the currently available Advantage processor, experimentally, we can embed only up to 119 qubits. This means we can solve portfolio optimization problems with 12 and 23 assets natively on 2000Q and Advantage processors, respectively.

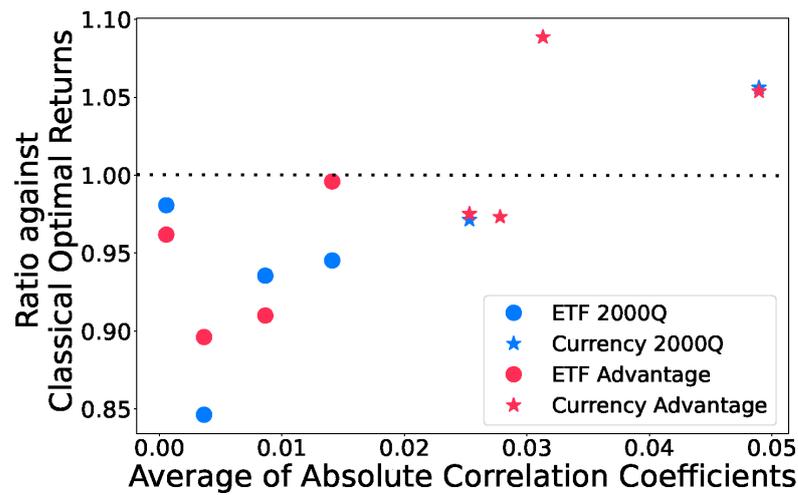


Figure 7. Final returns obtained from both quantum annealers against the average of the absolute correlation coefficients. The x-axis are the correlation coefficients of all N assets with each other computed using its daily returns from the chosen time periods and the y-axis is the ratio of the final returns against the classical optimal after Algorithm 1 converges using quantum annealers similar to Figure 5. The currency assets (stars) used in the tests all have higher correlation coefficients than those of the ETF assets, and generally yield better results.

On the 12 asset test shown in Figure 8, the 2000Q processor struggles to find the ground state as its embedding chain length reaches 16, while the Advantage processor provides results close to the simulated annealing and post-processed results. However, neither quantum annealer converges. Table 3 records the QUBO objective values of the last five iterations for the Advantage processor in this test. Although the objective values hardly differ, the solution quality is seemingly more sensitive to changes in the QUBO objective value for larger problems. A 0.1% change in the objective value leads to 30% difference in the portfolio variance. One potential reason is that larger problems have more assets that are less correlated, and as shown in Figure 7, smaller correlation coefficients generally equate to worse performance on quantum annealers. In this case, either the quantum annealers need to be more accurate to find the ground state, or our QUBO setup needs to be modified to account for higher asset counts.

Table 3. Objective values of last five iterations from simulated annealing and Advantage from the 12 asset test. This corroborates observations in Figure 8 that the Advantage processor is able to reach states with very good approximation ratios.

Last k Iteration	SA Objective	Advantage Objective	Difference
5	−1.026	−1.016	1.039%
4	−0.951	−0.950	0.092%
3	−0.879	−0.878	0.111%
2	−0.811	−0.809	0.170%
1	−0.746	−0.746	0.076%

For even larger problems of 23 assets, with the embedding chain lengths going up to 17, the Advantage processor fails to find the ground state by a large margin, as shown in Figure 9. Even though we can physically map a problem of this size, the results reflect the limitations of current-generation quantum annealers.

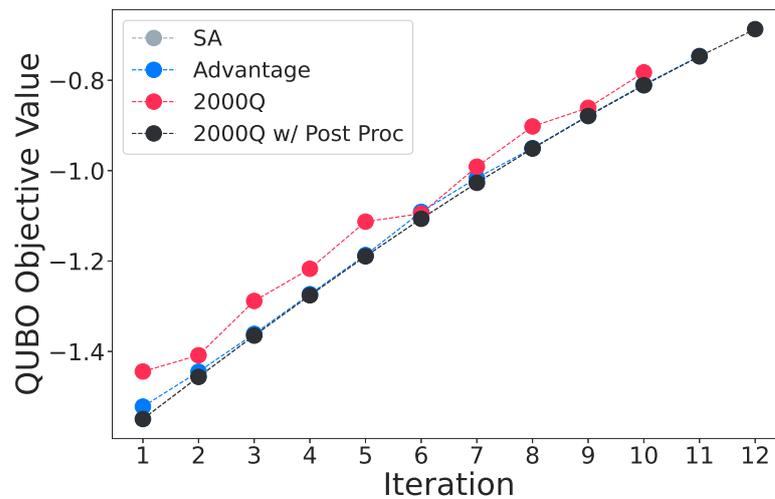


Figure 8. The objective comparison of the 12 asset test between all four backends. The solutions from 2000Q deviate from the ground states by a large margin, while the Advantage processor is able to keep up closely. Post-processing is able to improve the 2000Q results to once again match simulated annealing.

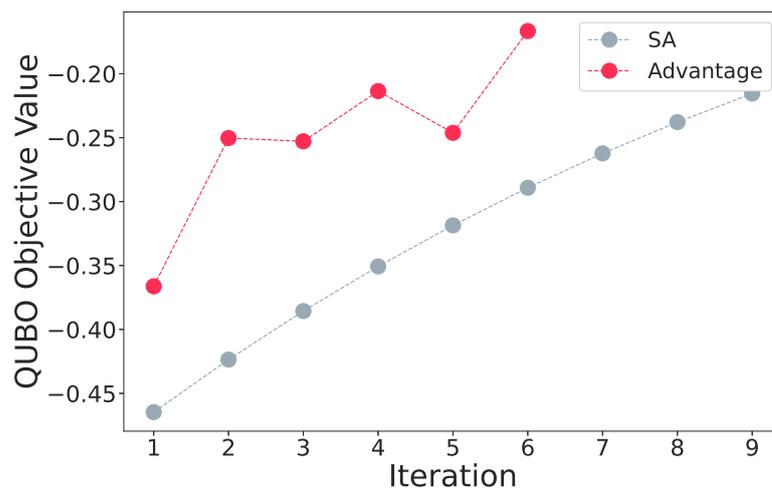


Figure 9. The objective comparison of the 23 asset test between simulated annealing and the Advantage processor. Due to the high chain lengths of the embedding, the Advantage processor fails to either reach the ground state or get close to it in all iterations, rendering the processor incapable of solving problems of such sizes.

5. Discussion

As newer quantum devices are released every year, it is important to design and benchmark algorithms across generations. As companies and researchers race to build the first quantum computer that can demonstrate quantum advantage on practical problems, different classes of quantum devices have emerged: general purpose quantum computers from IBM, Google, Honeywell, and others; the specialized quantum Ising machine from D-Wave; and quantum-inspired digital annealer from Fujitsu. These devices have different types of constraints due to different noise profiles, qubit connectivity, and/or implementable Hamiltonians, and none are perhaps at the scale and reliability needed to solve real-world problems at the edge of classical capability. Therefore, hybrid algorithms are needed to incorporate these quantum computers on practical problems with reasonable size.

In this paper, we have shown that it is not only possible to introduce such hybrid algorithm schemes that compute the optimal portfolios based on expected shortfall, but also highlighted where it is possible to reach working accuracy. We used a quantum annealer

to solve an asset allocation problem based on expected shortfall, employing a QUBO formulation of the Markowitz Optimization problem and interlacing it with a layer of classical decision-making. Here, we iteratively adjusted our problem Hamiltonian based on its feedback until the portfolio was within the desired risk threshold. The fact that both D-Wave 2000Q and Advantage quantum annealers performed reasonably well on the six-asset tests with portfolios' Sharpe ratios to above 80% of SA values is promising. Additionally encouraging is that the newer and more scalable Advantage processor achieved much better QUBO objective values on problems with 12 assets. Finally, we observed that both quantum annealers tended to obtain portfolios with higher returns on more correlated assets (Figure 7), which we believe should attract future research as it may help guide the application of quantum annealing on real-world applications in the near term.

Although the quantum annealers fell short on tests with more assets, we can remain optimistic about new hardware with more qubits, better connectivity, and lower noise in the near future. We also acknowledge the need to design algorithms that can scale with these new hardware, as we saw that the portfolio quality became increasingly sensitive to the QUBO objective values as we introduced more assets—results with 99.9% objective values of the optimal led to 30% more variance. Additionally, advances in gate-model quantum computers and combinatorial optimization algorithms [60,61] will provide other avenues for solving these problems. For example, it could be instructive to explore and compare to novel approaches, such as counterdiabatic techniques recently proposed for similar problems, but for gate-based systems [49].

Future research includes identifying subsets of problems that can be solved better on quantum devices, as we have discussed in Section 4. It is also important to find an efficient way to implement inequality constraints, as adding slack variables may not be the best choice in the QUBO. We also note that on specific test cases, the QUBO reformulation enables both simulated annealer and quantum annealers to find better portfolios than a classical convex optimizer *cvxpy*, by treating the constraints as soft. Other optimization problems might also benefit from QUBOs with soft constraints.

Author Contributions: Conceptualization, H.X., S.D. and A.B.; methodology, H.X. and S.D.; software, H.X.; validation, S.D., A.B. and A.P.; formal analysis, H.X., S.D., A.B. and A.P.; investigation, H.X.; resources, A.B.; data curation, H.X.; writing—original draft preparation, H.X.; writing—review and editing, H.X., S.D., A.B. and A.P.; visualization, H.X.; supervision, A.B. and A.P.; project administration, A.B. and A.P.; funding acquisition, A.B. All authors have read and agreed to the published version of the manuscript.

Funding: Funding for S.D. was supported in part by the US Department of Energy, Advanced Scientific Computing Research program office Quantum Algorithms Team project ERKJ335, through subcontract number 4000175762 to Purdue University from Oak Ridge National Laboratory managed by UT-Battelle, LLC acting under contract DE-AC05-00OR22725 with the Department of Energy. A.B. was supported by Purdue University, College of Science, Startup funds, and H.X. was supported by College of Science, Quantum Seed Grant. The access to D-Wave for this research was funded by the U.S. Department of Energy under Contract No. DE-AC05-00OR22725 through the Oak Ridge Leadership Computing Facility (OLCF) at the Oak Ridge National Laboratory (ORNL).

Institutional Review Board Statement: Not applicable.

Data Availability Statement: Data is available upon reasonable request.

Acknowledgments: We thank Travis Humble for overall suggestions and support for the project, and also Andrew King, Isil Ozfidan, and Erica Grant for helpful discussions. A.B. and S.D. thank the ORNL Quantum Computing Institute and the Purdue Quantum Science and Engineering Institute for their support.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Glover, F.; Kochenberger, G.; Du, Y. A Tutorial on Formulating and Using QUBO Models. *arXiv* **2019**, arXiv:1811.11538.
2. Pastorello, D.; Blanzieri, E. Quantum Annealing Learning Search for Solving QUBO Problems. *Quantum Inf. Process.* **2019**, *18*, 303. [CrossRef]
3. Kochenberger, G.; Hao, J.K.; Glover, F.; Lewis, M.; Lü, Z.; Wang, H.; Wang, Y. The Unconstrained Binary Quadratic Programming Problem: A Survey. *J. Comb. Optim.* **2014**, *28*, 58–81. [CrossRef]
4. Lucas, A. Ising Formulations of Many NP Problems. *Front. Phys.* **2014**, *2*, 5. [CrossRef]
5. Djidjev, H.N.; Chapuis, G.; Hahn, G.; Rizk, G. Efficient Combinatorial Optimization Using Quantum Annealing. *arXiv* **2018**, arXiv:1801.08653.
6. Ikeda, K.; Nakamura, Y.; Humble, T.S. Application of Quantum Annealing to Nurse Scheduling Problem. *Sci. Rep.* **2019**, *9*, 12837. [CrossRef] [PubMed]
7. Titiloye, O.; Crispin, A. Quantum Annealing of the Graph Coloring Problem. *Discret. Optim.* **2011**, *8*, 376–384. [CrossRef]
8. Yarkoni, S.; Raponi, E.; Bäck, T.; Schmitt, S. Quantum Annealing for Industry Applications: Introduction and Review. *Rep. Prog. Phys.* **2022**, *85*, 104001. [CrossRef]
9. Markowitz, H. Portfolio Selection. *J. Financ.* **1952**, *7*, 77–91. [CrossRef]
10. Black, F.; Litterman, R.B. Asset Allocation: Combining Investor Views with Market Equilibrium. *J. Fixed Income* **1991**, *1*, 7–18. [CrossRef]
11. McNeil, A.J.; Frey, R.; Embrechts, P. *Quantitative Risk Management: Concepts, Techniques and Tools—Revised Edition*; Princeton University Press: Princeton, NJ, USA, 2015.
12. Dasgupta, S.; Banerjee, A. Quantum Annealing Algorithm for Expected Shortfall Based Dynamic Asset Allocation. *arXiv* **2020**, arXiv:1909.12904.
13. Pardalos, P.M.; Vavasis, S.A. Quadratic Programming with One Negative Eigenvalue Is NP-hard. *J. Glob. Optim.* **1991**, *1*, 15–22. [CrossRef]
14. Vavasis, S.A. Complexity Theory: Quadratic Programming. In *Encyclopedia of Optimization*; Floudas, C.A., Pardalos, P.M., Eds.; Springer US: Boston, MA, USA, 2001; pp. 304–307. [CrossRef]
15. Grant, E.; Humble, T.S.; Stump, B. Benchmarking Quantum Annealing Controls with Portfolio Optimization. *Phys. Rev. Appl.* **2021**, *15*, 014012. [CrossRef]
16. O'Donnell, R. *Analysis of Boolean Functions*; Cambridge University Press: Cambridge, UK, 2014. [CrossRef]
17. Dattani, N. Quadraticization in Discrete Optimization and Quantum Mechanics. *arXiv* **2019**, arXiv:1901.04405.
18. Verma, A.; Lewis, M.; Kochenberger, G. Efficient QUBO Transformation for Higher Degree Pseudo Boolean Functions. *arXiv* **2021**, arXiv:2107.11695.
19. Mandal, A.; Roy, A.; Upadhyay, S.; Ushijima-Mwesigwa, H. Compressed Quadraticization of Higher Order Binary Optimization Problems. In Proceedings of the 2020 Data Compression Conference (DCC), Snowbird, UT, USA, 24–27 March 2020; pp. 383–383. [CrossRef]
20. Yahoo Finance iShares MSCI Emerging Markets ETF (EEM). Available online: <https://finance.yahoo.com/quote/EEM/history?p=EEM> (accessed on 18 May 2020).
21. Yahoo Finance Invesco QQQ Trust (QQQ). Available online: <https://finance.yahoo.com/quote/QQQ/history?p=QQQ> (accessed on 18 May 2020).
22. Yahoo Finance iShares Silver Trust (SLV). Available online: <https://finance.yahoo.com/quote/SLV/history?p=SLV> (accessed on 18 May 2020).
23. Yahoo Finance SPDR S&P 500 ETF Trust (SPY). Available online: <https://finance.yahoo.com/quote/SPY/history?p=SPY> (accessed on 18 May 2020).
24. Yahoo Finance ProShares UltraPro Short QQQ (SQQQ). Available online: <https://finance.yahoo.com/quote/SQQQ/history?p=SQQQ> (accessed on 18 May 2020).
25. Yahoo Finance Financial Select Sector SPDR Fund (XLF). Available online: <https://finance.yahoo.com/quote/XLF/history?p=XLF> (accessed on 18 May 2020).
26. Uryasev, S.; Rockafellar, R.T. Conditional Value-at-Risk: Optimization Approach. In *Stochastic Optimization: Algorithms and Applications*; Uryasev, S., Pardalos, P.M., Eds.; Applied Optimization, Springer US: Boston, MA, USA, 2001; pp. 411–435. [CrossRef]
27. Norton, M.; Khokhlov, V.; Uryasev, S. Calculating CVaR and bPOE for Common Probability Distributions with Application to Portfolio Optimization and Density Estimation. *Ann. Oper. Res.* **2021**, *299*, 1281–1315. [CrossRef]
28. Bertsimas, D.; Lauprete, G.J.; Samarov, A. Shortfall as a Risk Measure: Properties, Optimization and Applications. *J. Econ. Dyn. Control* **2004**, *28*, 1353–1381. [CrossRef]
29. Brooke, J.; Bitko, D.; F., T.; Rosenbaum.; Aeppli, G. Quantum Annealing of a Disordered Magnet. *Science* **1999**, *284*, 779–781. [CrossRef]
30. Santoro, G.E.; Martoňák, R.; Tosatti, E.; Car, R. Theory of Quantum Annealing of an Ising Spin Glass. *Science* **2002**, *295*, 2427–2430. [CrossRef]
31. King, A.D.; Raymond, J.; Lanting, T.; Harris, R.; Zucca, A.; Altomare, F.; Berkley, A.J.; Boothby, K.; Ejtemaee, S.; Enderud, C.; et al. Quantum Critical Dynamics in a 5000-Qubit Programmable Spin Glass. *arXiv* **2022**, arXiv:2207.13800.

32. Venegas-Andraca, S.E.; Cruz-Santos, W.; McGeoch, C.; Lanzagorta, M. A Cross-Disciplinary Introduction to Quantum Annealing-Based Algorithms. *Contemp. Phys.* **2018**, *59*, 174–197. [[CrossRef](#)]
33. Barahona, F. On the Computational Complexity of Ising Spin Glass Models. *J. Phys. A Math. Gen.* **1982**, *15*, 3241–3253. [[CrossRef](#)]
34. Zhang, Z. Computational Complexity of Spin-Glass Three-Dimensional (3D) Ising Model. *J. Mater. Sci. Technol.* **2020**, *44*, 116–120. [[CrossRef](#)]
35. Vinci, W.; Lidar, D.A. Non-Stoquastic Hamiltonians in Quantum Annealing via Geometric Phases. *Npj Quantum Inf.* **2017**, *3*, 1–6. [[CrossRef](#)]
36. Farhi, E.; Goldstone, J.; Gutmann, S.; Sipser, M. Quantum Computation by Adiabatic Evolution. *arXiv* **2000**, arXiv:quant-ph/0001106.
37. Albash, T.; Lidar, D.A. Adiabatic Quantum Computation. *Rev. Mod. Phys.* **2018**, *90*, 015002. [[CrossRef](#)]
38. Born, M.; Fock, V. Beweis des Adiabatenatzes. *Z. für Phys.* **1928**, *51*, 165–180. [[CrossRef](#)]
39. Inc., D.W.S. The Practical Quantum Computing Company. *ICE: Dynamic Ranges in h and J Values*; 2021.
40. Rosenberg, G.; Haghnegahdar, P.; Goddard, P.; Carr, P.; Wu, K.; de Prado, M.L. Solving the Optimal Trading Trajectory Problem Using a Quantum Annealer. *IEEE J. Sel. Top. Signal Process.* **2016**, *10*, 1053–1060. [[CrossRef](#)]
41. Venturelli, D.; Kondratyev, A. Reverse Quantum Annealing Approach to Portfolio Optimization Problems. *Quantum Mach. Intell.* **2019**, *1*, 17–30. [[CrossRef](#)]
42. Phillipson, F.; Bhatia, H.S. Portfolio Optimisation Using the D-Wave Quantum Annealer. In *Proceedings of the Computational Science—ICCS 2021*; Lecture Notes in Computer Science; Paszynski, M.; Kranzlmüller, D., Krzhizhanovskaya, V.V., Dongarra, J.J., Sloat, P.M.A., Eds.; Springer International Publishing: Cham, Switzerland, 2021; pp. 45–59. [[CrossRef](#)]
43. Metropolis, N.; Rosenbluth, A.W.; Rosenbluth, M.N.; Teller, A.H.; Teller, E. Equation of State Calculations by Fast Computing Machines. *J. Chem. Phys.* **1953**, *21*, 1087–1092. [[CrossRef](#)]
44. Kirkpatrick, S.; Gelatt, C.D.; Vecchi, M.P. Optimization by Simulated Annealing. *Science* **1983**, *220*, 671–680. [[CrossRef](#)]
45. Liu, Y.J.; Zhang, W.G. A Multi-Period Fuzzy Portfolio Optimization Model with Minimum Transaction Lots. *Eur. J. Oper. Res.* **2015**, *242*, 933–941. [[CrossRef](#)]
46. Vercher, E.; Bermúdez, J.D. Portfolio Optimization Using a Credibility Mean-Absolute Semi-Deviation Model. *Expert Syst. Appl.* **2015**, *42*, 7121–7131. [[CrossRef](#)]
47. Mansini, R.; Ogryczak, W.; Speranza, M.G. Twenty Years of Linear Programming Based Portfolio Optimization. *Eur. J. Oper. Res.* **2014**, *234*, 518–535. [[CrossRef](#)]
48. Schaerf, A. Local Search Techniques for Constrained Portfolio Selection Problems. *Comput. Econ.* **2002**, *20*, 177–190. [[CrossRef](#)]
49. Hegade, N.N.; Chandarana, P.; Paul, K.; Chen, X.; Albarrán-Arriagada, F.; Solano, E. Portfolio Optimization with Digitized-Counterdiabatic Quantum Algorithms. *Phys. Rev. Res.* **2022**, *4*, 043204. [[CrossRef](#)]
50. Inc., D.W.S. The Practical Quantum Computing Company. The D-Wave Advantage System: An Overview. 2021. Available online: https://www.dwavesys.com/media/s3qbjp3s/14-1049a-a_the_d-wave_advantage_system_an_overview.pdf (accessed on 18 May 2020).
51. Venturelli, D.; Mandrà, S.; Knysh, S.; O’Gorman, B.; Biswas, R.; Smelyanskiy, V. Quantum Optimization of Fully Connected Spin Glasses. *Phys. Rev. X* **2015**, *5*, 031040. [[CrossRef](#)]
52. Boothby, T.; King, A.D.; Roy, A. Fast Clique Minor Generation in Chimera Qubit Connectivity Graphs. *Quantum Inf. Process.* **2016**, *15*, 495–508. [[CrossRef](#)]
53. Pelofske, E.; Hahn, G.; Djidjev, H. Optimizing the Spin Reversal Transform on the D-Wave 2000Q. In *Proceedings of the 2019 IEEE International Conference on Rebooting Computing (ICRC)*, San Mateo, CA, USA, 6–8 November 2019; pp. 1–8. [[CrossRef](#)]
54. Lanting, T.; Amin, M.H.; Baron, C.; Babcock, M.; Boschee, J.; Boixo, S.; Smelyanskiy, V.N.; Foygel, M.; Petukhov, A.G. Probing Environmental Spin Polarization with Superconducting Flux Qubits. *arXiv* **2020**, arXiv:2003.14244.
55. Pudenz, K.L. Parameter Setting for Quantum Annealers. In *Proceedings of the 2016 IEEE High Performance Extreme Computing Conference (HPEC)*, Waltham, MA, USA, 13–15 September 2016; pp. 1–6. [[CrossRef](#)]
56. Markowitz, H.M. The Elimination Form of the Inverse and Its Application to Linear Programming. *Manag. Sci.* **1957**, *3*, 255–269. [[CrossRef](#)]
57. Jensen, F.V.; Lauritzen, S.L.; Olesen, K.G. Bayesian Updating in Causal Probabilistic Networks by Local Computations. *Comput. Stat. Q.* **1990**, *4*, 269–282.
58. Diamond, S.; Boyd, S.P. CVXPY: A Python-Embedded Modeling Language for Convex Optimization. *J. Mach. Learn. Res.* **2016**, *17*, 2909–2913.
59. Barbosa, A.; Pelofske, E.; Hahn, G.; Djidjev, H.N. Using Machine Learning for Quantum Annealing Accuracy Prediction. *Algorithms* **2021**, *14*, 187. [[CrossRef](#)]
60. Farhi, E.; Goldstone, J.; Gutmann, S. A Quantum Approximate Optimization Algorithm. *arXiv* **2014**, arXiv:1411.4028.
61. Hadfield, S.; Wang, Z.; O’Gorman, B.; Rieffel, E.G.; Venturelli, D.; Biswas, R. From the Quantum Approximate Optimization Algorithm to a Quantum Alternating Operator Ansatz. *Algorithms* **2019**, *12*, 34. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.