Multi-period Mean Expected-Shortfall Strategies: "Cut your Losses and Ride Your Gains"

Peter A. Forsyth[a] Kenneth R. Vetzal[b]

Sept. 19, 2022

Abstract

Dynamic mean-variance (MV) optimal strategies are inherently contrarian. Following periods of strong equity returns, there is a tendency to de-risk the portfolio by shifting into risk-free investments. On the other hand, if the portfolio still has some equity exposure, the weight on equities will increase following stretches of poor equity returns. This is essentially due to using variance as a risk measure, which penalizes both upside and downside deviations relative to a satiation point. As an alternative, we propose a dynamic trading strategy based on an expected wealth (EW), expected shortfall (ES) objective function. ES is defined as the mean of the worst β fraction of the outcomes, hence the EW-ES objective directly targets left tail risk. We use stochastic control methods to determine the optimal trading strategy. Our numerical method allows us to impose realistic constraints: no leverage, no shorting, infrequent rebalancing. For 5 year investment horizons, this strategy generates an annualized alpha of 180 bps compared to a 60:40 stock-bond constant weight policy. Bootstrap resampling with historical data shows that these results are robust to parametric model misspecification. The optimal EW-ES strategy is generally a momentum-type policy, in contrast to the contrarian MV optimal strategy.

Keywords: optimal control, expected shortfall, apparent alpha, tail risk, asset allocation, resampled backtests
JEL codes: G11, G22
AMS codes: 91G, 65N06, 65N12, 35Q93

1 Introduction

The Sharpe ratio is a commonly used measure of investment performance. However, the Sharpe ratio is easy to manipulate. Any strategy which includes non-linear payoffs (e.g. a portfolio including options) can produce an apparent outperformance (Dybvig and Ingersoll, 1982; Lhabitant, 2000; Goetzmann et al., 2002). As noted by Spurgon (2001), selling off the right side of the terminal wealth distribution can improve the Sharpe ratio. Such a strategy is simple to implement by owning the underlying asset and selling out of the money calls on the asset (covered call writing).

Of course, in a complete market options can be replicated by dynamically trading stocks and bonds. Consequently any portfolio containing options is equivalent to a dynamic trading strategy. Hence, Sharpe ratios can be maximized by using optimal stochastic control techniques, coupled with

[a]David R. Cheriton School of Computer Science, University of Waterloo, Waterloo ON, Canada N2L 3G1, paforsyt@uwaterloo.ca, +1 519 888 4567 ext. 44415.

[b]School of Accounting and Finance, University of Waterloo, Waterloo ON, Canada N2L 3G1, kvetzal@uwaterloo.ca, +1 519 888 4567 ext. 46518.

a suitable objective function. An interesting corollary to this observation is the use of stochastic control in fraud detection (Bernard and Vanduffel, 2014).

Except for extremely pathological cases, a strategy which maximizes a mean-variance objective function will also maximize the Sharpe ratio. In the dynamic trading context, a mean-variance optimal strategy can be determined by minimizing a quadratic target objective function (Li and Ng, 2000; Zhou and Li, 2000; Vigna, 2014). In a complete market, the optimal strategy never exceeds the target (Cui et al., 2012; Vigna, 2014; Bauerle and Grether, 2015; Dang and Forsyth, 2016). Consequently, even though the risk measure is symmetric (i.e. the variance), the optimal strategy produces a highly skewed terminal wealth distribution (Goetzmann et al., 2002; van Staden et al., 2021).

A dynamic investment strategy which maximizes the Sharpe ratio also produces an *apparent* alpha (Goetzmann et al., 2002) relative to a fixed proportion strategy. However, due to the highly skewed distributions produced by these strategies, it is not clear that the resulting wealth distribution is actually desired by the investor, in spite of these apparently good performance metrics. Of course, it could be argued that variance is a poor risk measure in any case. Investors are more concerned with downside risk measures. Indeed, upside volatility can be considered desirable.

As a result, in this paper, we propose using expected shortfall (ES) as the risk measure. ES is simply the average of the worst β fraction of outcomes, and hence is a downside tail risk measure. ES is basically the negative of the conditional value at risk (CVAR). Using stochastic control techniques, we determine the optimal investment strategies which are Pareto optimal, with reward given by expected terminal wealth (EW), and risk measured by ES. For medium-term investments (i.e. 2-5 years) the optimal EW-ES strategy generates a significant annualized alpha compared to constant weight policies. The optimal controls are determined using a parametric model of stock and bond processes, calibrated to 93 years of historical data. We verify that these strategies are robust to parametric model uncertainty, by testing the strategies on bootstrapped resampled historical data.

It is interesting to observe that the Sharpe ratio maximizing dynamic strategy (or equivalently quadratic shortfall minimizing) is fundamentally contrarian. Under this policy, when stocks increase in value, they are sold, and assets shifted to bonds. When stocks decrease in value, stocks are bought, and bond holdings reduced. However, we see the completely opposite policy in Pareto optimal EW-ES strategies. The optimal EW-ES policy has more of a momentum character: when stocks go up in value, stock holdings are increased. When stocks go down in value, stocks are sold, and assets shifted to bonds. This is simply a consequence of the ES penalty; when the investor's wealth is reduced, increasing the bond fraction protects against further moves to the downside. In fact, this strategy provides a mathematical rationale for the trader's maxim *"Cut your losses, ride your gains"*.

Since the Sharpe ratio maximizing and EW-ES optimal strategies have fundamentally different investment policies, and hence different terminal wealth distributions, each approach may appeal to different investors, at different stages in their investment lifecyles. However, for medium-term investors having wealth-preservation as a high priority, EW-ES strategies are worth considering.

We note that both the EW-ES and Sharpe ratio maximizing strategies are pre-commitment polices, which are not formally time consistent. However, this is really just a matter of interpretation, since for both strategies there is an equivalent induced objective function which generates the same controls, yet is time consistent (Strub et al., 2019). Hence, in both cases, these policies are implementable (Forsyth, 2020a). In addition, as noted by Bernard and Vanduffel (2014), if the EW-ES strategy is realized in an investment product sold to a retail investor then the optimal policy from the investor's point of view is in fact of pre-commitment type, since the retail client does not herself trade in the underlying assets during the lifetime of the contract.

In the case of an investment product sold to a retail investor, we can envision that the investor

repeatedly buys the product for $2-5$ year terms. Our results show that over a 5-year term, this product generates an alpha of 180 bps per year compared to a 60:40 stock-bond portfolio. Even for a 2-year investment horizon, the EW-ES strategy generates an alpha of 120 bps per year compared to a 60:40 stock-bond portfolio. This product would be appealing to investors who want to outperform a typical constant weight strategy, but who are also concerned with worst case tail risk over $2-5$ year periods.

To begin, in Section 2, we review the known results for Sharpe ratio maximizing strategies, and define the alpha of these strategies relative to fixed weight policies. We also provide some illustrative results, documenting the performance of these strategies. We subsequently proceed to formally define the EW-ES problem, describe our algorithm for determination of the optimal control, define the appropriate alpha, and discuss the numerical results.

2 Background: Maximizing Sharpe Ratios and Defining Alpha

Let r be the risk-free return, and T be the investment horizon. W_t is the wealth of a portfolio at time t. The continuously compounded Sharpe ratio is then defined as

$$\mathbb{S} = \frac{E[W_T] - W_0 e^{rT}}{std[W_T]}, \tag{2.1}$$

where $E[\cdot]$ and $std[\cdot]$ respectively denote expectation and standard deviation. Note that $\mathbb{S}$ is defined in terms of the terminal wealth and standard deviation at time T (Lhabitant, 2000; Goetzmann et al., 2002; Bernard and Vanduffel, 2014), in contrast to the *instantaneous* Sharpe ratio, which is defined in terms of averaging short period returns. Dynamic trading strategies will have different equity exposure over different short-term periods, and so averaging short period returns is not a meaningful metric.

Consider a market containing a stock index and a risk-free bond. Let the amount invested in the stock index be S_t, and the amount in the risk-free bond be B_t. We assume that

$$\begin{aligned} \frac{dS_t}{S_t} &= \mu\ dt + \sigma\ dZ \\ \frac{dB_t}{B_t} &= r\ dt\ , \end{aligned} \tag{2.2}$$

where μ is the stock drift rate, σ is the volatility, and dZ is the increment of a Wiener process. Let p be the fraction of the total portfolio W_t invested in the stock. Assuming continuous rebalancing, then the process for W_t is

$$\begin{aligned} \frac{dW_t}{W_t} &= p\frac{dS_t}{S_t} + (1-p)\frac{dB_t}{B_t} \\ &= (r + p(\mu - r))\ dt + p\sigma dZ\ . \end{aligned} \tag{2.3}$$

Given an initial investment W_0 at $t=0$, with terminal wealth W_T, we can pose the problem of determining the optimal control $p(W_t, t), t \in [0,T]$ in terms of a mean-variance objective. Defining a scalarization parameter $\kappa > 0$, the mean-variance problem can be formulated as

$$\sup_{p(\cdot)} E[W_T] - \kappa \mathrm{Var}[W_T]\ . \tag{2.4}$$

Varying κ in equation (2.4) traces out the efficient frontier. Problem (2.4) cannot be solved directly using dynamic programming. From (Zhou and Li, 2000; Li and Ng, 2000), we learn that we

can determine the control $p(\cdot)$ which maximizes objective function (2.4) by solving the alternative problem

$$\inf_{p(\cdot)} E[(W^* - W_T)^2] \ , \tag{2.5}$$

where we trace out the efficient frontier by varying W^*. Note that Problem 2.5 can be solved using dynamic programming.[1]

The optimal control for Problem 2.5 is given by (Zhou and Li, 2000; Li and Ng, 2000; Vigna, 2014)

$$\begin{aligned} p &= \frac{\xi}{\sigma W_t}\left(W^* e^{-r(T-t)} - W_t\right) \\ \text{where } \xi &= \frac{\mu - r}{\sigma} \ . \end{aligned} \tag{2.6}$$

Let W_T^{opt} denote the terminal wealth under strategy (2.6). The efficient frontier is then given by (Zhou and Li, 2000; Li and Ng, 2000)

$$\begin{aligned} E[W_T^{opt}] &= W_0 e^{rT} + \left(e^{\xi^2 T} - 1\right)^{1/2} \sqrt{Var(W_T^{opt})} \\ &= W_0 e^{rT} + \left(e^{\xi^2 T} - 1\right)^{1/2} std(W_T^{opt}) \ . \end{aligned} \tag{2.7}$$

where $Var[\cdot]$ denotes variance.

Recall that varying W^* will move us along the efficient frontier. Since equation (2.7) is a monotone increasing function of variance, for a fixed value of $E[W_T^{opt}]$ the strategy which minimizes $Var(W_T^{opt})$ also minimizes $std(W_T^{opt})$. Consequently, from equations (2.1) and (2.7), the optimal Sharpe ratio is

$$\mathbb{S}^{opt} = \left(e^{\xi^2 T} - 1\right)^{1/2} \ . \tag{2.8}$$

On the other hand, suppose we rebalance to a constant weight, i.e. $p = const.$ in equation (2.3). Let W_T^p denote the terminal wealth under a constant weight strategy p. Then we have

$$\begin{aligned} E[W_T^p] &= W_0 e^{(p(\mu - r) + r)T} \\ std[W_T^p] &= E[W_T^p]\left(e^{(\sigma p)^2 T} - 1\right)^{1/2} \ , \end{aligned} \tag{2.9}$$

with Sharpe ratio

$$\begin{aligned} \mathbb{S}^p &= \frac{1 - e^{-p(\mu - r)T}}{\left(e^{(\sigma p)^2 T} - 1\right)^{1/2}} \\ &\simeq \xi\sqrt{T} = \frac{(\mu - r)\sqrt{T}}{\sigma} \ ; \ T \to 0 \text{ or } p \to 0 \ . \end{aligned} \tag{2.10}$$

Note that the continuously compounded Sharpe ratio is a function of p in general, and approaches the instantaneous Sharpe ratio only in the limit as $T \to 0$ or $p \to 0$.

[1] In Vigna (2014), it is shown that $W_t < W^*, \forall t$ under the optimal control.

Let

$$E[W_T^{opt,p}] = W_0 e^{rT} + \left(e^{\xi^2 T} - 1 \right)^{1/2} std(W_T^p) \ . \tag{2.11}$$

This can be interpreted as follows. Given a constant weight strategy with equity fraction p, which generates $std(W_T^p)$, $E[W_T^{opt,p}]$ is the expected terminal wealth under control (2.6) which has the same the standard deviation $std(W_T^p)$ (this follows from equation (2.7)).

A convenient way to compare these strategies is through the *apparent* annualized α, which we define as

$$\alpha^p = \frac{\log(E[W_T^{opt,p}]) - \log(E[W_T^p])}{T} \ . \tag{2.12}$$

This is the extra annualized expected return generated by strategy (2.6) compared to the constant weight strategy with equity fraction p, given that both strategies have the same risk, as measured by standard deviation. Consistent with Goetzmann et al. (2002), we call α^p the apparent alpha, since there is no stock-picking skill involved here, merely use of a dynamic control.

Before proceeding to some illustrative numerical examples, we note that the MV optimal strategy with control (2.6) does not restrict the equity weight p in any way, permitting unlimited leverage and short-selling. As a result, the strategy is typically very aggressive, with heavy use of leverage early on. However, a rigorous solution of Problem 2.4 with no-shorting and no-leverage constraints requires numerical solution of a Hamilton-Jacobi-Bellman (HJB) equation (Wang and Forsyth, 2010). It is more instructive for our purposes to approximate the constrained control using the approach in Vigna (2014). We constrain the unconstrained control so that there is no-shorting and no-leverage, by letting p^* be the unconstrained solution to (2.6) and setting

$$p = \max(0.0, \min(p^*, 1.0)) \ . \tag{2.13}$$

In the following, we will refer to the strategy that uses this constrained approach as the Clipped MV Optimal strategy, in contrast to the unconstrained MV Optimal strategy (2.6).

We conclude this section with a couple of comments about options trading and pre-commitment and time consistency. With regard to options, we note that any portfolio that uses covered call writing (or equivalently, cash-secured put writing) can be replicated by continuous trading in a portfolio which has a risk-free bond and the underlying stock that satisfies the no-shorting, no-leverage constraints as in the Clipped MV Optimal strategy (see Appendix A). More generally, in the complete market case dynamic trading in the bond and stock is equivalent to using options in the trading strategy. Hence, even if options are not directly included in, for example, the MV Optimal strategy (2.6), this is clearly equivalent to the use of derivatives. As a result, we can think of any financial product based on a dynamic trading strateyg as a structured product. In the presence of constraints, the market may be incomplete. However, with some abuse of common terminology, we will still refer to such packaged investment vehicles as structured products, even in an incomplete market.

With respect to pre-commitment and time consistency, we make the following two observations:

Remark 2.1 (Pre-commitment policy)**.** *Strategy (2.6) is the pre-commitment solution, which is not formally time consistent. However, consider the case of a retail investor, who purchases a financial product from a financial institution. The investor does not trade herself in the assets underlying the product during the life of the contract. Hence, the performance of the product is evaluated in terms of the initial and final wealth. Consequently, as noted in Bernard and Vanduffel (2014), the pre-commitment policy is appropriate in this case.*

Remark 2.2 (Pre-commitment strategies equivalence to induced time consistent strategy)**.** *The control (2.6) is formally the pre-commitment policy. However, the time zero strategy based on the pre-commitment policy solution of Problem 2.4 is identical to the strategy for an induced time consistent policy, and hence it is implementable.*[2] *The induced time consistent strategy in this case is a target based shortfall, Problem 2.5, with a fixed value of* $W^* \ \forall t > 0$*. The concept of induced time consistent strategies is discussed in Strub et al. (2019). The relationship between pre-commitment and implementable target-based schemes in the mean-variance context is discussed in Vigna (2014; 2022) and Menoncin and Vigna (2017).*

2.1 Numerical Examples

We use data from the Center for Research in Security Prices (CRSP) on a monthly basis over the 1926:1-2019:12 period.[3] Our base case tests use the CRSP 30 day T-bill for the bond asset and the CRSP value-weighted total return index for the stock asset. This latter index includes all distributions for all domestic stocks trading on major U.S. exchanges. All of these various indexes are in nominal terms, so we adjust them for inflation by using the U.S. CPI index (also obtained from CRSP). Since we are considering a multi-year investment horizon, it is important to use real (i.e. inflation-adjusted) returns.

With constant parameters, specification (2.2) assumes geometric Brownian motion with drift μ and volatility σ for the stock index, and a constant risk-free rate r. Table 2.1 gives maximum likelihood estimates for μ and σ, as well as the long-run sample average value of the 30 day T-bill rate which we take as a proxy for r. This table also summarizes our investment scenario, with initial wealth $W_0 = 1000$, an investment horizon of $T = 5$ years, and monthly rebalancing.

Stochastic Model Parameters		Investment Scenario	
μ	.0822	W_0	1000
σ	.1842	T	5 years
r	.0044	Rebalancing	Monthly

TABLE 2.1: *Investment scenario and estimated annualized parameters for processes (2.2), based on inflation-adjusted value-weighted CRSP index and 30 day T-bills. Sample period 1926:1 to 2019:12.*

The illustrative examples both here and in subsequent sections of the paper are based on two different simulation procedures. The first, which we call the synthetic market, uses standard Monte Carlo simulation techniques for asset returns, assuming that the model is correctly specified (i.e. in this case, we simulate geometric Brownian motion for stock index returns with $\mu = .0822$ and $\sigma = .1842$, and we set the risk-free rate r to the constant value of .0044). The second simulation procedure, which we refer to as the historical market, relies on the stationary block bookstrap method (Politis and Romano, 1994; Politis and White, 2004; Patton et al., 2009; Dichtl et al., 2016). In particular, we use monthly resampled returns from the historical data set, with data drawn with replacements in blocks of various sizes simultaneously for the stock and bond indexes. Sampling the data in blocks incorporates potential serial correlation in the real return data. The blocksizes are

[2] An implementable strategy has the property that the investor has no incentive to deviate from the strategy computed at time zero at later times (Forsyth, 2020a).

[3] More specifically, results presented here were calculated based on data from Historical Indexes, ©2020 Center for Research in Security Prices (CRSP), the University of Chicago Booth School of Business. Wharton Research Data Services was used in preparing this article. This service and the data available thereon constitute valuable intellectual property and trade secrets of WRDS and/or its third-party suppliers.

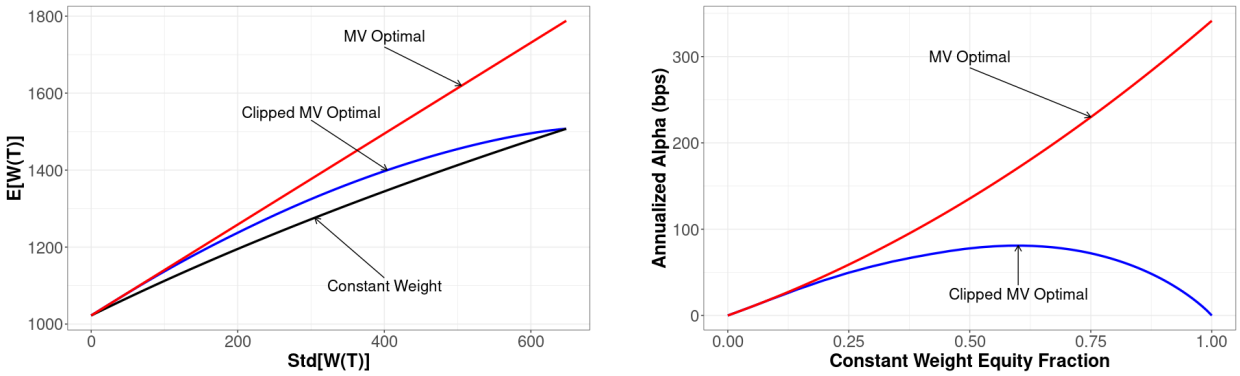

(a) *Efficient frontiers.* (b) *Annualized* α.

FIGURE 2.1: *Comparison of constant weight and MV optimal strategies in the synthetic market with 640,000 paths and parameters from Table 2.1.*

randomly drawn, based on an expected blocksize parameter $\hat{b} = 1/v$ where the block size follows a geometric distribution with $Prob(b = k) = (1 - v)^{k-1}v$. Enough blocks are then pasted together to construct simulated paths over the investment horizon of $T = 5$ years. The optimal value of $\hat{b}$ can be estimated using an algorithm from Patton et al. (2009). However, applying this algorithm separately to the real stock index and 30-day T-bill series results in quite different estimates: about 3 months for the former, and 50 months for the latter. If we take the average estimate from the two series, we get about two years. However, we provide results for a range of expected blocksizes as a check on the robustness of the historical market results. We emphasize that these historical market simulations make no assumptions about the stochastic processes governing bond and stock index returns.[4]

Figure 2.1 shows synthetic market results for constant weight strategies (with $p \in [0,1]$), along with MV Optimal and Clipped MV Optimal strategies, with 640,000 paths and monthly rebalancing. Panel (a) plots the efficient frontiers, i.e. $E[W_T]$ vs. $std[W_T]$. Given that the risk-free rate is assumed to be non-stochastic, it is possible to achieve a standard deviation of zero by investing entirely in the risk-free asset. The Clipped MV Optimal strategy is clearly constrained to match the constant weight strategy for constant equity weights of both $p = 0$ and $p = 1$, but it offers some outperformance compared to the constant weight strategy using intermediate values of p. The MV Optimal strategy provides strong outperformance, particularly for high levels of risk as measured by standard deviation. Panel (b) plots the annualized alpha (2.12) in bps for the MV Optimal and Clipped MV Optimal strategies, relative to constant weight strategies. The MV Optimal strategy gives very high apparent alpha (over 300 bps) as the constant equity weight approaches 1, but this unconstrained strategy is arguably quite unrealistic, having no limit whatsoever on leverage. In contrast, the Clipped MV Optimal strategy gives a maximum apparent alpha of about 80 bps relative to a constant weight strategy with $p \approx 0.6$, but as noted above the constraints imposed imply that the Clipped MV Optimal strategy cannot offer any outperformance compared to using a constant weight of $p = 1$.[5]

[4]Detailed pseudo-code for block bootstrap resampling can be found in Forsyth and Vetzal (2019).

[5]The results provided in Figure 2.1 are based on monthly rebalancing and 640,000 paths, but the actual control for the MV Optimal strategy (2.6) assumes continuous rebalancing. Obviously, in practice rebalancing must be discrete. Additional simulations were run increasing the number of paths by a factor of 4 and increasing the rebalancing frequency to 10 times per month. The expected value of terminal wealth and its standard deviation in all cases changed by less than 1% from the values given in Figure 2.1(a). Somewhat larger changes (up to almost 5%) were observed for the apparent alphas in Figure 2.1(b), in cases where there was a significant equity weight.

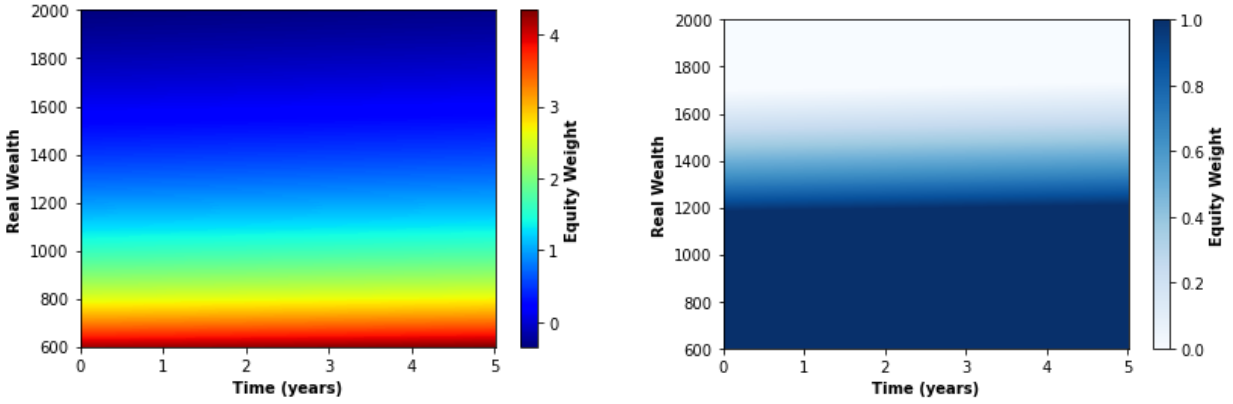

(a) *MV Optimal strategy.* (b) *Clipped MV Optimal strategy.*

FIGURE 2.2: *Heatmaps of optimal equity weights based on parameters from Table 2.1 with* $W^* = 1736$.

The MV Optimal control (2.6) depends on time only through the present value factor $e^{-r(T-t)}$. Since we are working in real terms and the long run real interest rate r at .0044 is not much above zero, the control strategy at any point in time should depend almost entirely on the level of real wealth then. This is seen in Figure 2.2, which provides heatmaps indicating the optimal fraction invested in equities at various wealth levels over the 5-year investment horizon for the MV Optimal and Clipped MV Optimal strategies.[6] The contrarian behavior of these strategies alluded to above is clear from both panels of Figure 2.2, as the portfolio has reduced equity exposure for high levels of real wealth, which would follow from strong prior market returns. Conversely, the strategies increase equity exposure at low real wealth, which would result from poor previous returns. The two panels are plotted with different color scales to highlight the extremely aggressive nature of the MV Optimal strategy. At the initial level $W_0 = 1000$ the equity weight is about 1.60, and this would increase to more than 4 if real wealth were to decline by about 40%.[7]

Further insight into the properties of both MV optimal strategies can be gleaned from examining the evolution of the densities of real wealth over time, as shown in Figure 2.3. These plots use the same W^* of 1763 with 640,000 paths and monthly rebalancing, although the densities are plotted on a quarterly basis. In addition to the densities, both plots show a single vertical black line at $W_0 = 1000$, as well as three colored vertical lines, which mark the 5th, 50th, and 95th percentiles of the distribution at each point in time shown. Since we are rebalancing discretely, there is a non-zero probability of exceeding W^*. Both strategies have a large left skew. This is particularly true for the MV Optimal strategy in panel (a). In contrast, the Clipped MV Optimal strategy in panel (b) cuts off part of the extreme left tail.

We next consider historical market simulations. Recall that this entails resampling from the actual observed data, rather than a Monte Carlo simulation of the assumed stochastic model. We again use 640,000 paths, and start by considering an expected blocksize of 60 months. Figure 2.4 is analogous to Figure 2.1 above, showing efficient frontiers and apparent alpha. In this case, there are some non-Pareto optimal points which have been removed for plotting purposes. For example, the frontier for the constant weight strategy in panel (a) bends back, with higher $std[W_T]$ and lower $E[W_T]$ if we consider weights closer to zero. Unlike the synthetic market with a constant interest rate where risk could be pushed all the way to zero, that is impossible here since the interest rate

[6]Here we specify $W^* = 1736$, a value which results in the Clipped MV Optimal strategy having $std[W_T]$ that is approximately equal to that produced by a constant weight strategy with $p = 0.6$.

[7]Note that control (2.6) implies that leverage is unbounded as $W_t \to 0^+$.

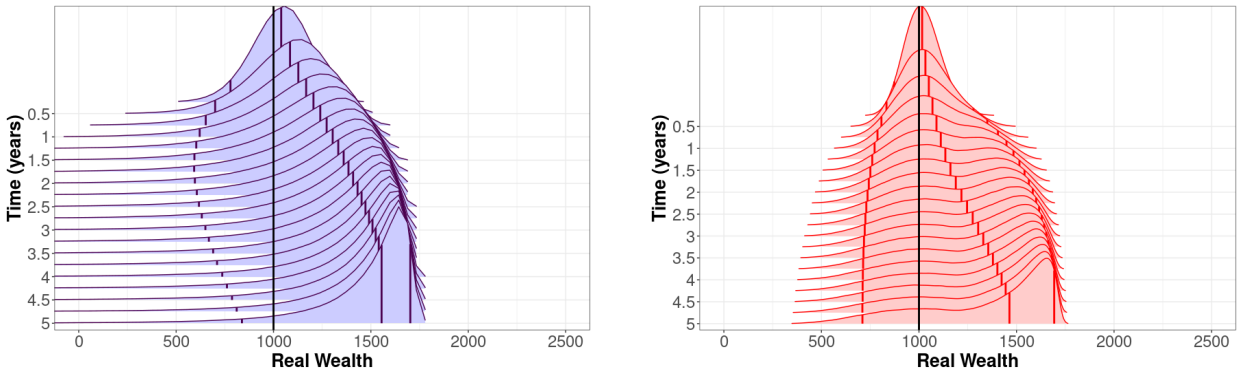

(a) *MV Optimal strategy.* (b) *Clipped MV Optimal strategy.*

FIGURE 2.3: *Densities of real wealth over time based on data from Table 2.1 with* $W^* = 1736$*, 640,000 simulated synthetic market paths, and monthly rebalancing. Densities plotted quarterly. The black vertical line indicates the initial real wealth* $W_0 = 1000$*. The colored vertical lines in each density represent the 5th, 50th, and 95th percentiles of the distribution.*

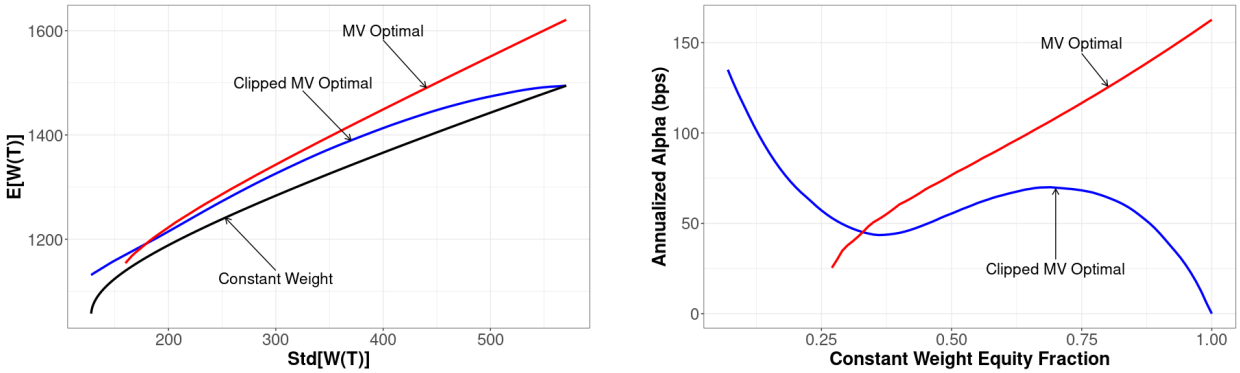

(a) *Efficient frontiers.* (b) *Annualized* α*.*

FIGURE 2.4: *Comparison of constant weight and MV optimal strategies in the historical market with 640,000 paths, expected blocksize of 60 months, and parameters from Table 2.1. Non-Pareto optimal points removed for each strategy.*

is stochastic. Panel (b) shows the expected alpha of zero for the Clipped MV Optimal strategy relative to the constant weight strategy with $p = 1$, and a corresponding significant alpha for the unconstrained MV Optimal strategy. However, this is much reduced from the level found in Figure 2.1(b). Panel (b) also indicates high alpha for the Clipped MV Optimal strategy compared to constant weight strategies that are heavily invested in bonds, but from panel (a) this is largely because the constant weight strategy performs poorly in these cases. As might be expected, the performance of the two MV Optimal strategies is worse under the conditions of these historical market tests, relative to the synthetic market which uses return data based on the geometric Brownian motion model that is assumed when deriving the control strategy.

These conclusions are reinforced by Figure 2.5, which repeats the exercise but this time draws the historical data using an expected blocksize of 24 months. This is an even more stringent test, since there can be many paths with very large changes in interest rates. For example, a period of sampling from the early 1980s when interest rates were very high may be followed on the same simulated path by a stretch of sampling from the 2010s when interest rates were quite low. While this type of situation could happen with the larger expected blocksize of 60 months considered

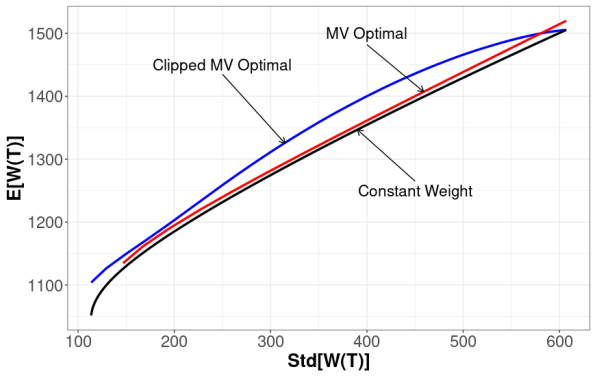

(a) *Efficient frontiers.*

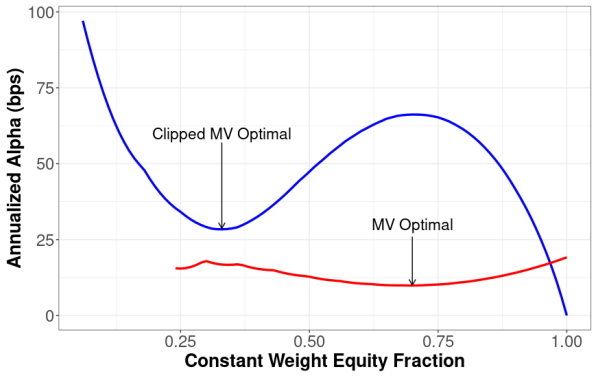

(b) *Annualized* α.

FIGURE 2.5: *Comparison of constant weight and MV optimal strategies in the historical market with 640,000 paths, expected blocksize of 24 months, and parameters from Table 2.1. Non-Pareto optimal points removed for each strategy.*

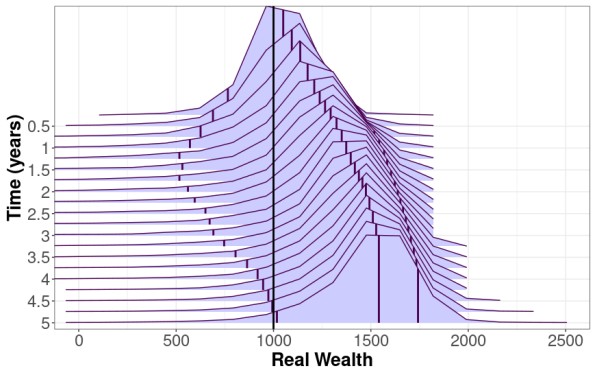

(a) *MV Optimal strategy.*

(b) *Clipped MV Optimal strategy.*

FIGURE 2.6: *Densities of real wealth over time based on data from Table 2.1 with* $W^* = 1736$*, 640,000 simulated historical market paths, expected blocksize of 60 months, and monthly rebalancing. Densities plotted quarterly. The black vertical line indicates the initial real wealth* $W_0 = 1000$*. The colored vertical lines in each density represent the 5th, 50th, and 95th percentiles of the distribution.*

above, it is much more likely here with the shorter expected blocksize. Compared to Figure 2.4, we observe even worse performance, especially for the MV Optimal strategy. In fact, here the MV Optimal strategy is worse than the Clipped MV Optimal strategy, except for cases with very high $std[W_T]$. We conjecture that this is because the aggressive nature of the MV Optimal strategy makes it more sensitive to model mis-specification.

Finally, in Figure 2.6 we show the evolution of the density of real wealth over time for the MV strategies, in the historical market with an expected blocksize of 60 months. While there are obvious general similarities with the corresponding Figure 2.3 in the synthetic market, we can observe that the real wealth densities are not as smooth here, and the chance of achieving higher real final wealth is increased. This is because the strategy of de-risking is still based on the assumption of constant interest rates, but here rates are stochastic (i.e. there are cases where the de-risked portfolio earns higher than expected returns, leading to higher final wealth).

2.2 Deficiencies of MV (Sharpe Ratio) Criteria

The apparent alpha generated by an MV optimal strategy comes from skewing the terminal wealth distribution. The right side of the distribution is cut-off (eliminating very large gains), and at the same time, an increase in left tail risk occurs (Lhabitant, 2000; Goetzmann et al., 2002; Forsyth and Vetzal, 2017a;b; 2019). As shown above in the synthetic market Figure 2.3, the left tail risk is somewhat reduced when constraints are imposed. This can also be seen as a natural consequence of the contrarian flavor of the strategy, which increases the weight in stocks when wealth decreases, and decreases the weight in stocks when wealth increases, i.e. buy when the market goes down, sell when the market goes up (see Figure 2.2). This implies that the investor is fully invested in bonds after stocks do well, and will not participate in further gains. On the other hand, the investor increases holdings in stocks when stocks perform poorly. This means that poor results can be expected if the market trends downward over the entire investment horizon. In this case (downward trending stocks in $[0,T]$) control (2.6) will generate a worse result than a constant weight strategy, which keeps at least some proportion of wealth always invested in bonds. The opposite is true at large values of wealth. In this case, an MV optimal strategy is always fully invested in bonds, while the constant weight strategy has some investment in stocks, and can participate in further equity market gains.

In summary, we can see that a major problem with dynamic MV (Sharpe ratio maximizing) strategies is that variance is a symmetric risk measure, which penalizes both the upside as well as the downside. An easy way to improve Sharpe ratios is to sell off the upside, which trivially reduces variance. On the other hand, somewhat counterintuitively, Sharpe ratio maximizing strategies also increase left tail risk, compared to a benchmark constant weight strategy.[8] This motivates us to consider below strategies that rely on a downside risk measure. In addition, the results above show the potential importance of imposing constraints: the MV Optimal strategy uses large amounts of leverage, which is unrealistic. However, below we specify constraints directly as part of the optimization problem, rather than being imposed afterwards in an ad hoc manner, as in (2.13).

3 Mean-Expected Shortfall Strategies

Based on our analysis of dynamic MV strategies, it seems clear that an asymmetric risk measure might be more useful than variance. If we directly target an asymmetric risk measure in the objective function, this will shape the distribution of the terminal wealth in a way which corresponds to our intuitive concept of risk (i.e. downside not upside).

Let $g(W_T)$ be the probability density function of terminal wealth W_T at $t = T$, and let

$$\int_{-\infty}^{W_\beta^*} g(W_T)\, dW_T = \beta, \tag{3.1}$$

so that $Prob[W_T > W_\beta^*] = 1 - \beta$. We can interpret W_β^* as the Value at Risk (VAR) at level β. The Expected Shortfall (ES) at level β is then

$$\mathrm{ES}_\beta = \frac{\int_{-\infty}^{W_\beta^*} W_T\, g(W_T)\, dW_T}{\beta}, \tag{3.2}$$

which is the mean of the worst β fraction of outcomes. Typically, $\beta \in \{.01, .05\}$. Note that the definition of ES in equation (3.2) uses the probability density of the final wealth distribution, not

[8] Note that in some cases an asset allocation strategy based on objective function (2.4) may be desirable based on the CDF of the final wealth distribution (see, e.g. Forsyth and Vetzal, 2019). However, this is best understood in terms of objective function (2.5), rather than Sharpe ratio maximization.

the density of *loss*. Hence, in our case a larger value of ES (i.e. a larger value of average worst case terminal wealth) is desired.[9] It will be convenient to use the equivalent definition of ES_β from Rockafellar and Uryasev (2000):

$$\mathrm{ES}_\beta = \sup_{W^*} E\left[W^* + \frac{\min(W_T - W^*,0)}{\beta}\right] . \tag{3.3}$$

As a measure of reward, we will simply use expected wealth $E[W_T]$, denoted by EW. Since ES_β and EW are conflicting measures, we find Pareto optimal strategies by using a scalarization parameter $\kappa > 0$ and then maximizing the objective function

$$\mathrm{ES}_\beta + \kappa \mathrm{EW} . \tag{3.4}$$

Varying κ traces out an efficient frontier in the $(\mathrm{EW}, \mathrm{ES})$ plane.

We will determine optimal strategies which are discretely rebalanced with no-shorting and no-leverage constraints. We next outline our assumptions concerning the stochastic processes of the underlying investments and some notational conventions.

4 Investment Market

We assume that the investor has access to two funds: a broad market stock index fund and a constant maturity bond index fund. The investment horizon is T. Let S_t and B_t respectively denote the real amounts invested in the stock index and the bond index respectively. In general, these amounts will depend on the investor's strategy as well as changes in the real unit prices of the assets. In the absence of an investor determined control (i.e. cash injections or rebalancing), all changes in S_t and B_t result from changes in asset prices.

We model the stock index as following a jump diffusion process. Let $S_{t^-} = S(t-\epsilon), \epsilon \to 0^+$, i.e. t^- is the instant of time before t, and let ξ^s be a random jump multiplier. When a jump occurs, $S_t = \xi^s S_{t^-}$. We assume that $\log(\xi^s)$ follows a double exponential distribution (Kou, 2002; Kou and Wang, 2004). The probability of an upward jump is p_u^s, while $1 - p_u^s$ is the chance of a downward jump. The density function for $y = \log(\xi^s)$ is

$$f^s(y) = p_u^s \eta_1^s e^{-\eta_1^s y} \mathbf{1}_{y \geq 0} + (1 - p_u^s)\eta_2^s e^{\eta_2^s y} \mathbf{1}_{y<0}. \tag{4.1}$$

Define

$$\kappa_\xi^s = E[\xi^s - 1] = \frac{p_u^s \eta_1^s}{\eta_1^s - 1} + \frac{(1 - p_u^s \eta_2^s}{\eta_2^s + 1} - 1. \tag{4.2}$$

In the absence of control,

$$\frac{dS_t}{S_{t^-}} = \left(\mu^s - \lambda_\xi^s \kappa_\xi^s\right)\, dt + \sigma^s\, dZ^s + d\left(\sum_{i=1}^{\pi_t^s} (\xi_i^s - 1)\right), \tag{4.3}$$

where μ^s is the (uncompensated) drift rate, σ^s is the diffusive volatility, Z^s is a Brownian motion, π_t^s is a Poisson process with positive intensity parameter λ_ξ^s, and ξ_i^s are i.i.d. positive random variables having distribution (4.1). Moreover, ξ_i^s, π_t^s, and Z^s are assumed to all be mutually independent.

In addition, we directly model the returns of the constant maturity bond index as a stochastic process, following MacMinn et al. (2014) and Lin et al. (2015). As in MacMinn et al. (2014), we

[9] The negative of ES is often called Conditional Value at Risk (CVAR).

assume that the constant maturity bond index follows a jump diffusion process, similar in form to the process assumed above for the stock index. In particular, $B_{t^-} = B(t-\epsilon), \epsilon \to 0^+$. In the absence of control, B_t evolves as

$$\frac{dB_t}{B_{t^-}} = \left(\mu^b - \lambda^b_\xi \kappa^b_\xi + \mu^b_c \mathbf{1}_{\{B_{t^-}<0\}}\right) dt + \sigma^b \, dZ^b + d\left(\sum_{i=1}^{\pi^b_t} (\xi^b_i - 1)\right), \tag{4.4}$$

where the terms in equation (4.4) are defined analogously to equation (4.3). In particular, π^b_t is a Poisson process with positive intensity parameter λ^b_ξ, and ξ^b_i has distribution

$$f^b(y = \log \xi^b) = p^b_u \eta^b_1 e^{-\eta^b_1 y} \mathbf{1}_{y \geq 0} + (1 - p^b_u) \eta^b_2 e^{\eta^b_2 y} \mathbf{1}_{y<0}, \tag{4.5}$$

and $\kappa^b_\xi = E[\xi^b - 1]$. ξ^b_i, π^b_t, and Z^b are assumed to all be mutually independent. The term $\mu^b_c \mathbf{1}_{\{B_{t^-}<0\}}$ in equation (4.4) represents an additional cost of borrowing ($B_t < 0$), i.e. a spread between borrowing and lending rates. We assume that the diffusive components of S_t and B_t are correlated, i.e. $dZ^s \cdot dZ^b = \rho_{sb} \, dt$. However, the jump process terms for these two indexes are assumed to be mutually independent.[10]

By using jump processes and random interest rates, we have generalized the environment considered in Section 2. In principle, equations (4.3) and (4.4) could be extended further to include stochastic volatility. However, Ma and Forsyth (2016) have shown that stochastic volatility is unimportant if the time horizon of the investment is larger than the mean reversion time of the volatility process. Based on historical data, the half-life of a volatility shock is 1-2 months (Ma and Forsyth, 2016). That said, we will conduct tests below similarly to Section 2.1: we will determine the optimal controls using the parametric model based on equations (4.3) and (4.4), and then apply these controls both in the synthetic market (i.e. with Monte Carlo simulations of the same parametric model) and in the historical market, i.e. using resampled historical data, which is devoid of any specific assumptions about the underlying stochastic processes.

We define the investor's total wealth at time t as $W_t \equiv S_t + B_t$. We generally impose the constraints that (assuming solvency) shorting stock and using leverage (i.e. borrowing) are not allowed, However, in some of our examples we will allow limited use of leverage. In such cases, we assume that the cost of borrowing is the return of the constant maturity bond index plus the spread component μ^b_c.

5 Notational Conventions

We specify a set of discrete rebalancing times $\mathcal{T}$

$$\mathcal{T} = \{t_0 = 0 < t_1 < t_2 < \ldots < t_M = T\} \tag{5.1}$$

where it is assumed that $t_i - t_{i-1} = \Delta t = T/M$ is constant for simplicity. To reduce subscripts, we will sometimes use the notation $S_t \equiv S(t), B_t \equiv B(t)$ and $W_t \equiv W(t)$. More specifically, let the inception time of the investment be $t_0 = 0$. At each time t_i, $i = 0, 1, \ldots, M-1$, the investor rebalances the portfolio. At $t_M = T$, the portfolio is liquidated.

Given a time dependent function $f(t)$, we will use the shorthand notation $f(t_i^+) \equiv \lim_{\epsilon \to \to 0^+} f(t_i + \epsilon)$ and $f(t_i^-) \equiv \lim_{\epsilon \to 0^+} f(t_i - \epsilon)$.

[10] See Forsyth (2020b) for a discussion of the evidence for stock and bond price jump independence.

We assume no taxes are triggered on rebalancing.[11] We also assume that transaction costs are small enough to be ignored in our analysis. This is not unreasonable as we assume discrete and relatively infrequent rebalancing, and we can also imagine that the investor holds units of a large pooled contract. In addition, the basic underlying investments are assumed to be broad stock and bond index ETFs, which are very liquid. This then implies that the condition

$$W(t_i^+) = W(t_i^-) \tag{5.2}$$

holds.[12]

The multi-dimensional controlled underlying process is denoted $X(t) = (S(t), B(t))$, with $t \in [0,T]$. The realized state of the system is $x = (s,b)$. Let the rebalancing control $p_i(\cdot)$ be the fraction invested in the stock index at rebalancing date t_i, i.e.

$$p_i\left(X(t_i^-)\right) = p\left(X(t_i^-),t_i\right) = \frac{S(t_i^+)}{S(t_i^+) + B(t_i^+)}. \tag{5.3}$$

The controls depend on the state of the investment portfolio before the rebalancing occurs, i.e. $p_i(\cdot) = p\left(X(t_i^-),t_i)\right) = p\left(X_i^-, t_i\right)$, $t_i \in \mathcal{T}$, the set of rebalancing times. We determine the optimal strategies amongst all strategies with constant wealth (before and after rebalancing), so that

$$\begin{aligned} p_i(\cdot) &= p(W(t_i^+), t_i) \\ W(t_i^+) &= S(t_i^-) + B(t_i^-) \\ S(t_i^+) &= S_i^+ = p_i(W_i^+)\, W_i^+ \\ B(t_i^+) &= B_i^+ = (1 - p_i(W_i^+))\, W_i^+. \end{aligned} \tag{5.4}$$

Let $\mathcal{Z}$ represent the set of admissible values of the control $p_i(\cdot)$. An admissible control $\mathcal{P} \in \mathcal{A}$, where $\mathcal{A}$ is the admissible control set, can be written as

$$\mathcal{P} = \{p_i(\cdot) \in \mathcal{Z} \ : \ i = 0, \ldots, M-1\}. \tag{5.5}$$

No-shorting and no-leverage constraints are imposed by specifying

$$\mathcal{Z} = [0,1]. \tag{5.6}$$

Finally, we define $\mathcal{P}_n \equiv \mathcal{P}_{t_n} \subset \mathcal{P}$ as the tail of the set of controls in $[t_n, t_{n+1}, \ldots, t_{M-1}]$, i.e.

$$\mathcal{P}_n = \{p_n(\cdot), \ldots, p_{M-1}(\cdot)\}. \tag{5.7}$$

6 Problem Definition

We now specify the EW-ES problem which was discussed informally in Section 3. Since expected wealth (EW) and expected shortfall (ES) are conflicting measures, we use a scalarization technique to find the Pareto points for this multi-objective optimization problem. For a given scalarization parameter $\kappa > 0$, we seek the control $\mathcal{P}_0$ that maximizes

$$\mathrm{ES}_\beta(X_0^-, t_0^-) + \kappa \mathrm{EW}(X_0^-, t_0^-)\ . \tag{6.1}$$

[11] If the contract is held in a tax-advantaged savings account, then no taxes are paid on rebalancing. In addition, in Canada rebalancing can occur without triggering taxes in a corporate class mutual fund. In the US, ETFs can defer taxes on rebalancing through the use of *heartbeat* trades (Moussawi et al., 2022).

[12] Transaction costs can be incorporated with increased computational cost. In addition, large liquid ETFs have very low transaction costs, which have little effect with infrequent trading (van Staden et al., 2018).

More precisely, we define the pre-commitment EW-ES problem ($PCEE_{t_0}(\kappa)$) problem in terms of the value function $J(s,b,t_0^-)$, where we use the definition of ES_β from equation (3.3).

$$(PCEE_{t_0}(\kappa)): \qquad J\left(s,b,t_0^-\right) = \sup_{\mathcal{P}_0\in\mathcal{A}} \sup_{W^*} \left\{ E_{\mathcal{P}_0}^{X_0^+,t_0^+} \left[W^* + \frac{1}{\beta}\min(W_T - W^*, 0) \; + \; \kappa W_T \; \middle| X(t_0^-) = (s,b) \right] \right\} \tag{6.2}$$

$$\text{subject to } \begin{cases} (S_t, B_t) \text{ follow processes (4.3) and (4.4)}; \quad t \notin \mathcal{T} \\ W_\ell^+ = S_\ell^- + B_\ell^- \; ; \; X_\ell^+ = (S_\ell^+, B_\ell^+) \\ S_\ell^+ = p_\ell(\cdot) W_\ell^+ \; ; \; B_\ell^+ = (1 - p_\ell(\cdot)) W_\ell^+ \\ p_\ell(\cdot) \in \mathcal{Z}(W_\ell^+, t_\ell) \\ \ell = 0, \ldots, M \; ; \; t_\ell \in \mathcal{T} \end{cases} . \tag{6.3}$$

Interchange the sup sup [13] in equation (6.2), so that value function $J\left(s,b,t_0^-\right)$ can be written as

$$J\left(s,b,t_0^-\right) = \sup_{W^*} \sup_{\mathcal{P}_0\in\mathcal{A}} \left\{ E_{\mathcal{P}_0}^{X_0^+,t_0^+} \left[W^* + \frac{1}{\beta}\min(W_T - W^*, 0) + \kappa W_T \middle| X(t_0^-) = (s,b) \right] \right\} . \tag{6.4}$$

Appendix B notes that the control determined from the pre-commitment policy (6.4) at time zero is identical to the control determined from a time consistent linear shortfall policy. However, as discussed in Section 2, the pre-commitment strategy is appropriate for a financial product purchased by a retail customer, and we will not discuss the induced time consistent policy further. We use the method described in Forsyth (2020a) to solve Problem 6.2. For the convenience of the reader, we provide a brief description of this method in Appendix C.

7 Parameter Estimates and Investment Scenario

We estimate the parameters for the double exponential jump diffusion processes (4.3) and (4.4) using the data described in Section 2.1. Recall that this is monthly data from 1926 to 2019 for inflation-adjusted returns for the CRSP value-weighted total return and 30-day US T-bill indexes. We use the threshold technique (Mancini, 2009; Cont and Mancini, 2011; Dang and Forsyth, 2016). Table 7.1 shows the results of calibrating the models to the historical data. The correlation ρ_{sb} is computed by removing any returns which occur at times corresponding to jumps in either series, and then using the sample covariance. Further discussion of the validity of assuming that the stock and bond jumps are independent is given in Forsyth (2020b).

Table 7.2 shows our base case investment scenario. We consider $T = 5$ years, with an initial investment of 1000. Rebalancing occurs discretely, on a quarterly basis. Our default constraints (see equation (5.5)) are $\mathcal{Z} = [0,1]$, i.e. no shorting and no-leverage. This clearly implies that no trading occurs under insolvency, which can only occur if there is a jump to zero for both assets. In Table 7.2 we have set the borrowing spread $\mu_c^b = .02$. There is no borrowing in our base case, but we do allow leverage in some of our later examples.

[13] Let $F = \sup_{(a,b)\in A\times B} f(a,b)$, then $\forall \epsilon > 0$, $\exists (a^*,b^*) \in A \times B$, s.t. $f(a^*,b^*) > F - \epsilon$. Then $F \geq \sup_{a\in A}\sup_{b\in B} f(a,b) \geq \sup_{b\in b} f(a^*,b) \geq f(a^*,b^*) > F - \epsilon$. Hence, $\epsilon \to 0$ implies $\sup_{a\in A}\sup_{b\in B} f(a,b) = F$. Similarly $\sup_{b\in B}\sup_{a\in A} f(a,b) = F$.

CRSP	μ^s	σ^s	λ^s	u^s	η_1^s	η_2^s	ρ_{sb}
	0.0877	0.1459	0.3191	0.2333	4.3608	5.504	0.08228
30-day T-bill	μ^b	σ^b	λ^b	u^b	η_1^b	η_2^b	ρ_{sb}
	0.0045	0.0130	0.5106	0.3958	65.85	57.75	0.08228

TABLE 7.1: *Estimated annualized parameters for double exponential jump diffusion model. Value-weighted CRSP index, 30 day US T-bill index deflated by the CPI. Sample period 1926:1 to 2019:12.*

Investment horizon T (years)	5.0
Equity market index	CRSP Cap-weighted index (real)
Bond index	30-day T-bill (US) (real)
Initial portfolio value W_0	1000
Rebalancing times	$t = 0, 0.25, 0.5, \ldots, 4.5, 4.75$
Equity fraction range	$\mathcal{Z} \in [0,1]$
Borrowing spread μ_c^b	0.02
Rebalancing interval (years)	0.25
Market parameters	See Table 7.1

TABLE 7.2: *Input data for examples.*

8 EW-ES alpha

Suppose we have a set of points ($E[W_T^{opt}]$,ES_β^{opt}) on the efficient frontier, determined using the optimal policy from Problem 6.2. In addition, we also compute points $(E[W_T^p], ES_\beta^p)$ with a constant weight strategy p where p is the (constant) fraction in equities, reset at each discrete rebalancing date. We can write the efficient EW-ES frontier, determined by solving Problem 6.3, as a function $F(\cdot)$, i.e.

$$E[W_T^{opt}] = F(ES_\beta^{opt}) \ . \tag{8.1}$$

For each value of p, we then determine

$$E[W_T^{opt,p}] = F(ES_\beta^p) \ , \tag{8.2}$$

which is the expected wealth from the optimal strategy having the same risk (measured by ES) as the constant weight strategy with weight p. The annualized alpha for this value of p is then

$$\alpha^p = \frac{\log(E[W_T^{opt,p}]) - \log(E[W_T^p])}{T} \ . \tag{8.3}$$

Remark 8.1 (Optimal Efficient Frontier)**.** *In practice, we compute the efficient frontier function $F(ES_\beta^{opt})$ at a finite number of points, and approximate the efficient frontier function $F(\cdot)$ by linear interpolation.*

9 Base Case: Synthetic Market

We compute and store the control determined by solving Problem 6.3 in the synthetic market using the parameters from Table 7.1 for the investment scenario given in Table 7.2. We then use the

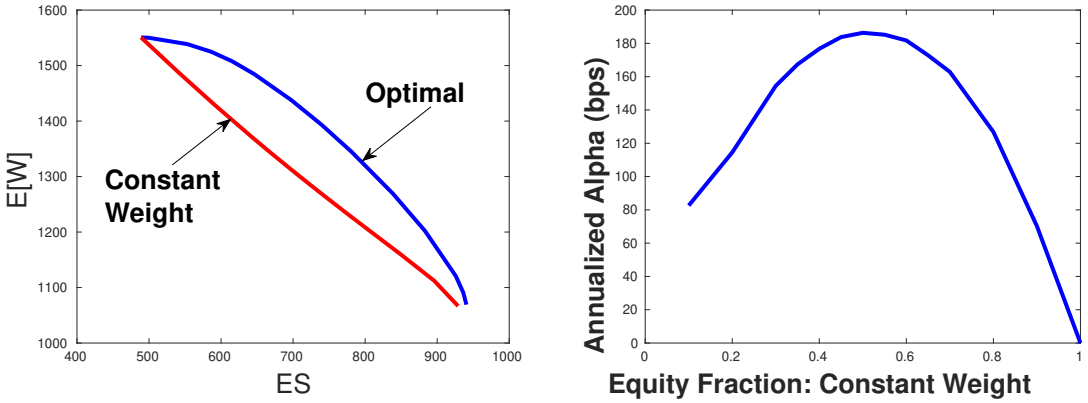

(a) *Efficient frontiers, constant weight compared to optimal control.*

(b) *Annualized α, optimal vs. constant weight strategy.*

FIGURE 9.1: *Scenario from Table 7.2, synthetic market. Optimal strategy determined by solving Problem 6.3 in the synthetic market, with parameters in Table 7.1. Control stored and then used to compute the final results with* 2.56×10^6 *Monte Carlo simulations.*

stored control in Monte Carlo simulations to generate summary statistics. Figure 9.1(a) shows the EW-ES frontiers for the optimal strategy (6.3) and for the constant weight strategy. [14] as a function of EW. Figure 9.1(b) shows the annualized alpha compared with the benchmark constant weight strategy, based on equation (8.3). The annualized alpha reaches a maximum of about 180 bps, when compared to a benchmark 60:40 stock bond portfolio.

Figure 9.2 shows the probability density of the internal rate of return for the optimal strategy in the synthetic market, for $\kappa = 1.0$. We also specify $W^* = 788$, which results in an ES approximately equal to that of a constant weight strategy with $p = 0.6$. Note the rapid decrease in the density near the 5^{th} percentile and the long right tail. We can see that this density protects the downside (in terms of ES) and maximizes EW through the right skew.

Figure 9.3 compares the cumulative distribution functions (CDFs) for the optimal EW-ES strategy ($\kappa = 1.0$ and $W^* = 786$) and a constant weight strategy with $p = 0.6$. Both strategies have approximately the same ES(5%). The rapid decrease in the CDF for the optimal EW-ES strategy (near $W = 800$) results in the same tail risk as for the constant proportion strategy. The EW-ES policy then gives up some performance between final wealth values in the range [800,1260]. However, the EW-ES strategy then gains in performance for $W_T > 1260$. This results in a larger expected final wealth value for the same tail risk. Note that there is no magic bullet in terms of strategies. If we constrain both strategies to have same left tail risk, then the EW-ES policy gives up some performance in the probability ranges from [0.05,0.55], in order to generate superior performance in the upper 45% of the outcomes. In other words, ignoring the left tail behavior (which is roughly the same), then the EW-ES strategy has slightly worse performance in 50% of the outcomes, which is counterbalanced by a large outperformance in 45% of the outcomes. This is essentially the opposite strategy compared to the multi-period pre-commitment Sharpe ratio maximizing (or quadratic shortfall) policy (Forsyth and Vetzal, 2017a;b; 2019). To emphasize this, Figure 9.4 provides another comparison of the constant weight (panel (a)) and EW-ES strategies (panel (b), showing the

[14] Detailed tables containing the results used to generate Figure 9.1 are provided in Appendix E. Note that here expected terminal wealth as a function of ES is downward-sloping because higher ES represents lower risk. This is in contrast to cases like Figure 2.1(a) where risk (measured by standard deviation) was increasing along the horizontal axis.

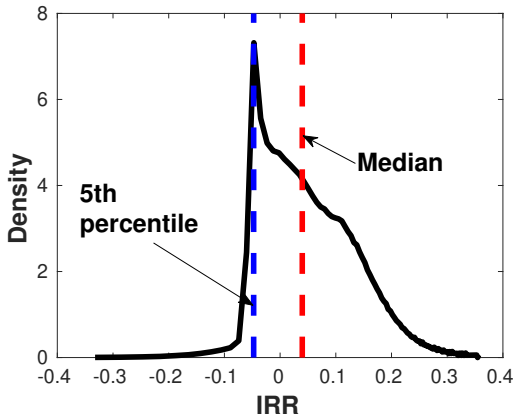

FIGURE 9.2: *Density of internal rate of return (IRR). Scenario from Table 7.2, synthetic market. Optimal strategy determined by solving Problem 6.3 with* $\kappa = 1.0$ *and* $W^* = 788$ *in the synthetic market, with parameters in Table 7.1. Control stored and then used to compute the final results with* 6.4×10^5 *Monte Carlo simulations. The IRR has a mean of .0724 and a median of .0396.*

densities of real wealth over time. Both strategies have approximately the same ES, but the EW-ES strategy has a large right skew. This is in stark contrast to the large left skew seen for the MV Optimal and Clipped MV Optimal strategies in Figure 2.3.

Figure 9.5 shows the percentiles of the fraction invested in stocks over time and the percentiles of total wealth for the EW-ES strategy. The median fraction invested in stocks is quite high, between $.80-.90$, but there is a large gap between the median and 5^{th} percentile, indicating that the strategy reacts very strongly to decreasing wealth, as shown in Figure 9.6 which provides a heatmap of the optimal controls. Initially, with $W_0 = 1000$ at $t = 0$ the equity weight is about 0.85. If the portfolio does well, the fraction in stocks remains high. However, if the portfolio performs poorly, the fraction in stocks is rapidly reduced. Further poor returns will result in a large fraction in bonds (in the blue zone of the heatmap) so as to protect the ES. Conversely, if stocks do well, the strategy will increase the allocation to stocks. Hence this is a momentum-type policy. However, if the portfolio suffers a very large sudden loss (basically jumping down through the blue zone in Figure 9.6), the strategy will once again increase the stock position. This can be seen as a last ditch attempt to recover, by taking a large equity position and hoping for strong returns. Similar behavior happens in the case of very poor investment returns for the MV Optimal strategy, as seen in Figure 2.2. On the whole, the momentum type of strategy shown in Figure 9.6 displays a much more varied pattern across both real wealth and time compared to the contrarian MV Optimal strategy in Figure 2.2.

10 Historical market

We next present historical market results. Recall that the procedure used here is to determine the optimal controls based on an assumed parametric model (in this case the double exponential jump diffusion processes (4.3) and (4.4)), and then apply the controls to resampled historical return data as described above in Section 2.1.

Figure 10.1 shows the efficient EW-ES frontiers and the apparent alpha for the bootstrapped historical market. The results are shown for various expected blocksizes (Blk) in years.[15] We

[15] Appendix F provides detailed tables of the results used to generate Figure 10.1. To save space, we only provide the results for the case with an expected blocksize of 5 years.

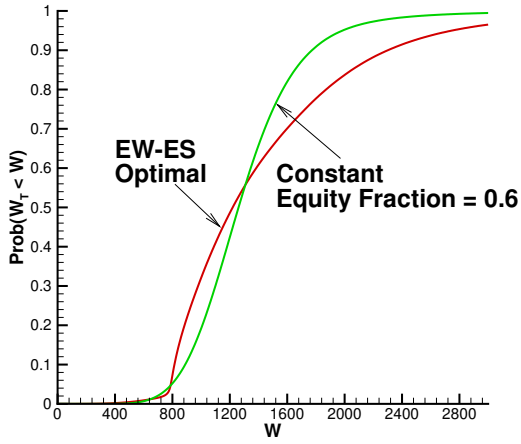

FIGURE 9.3: *Cumulative distribution functions of terminal real wealth for the optimal EW-ES* ($\kappa = 1.0, W^* = 786$) *strategy compared to a constant weight strategy* ($p = 0.6$). *Both strategies have approximately the same ES* $\in$ (696,698). *Scenario from Table 7.2, synthetic market. Optimal strategy determined by solving Problem 6.3 in the synthetic market, with parameters in Table 7.1. Control stored and then used to compute the final results with* 2.56×10^6 *Monte Carlo simulations.*

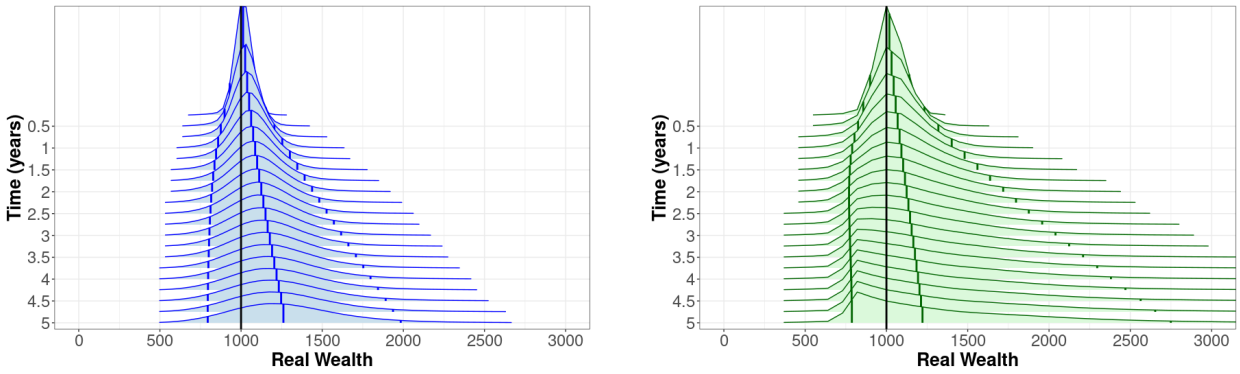

(a) *Constant weight strategy,* $p = 0.6$.

(b) *EW-ES strategy,* $\kappa = 1$, $W^* = 786$.

FIGURE 9.4: *Densities of real wealth over time for constant weight* ($p = 0.6$) *and EW-ES* ($\kappa = 1$, $W^* = 786$) *strategies. Investment scenario from Table 7.2. 640,000 simulated synthetic market paths with quarterly rebalancing. Densities plotted quarterly. The black vertical line indicates the initial real wealth* $W_0 = 1000$. *The colored vertical lines in each density represent the 5th, 50th, and 95th percentiles of the distribution.*

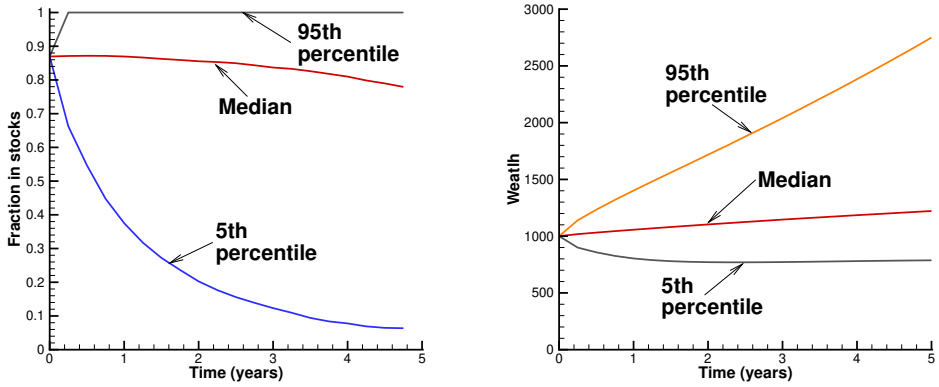

(a) *Percentiles fraction in stocks.* (b) *Percentiles wealth*

FIGURE 9.5: *Scenario in Table 7.2, synthetic market. Optimal strategy determined by solving Problem 6.3 in the synthetic market, parameters in Table 7.1. Control stored and then used to compute the final results with* 6.4×10^5 *Monte Carlo simulations. Parameters based on CRSP inflation adjusted data, 1926:1-2019:12.* $\kappa = 1.0$.

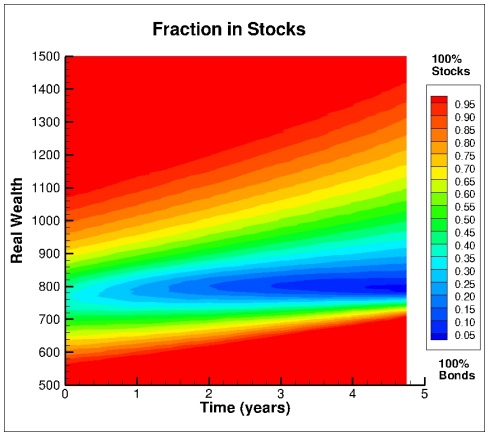

FIGURE 9.6: *Heatmap of fraction in stocks. Scenario from Table 7.2. Optimal strategy determined by solving Problem 6.3 with* $\kappa = 1.0$ *and* $W^* = 786$ *in the synthetic market, with parameters in Table 7.1.*

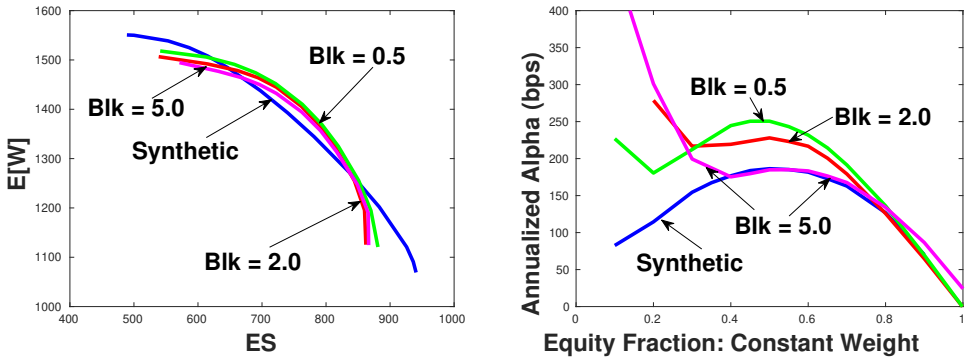

(a) *Efficient frontiers, historical market. Synthetic market frontier also shown.*

(b) *Annualized α, optimal vs. constant weight strategy, historical market. Synthetic market result also shown.*

FIGURE 10.1: *Optimal strategy determined by solving Problem 6.3 in the synthetic market, parameters in Table 7.1. Control stored and then tested in bootstrapped historical market. Non-Pareto points eliminated. Expected blocksize (Blk, years) used in the bootstrap resampling method also shown.*

also show the synthetic market results as well for comparison. Figure 10.1(a) indicates that the control computed using the parametric model (synthetic market) performs quite well in the historical market, except for large values of ES. This indicates that the optimal controls are fairly robust to parametric model misspecification, over a wide range of interesting values of ES (i.e. ES in the range $650 - 850$). The results are also fairly insensitive to the choice of expected blocksize.

On the other hand, the apparent alpha based on the synthetic market control, tested in the historical market, is generally superior to the synthetic market controls tested in the synthetic market, especially for the interesting constant weight benchmark portfolios for p in the range $0.4 - 0.6$. However, while Figure 10.1(a) shows very little sensitivity to expected blocksize, Figure 10.1(b) is sensitive to blocksize. This is because the constant weight strategies are much more sensitive to the blocksize in the resampling algorithm compared to the EW-ES strategies, suggesting that constant weight strategies are less robust than EW-ES strategies.

The apparent alpha in Figure 10.1(b) is surprisingly large for small values of constant p for the benchmark portfolio. This results from the very poor performance of constant weight portfolios for small equity weights in the historical market.[16] This can be verified by examining the case with $p = 0.0$ in Table F.2, which has an ES of about 764. In other words, short term T-bills have an expected loss in the worst 5% of cases of about 25% in real terms over five years. This is due to negative real short term interest rates in times of inflation.

Figure 10.2 depicts the evolution of the densities of real wealth over time in the historical market with an expected blocksize of 5 years, for both the constant weight ($p = 0.6$) and EW-ES strategies. The patterns observed here are quite similar to those seen above in Figure 9.4 in the synthetic market, although the results here are clearly not as smooth. The main features of the EW-ES strategy are preserved quite well in the historical market, as there is clear downside protection and a strong right skew indicating upside potential.

[16] Recall that we noted similar behavior for the Clipped MV Optimal strategy in Figure 2.4.

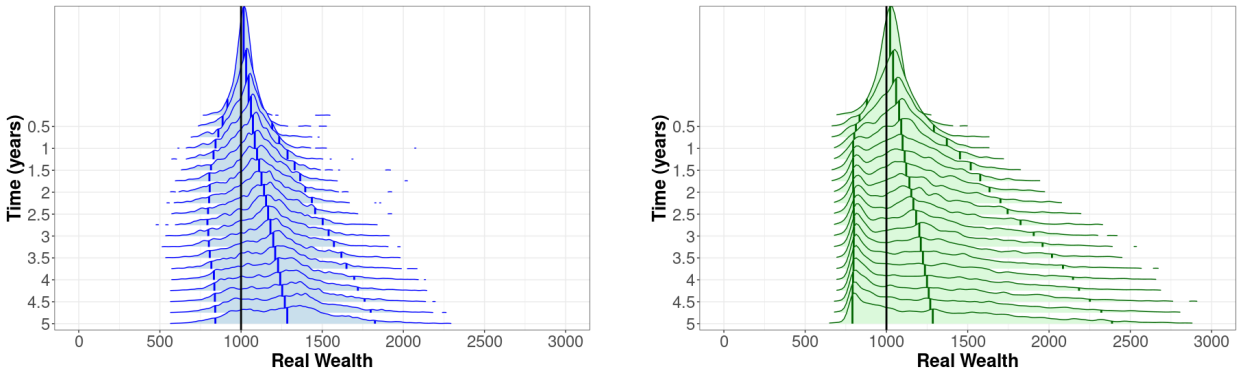

(a) *Constant weight strategy,* $p = 0.6$.

(b) *EW-ES strategy,* $\kappa = 1$, $W^* = 786$.

FIGURE 10.2: *Densities of real wealth over time for constant weight* ($p = 0.6$) *and EW-ES* ($\kappa = 1$, $W^* = 786$) *strategies. Investment scenario from Table 7.2. 640,000 simulated historical market paths with expected blocksize 60 months and quarterly rebalancing. Densities plotted quarterly. The black vertical line indicates the initial real wealth* $W_0 = 1000$. *The colored vertical lines in each density represent the 5th, 50th, and 95th percentiles of the distribution.*

11 Alternative Investment Scenarios

In this section we expore the effects of some departures from the base case assumptions presented in Table 7.2. We first study the rebalancing frequency and the potential use of leverage, and then turn to a shorter investment horizon. We concentrate here exclusively on the synthetic market.

11.1 Rebalancing Frequency and Leverage

Figure 11.1(a) shows the effect of changing the rebalancing frequency on the efficient frontiers (optimal strategy, synthetic market). There is very little discernable effect of changing the rebalancing frequency from one month to six months.

Recall from Appendix A that continuous trading in the stock and bond with no leverage can replicate a covered call strategy. Since it appears from Figure 11.1(a) that the effect of the rebalancing frequency is small, this suggests that our discretely rebalanced portolio can approximate a covered call strategy.

However, there may be other strategies which use options that allow the investor to use leverage with limited downside. Formally, these other option strategies would require continuous trading. However, we can approximate the effect of other strategies which include options by allowing the use of leverage.

Figure 11.1(b) shows the results for allowing leverage, both with and without a borrowing spread.[17] For values of ES > 750, which is arguably the region of interest, the gain in using leverage is quite small. Figure 11.1(b) suggests that allowing the use of leverage is not very important for such values of ES. Hence there is no particular advantage to taking direct positions in options as compared to dynamic trading, unless the value of ES considered is quite small (i.e. the investor wants a very risky strategy).

[17] We enforce the constraint that trading ceases if insolvency occurs, and debt accumulates at the borrowing rate. However, the probability of this occurring over five years is very small.

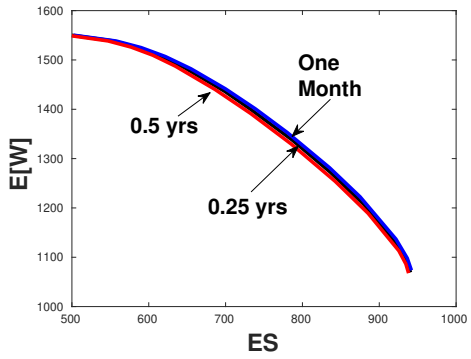

(a) *Effect of rebalancing frequency*

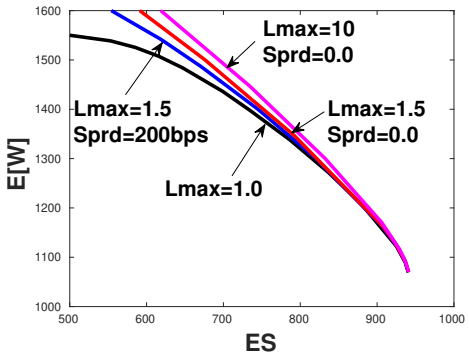

(b) *Effect of leverage*

FIGURE 11.1: *Synthetic market, optimal strategy, effect of rebalancing frequency and use of leverage. Lmax is the maximum value of p allowed. Optimal strategy determined by solving Problem 6.3 in the synthetic market, parameters in Table 7.1.*

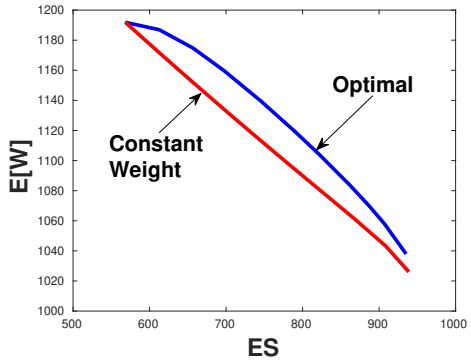

(a) *Efficient frontiers, constant weight compared to optimal control.*

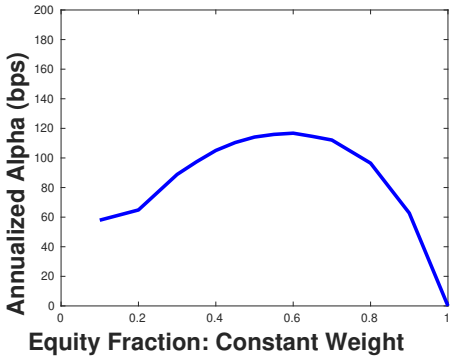

(b) *Annualized α, optimal vs. constant weight strategy.*

FIGURE 11.2: *Scenario in Table 7.2, synthetic market. Optimal strategy determined by solving Problem 6.3 in the synthetic market, parameters in Table 7.1, with the exception that $T = 2.0$ years, and the portfolio is rebalanced monthly. Control stored and then used to compute the final results with 6.4×10^5 Monte Carlo simulations.*

11.2 Two Year Time Horizon

It is also interesting to examine the effect of a shorter investment horizon. Up to now, we have been using $T = 5.0$ years. Figure 11.2 shows the results for the efficient EW-ES frontier and the apparent alpha for $T = 2.0$ years in the synthetic market with monthly rebalancing. It is interesting to see that even for this comparatively short time period, the apparent annualized alpha is roughly 120 bps, compared to a constant weight strategy with $p = 0.6$. This compares with a maximum alpha of about 180 bps for the $T = 5$ years case.

Figure 11.3 shows the heatmap of the optimal controls (fraction in stocks) for the $T = 2$ years case with $\kappa = 2.25$. In contrast to the 5-year horizon case in Figure 9.6, here it can be seen that if stocks drop in value very early on, the optimal strategy is to quickly shift out of stocks into bonds. The optimal strategy ($T = 2$ years) is to heavily invest in bonds when the total wealth approaches $W^* = 828$.

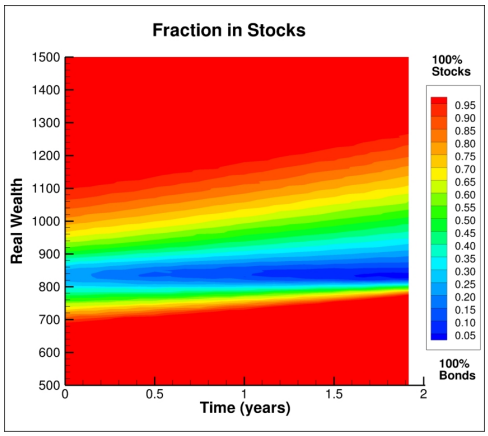

FIGURE 11.3: *Heat map of fraction in stocks, scenario in Table 7.2, parameters from Table 7.1. Here* $T = 2.0$ *years, monthly rebalancingm* $\kappa = 2.25$, $W^* = 828$, $ES(5\%) = 745$, $E[W_T] = 1140$. *Compare with Figure 9.6 (*$T = 5$ *years).*

12 Discussion

Sharpe ratio maximizing strategies exploit the symmetry of the standard deviation risk measure by selling off large gains. In some circumstances, such as saving for retirement in a DC savings account, multi-period Sharpe ratio maximization can in fact be useful. However, this is due to the fact that multi-period Sharpe ratio maximizing strategies are equivalent to an induced time consistent target based quadratic shortfall minimization (Vigna, 2014). It is this property, rather than Sharpe ratio maximization *per se*, which is desirable for DC plan investments.

In order to align the objective function more precisely with an investor's intuitive concept of risk, we propose using an asymmetric risk measure that is based on expected shortfall (ES). Although this strategy is formally a pre-commitment policy, we argue that the lack of time consistency in this case does not cause conceptual difficulties. We can think of this strategy as part of a packaged product sold to a retail customer who does not trade in the underlying securities during the lifetime of the contract. This packaged product is essentially a *black box* to the retail customer, who evaluates success or failure of this policy based on the initial investment and final value.

The optimal policy is based on determining the expected wealth, expected shortfall (EW-ES) frontiers. This strategy cuts off the left tail of the distribution (protecting worst case wealth outcomes), with a skewed right tail, which maximizes EW. This strategy might be regarded as having some of the features of value investing, where preservation of capital is given the highest priority. However, in this case, capital is preserved by dynamic trading rather than stock picking.

Note that Sharpe ratio maximizing strategies have a contrarian flavor. They can be summarized as *"buy stocks and sell bonds when stocks go down, sell stocks and buy bonds when stocks go up."* The EW-ES policy is essentially the opposite, sell stocks and buy bonds when stocks go down; buy stocks and sell bonds when stocks go up: *"cut your losses, ride your gains"*. This is essentially a momentum type of strategy.

A negative feature with this EW-ES strategy (which it shares with the constant proportion portfolio insurance strategy (CPPI) (Black and Jones, 1987)) is the possibility of cashing out, i.e. if stocks drop early in the investment process, the portfolio will move to largely bond holdings, with little possibility of recovery. This strategy will be particularly ineffective in the case of a large market drop, followed by a rapid recovery, such as the recent pandemic market crash and recovery. This is, however, a very unusual return sequence based on U.S. market history.

13 Conclusion

We have reviewed the known results concerning dynamic strategies which maximize the Sharpe ratio. These strategies essentially reduce risk by selling off the right tail of the distribution. While this approach may be desirable in some circumstances (e.g. when saving for retirement where a target-based strategy can be used to generate cash flows to replace employment income (Forsyth and Vetzal, 2019)), this may not be a suitable approach for all investors.

Denoting the expected terminal wealth by EW, and the expected shortfall by ES (ES is the mean of the worst β fraction of outcomes), we propose an objective function based on EW-ES criteria. Measuring risk by ES (expected shortfall) fundamentally means that the investor is concerned with the left tail risk. This amounts to preserving a desired minimum value of terminal wealth, with high probability. We determine the optimal EW-ES dynamic trading strategy using optimal stochastic control techniques, based on a parametric model for stock and bond processes, fit to historical data. We impose realistic constraints on the trading strategy: no-leverage, no-shorting, and infrequent rebalancing.

We define the alpha of this investment strategy relative to a benchmark by using ES as the measure of risk. We focus on medium-term investments (i.e. 2-5 years), where wealth preservation can be regarded as of high importance. Compared to a 60:40 stock-bond constant weight strategy, the optimal EW-ES policy gives an annualized alpha of about 180 bps over a 5-year investment horizon.

We emphasize again that optimal EW-ES strategies are fundamentally different compared to Sharpe ratio maximizing strategies. Optimal Sharpe ratio strategies are contrarian: increase the weight in stocks when stocks do poorly and decrease the weight in stocks when stocks to do well. EW-ES optimal strategies do the opposite: decrease the weight in stocks when stocks perform poorly, increase the weight in stocks when stocks do well. This policy protects the left tail, while taking advantage of the large possible gains emanating from the right tail of the distribution. Finally, it is worth noting that bootstrap resampling tests using historical data show that the optimal EW-ES policy is fairly robust to parametric model misspecification.

14 Acknowledgment

Forsyth's work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) grant RGPIN-2017-03760.

15 Disclosure

The authors report there are no competing interests to declare.

Appendices

A Replication of Covered Call and Cash-Secured Put Writing

This appendix outlines the assumptions required to ensure that dynamic trading in stocks and risk-less bonds can replicate covered call and cash-secured put writing, which are two popular approaches to generating apparent alpha (Lhabitant, 2000; Goetzmann et al., 2002; Ungar and Moran, 2009).

Assumption A.1. *Let $V(S,t)$ be the price of a call or put option, with S being the unit price of the stock.* [18] *We make the following assumptions:*

(i) Trading in the stock with price $S(t)$ and the riskless bond $B(t)$ occur continuously, with no market frictions.

(ii) Call and put option prices are convex.

(iii) Call and put options can be perfectly replicated by the portfolio of $\alpha(S,t)$ units stock and an amount $B(S,t)$ in the riskless bond which pays an interest rate of r, so that the total value of the replicating portfolio is

$$V(S,t) = \alpha S + B \tag{A.1}$$

with $\alpha = V_S$. At points where V_S does not exist (which can only occur at a countable number of points), we take the appropriate left or right limits.

(iv) V_S satisfies the bounds

$$\begin{aligned} 0 \le V_S \le 1 \quad & Call \\ -1 \le V_S \le 0 \quad & Put \end{aligned}$$

(v) The value of a call option at $S = 0$ is zero. The value of a European put at $S = 0$ is $Ke^{-r(T-t)}$, where K is the strike, and T is the expiry time. The value of an American put at $S = 0$ is K.

Some discussion concerning the conditions under which assumptions (i)-(iv) hold is given in Bergman et al. (1996). Suffice to say, all these assumptions are met in a Black-Scholes market, as used for the background example in Section 2.

Proposition A.1. *Under Assumptions A.1, covered call writing and cash-secured put writing can be replicated by a portfolio consisting of the underlying asset and riskless bond with non-negative amounts in both stock and bond, i.e. the portfolio satisfies the* no-shorting, no leverage *condition $0 \le p \le 1$ where p is the fraction of the portfolio wealth held in the risky asset.*

Proof. We consider first the case of a stock which does not pay dividends.
(a) **Covered Call**: Consider a covered call long one unit of stock, short one call option and long the option premium. By definition $S \ge 0$. At $t = 0$, the covered call portfolio Π^{cc} is

$$\begin{aligned} \Pi^{cc}(S,0) &= S \; - \overbrace{V(S,0)}^{\text{short call}} \; + \overbrace{V_0}^{\text{cash premium}} \\ V_0 &= V(S,0) \; . \end{aligned} \tag{A.2}$$

Replace $V(S,0)$ by the replicating portfolio

$$\begin{aligned} V(S,0) &= \alpha(S,0)S + B(S,0) \\ \alpha(S,0) &= V_S(S,0), \end{aligned} \tag{A.3}$$

where the amount in the riskless bond is

$$B(S,0) = V(S,0) - V_S(S,0)S. \tag{A.4}$$

[18] With an abuse of notation, in this appendix S and B are the unit prices of the stock and bond respectively, not the amounts invested in them.

Then equation (A.2) becomes

$$\Pi^{cc}(S,0) = S - (\alpha(S,0)S + B(S,0)) + V_0 \ . \tag{A.5}$$

In general for $t > 0$ we have

$$\Pi^{cc}(S,t) = S - \alpha(S,t)S - B(S,t) + V_0 e^{rt} \tag{A.6}$$

$$B(S,t) = V(S,t) - \alpha(S,t)S \ , \tag{A.7}$$

where r is the continuously compounded risk-free rate, and so

$$\Pi^{cc}(S,t) = \overbrace{S(1-\alpha)}^{\text{stock position}} \ + \ \overbrace{V_0 e^{rt} - B(S,t)}^{\text{cash}} \ . \tag{A.8}$$

Consider the line in the (V,S) plane

$$\alpha(S^*,t)S + C \ , \tag{A.9}$$

where the constant C is determined from

$$\alpha(S^*,t)S^* + C = V(S^*,t) \ . \tag{A.10}$$

The line (A.9) is then tangent to the option value $V(S,t)$ at $S = S^*$. Since V is convex, it lies at or above its tangent line everywhere, including at $S = 0$ where $V(0,t) = 0$, which implies that $C \leq 0$. Since this is true for any S^*, then $(-B(S,t)) = \alpha(S,t)S - V(S,t) \geq 0$. For a call $0 \leq \alpha \leq 1$, so from equation (A.8) the stock position is also non-negative.

(b) **Cash-Secured Put**: Now consider writing a cash-secured put with strike K. If the put is European-style, the writer is long Ke^{-rT} in cash, short one put option, and long the option premium. By put-call parity and the results from (a), the cash-secured put also has non-negative positions in both the cash and stock.

If the put is American-style, the writer deposits K with the broker. At $t = 0$, the cash-secured put portfolio $\Pi^{cp}(S,t)$ is

$$\Pi^{cp}(S,0) = K \ - \ \overbrace{V(S,0)}^{\text{short put}} \ + \ \overbrace{V_0}^{\text{cash premium}}$$

$$V_0 = V(S,0) \ . \tag{A.11}$$

Replace $V(S,0)$ by the replicating portfolio

$$V(S,0) = \alpha(S,0)S + B(S,0)$$

$$\alpha(S,0) = V_S(S,0) \ , \tag{A.12}$$

where the amount in the riskless bond is

$$B(S,0) = V(S,0) - V_S(S,0)S \ . \tag{A.13}$$

Then equation (A.11) becomes

$$\Pi^{cp}(S,0) = K - (\alpha(S,0)S + B(S,0)) + V_0 \tag{A.14}$$

In general, for $t > 0$ we have

$$\Pi^{cp} = Ke^{rt} - \alpha(S,t)S - B(S,t) + V_0 e^{rt} \tag{A.15}$$

$$B(S,t) = V(S,t) - \alpha(S,t)S \ , \tag{A.16}$$

or

$$\Pi^{cp} = \overbrace{-\alpha S}^{\text{stock position}} + \overbrace{V_0 e^{rt} + K e^{rt} - B(S,t)}^{\text{cash}} . \tag{A.17}$$

Consider the line in the (V,S) plane

$$\alpha(S^*,t)S + C , \tag{A.18}$$

where the constant C is determined by

$$\alpha(S^*,t)S^* + C = V(S^*,t) , \tag{A.19}$$

so that (A.18) is tangent to $V(S,t)$ at $S = S^*$.

Since V is convex, it lies at or above its tangent line everywhere, including at $S = 0$ where $V(0,t) = K$, which implies that $C \le K$. Hence $V(S^*,t) - \alpha(S^*,t)S^* = B(S^*,t) = C \le K$. This is true for any S^*, so that $(Ke^{rt} - B(S,t)) \ge 0$. Since $-1 \le \alpha \le 0$ for a put, the stock position is non-negative.

(c) **Dividends**: Now consider the case of a covered call written on a stock which pays a non-proportional dividend of $\min(D,S)$ at $t = t_d$. Let t_d^- be the instant before t_d. Then

$$\Pi^{cc}(S^-,t_d^-) = \overbrace{S^-(1-\alpha^-)}^{\text{stock position}} - \overbrace{B(S^-,t_d^-) + V_0 e^{rt}}^{\text{cash}} , \tag{A.20}$$

where by the arguments in (a) above the stock position is non-negative and $-B(S^-,t_d^-) = -B^- \ge 0$. At t_d^+, after the dividend is paid, the new stock and cash positions are

$$\begin{aligned} (\text{stock position})^+ &= (S^- - \min(S^-,D))(1-\alpha^-) \\ (\text{cash position})^+ &= -B^- + \min(S^-,D)(1-\alpha^-) \end{aligned} \tag{A.21}$$

From (a) above, $(1-\alpha^-) \ge 0$ and $-B^- \ge 0$, so both positions are nonnegative after the dividend is paid. After rebalancing the replicating portfolio after the dividend payment, the stock and bond positions remain non-negative.[19] The cash-secured put case can also be proven using similar arguments. □

B Induced Time Consistent Policy

Noting that the inner supremum in equation (6.4) is a continuous function of W^*, define

$$\mathcal{W}^*(s,b) = \arg\max_{W^*} \left\{ \sup_{\mathcal{P}_0 \in \mathcal{A}} \left\{ E_{\mathcal{P}_0}^{X_0^+,t_0^+} \left[W^* + \frac{1}{\beta}\min(W_T - W^*, 0) + \kappa W_T \middle| X(t_0^-) = (s,b) \right] \right\} \right. . \tag{B.1}$$

We refer the reader to Forsyth (2020a) for an extensive discussion concerning pre-commitment and time consistent ES strategies. We summarize the relevant results from that research here. Denote the investor's initial wealth at t_0 by W_0^-. Then we have the following result:

Proposition B.1 (Pre-commitment strategy equivalence to a time consistent policy for an alternative objective function)**.** *The pre-commitment EW-ES strategy $\mathcal{P}^*$ determined by solving $J(0,W_0,t_0^-)$*

[19] Since $V(S^-,t) = V(S^- - \min(D,S^-),t^+)$, then $V_S(S^-,t) = V_S(S^+,t^+)$ and we don't actually need to rebalance across the dividend date.

(with $\mathcal{W}^*(0,W_0^-)$ from equation (B.1)) is the time consistent strategy for the equivalent problem TCEQ (with fixed $\mathcal{W}^*(0,W_0^-)$), with value function $\tilde{J}(s,b,t)$ defined by

$$TCEQ_{t_n}(\kappa\beta): \qquad \tilde{J}\left(s,b,t_n^-\right) = \sup_{\mathcal{P}_n\in\mathcal{A}}\left\{E_{\mathcal{P}_n}^{X_n^+,t_n^+}\left[\min(W_T-\mathcal{W}^*(0,W_0^-),0) \;+\; (\kappa\beta)W_T \;\middle|\; X(t_n^-)=(s,b)\right]\right\}. \tag{B.2}$$

Proof. This follows similar steps as in Forsyth (2020a), proof of Proposition 6.2. □

Remark B.1 (An Implementable Strategy)**.** *Given an initial level of wealth W_0^- at t_0, the optimal control for the pre-commitment problem (6.2) is the same optimal control for the time consistent problem ($TCEQ_{t_n}(\kappa\beta)$) (B.2), $\forall t>0$. Hence we can regard problem ($TCEQ_{t_n}(\kappa\beta)$) as the* EW-ES induced time consistent strategy*. The induced strategy is implementable, in the sense that the investor has no incentive to deviate from the strategy computed at time zero at later times (Forsyth, 2020a).*

C Algorithm for EW-ES Strategy

We use the method described in Forsyth (2020a) to solve Problem 6.2. We give a brief description of this technique below. We write equation (6.4) as

$$J\left(s,b,t_0^-\right) = sup_{W^*}V(s,b,0^-)\;, \tag{C.1}$$

where the auxiliary function $V(s,b,t)$ is defined as

$$V(s,b,W^*,t_n^-) = \sup_{\mathcal{P}_n\in\mathcal{A}_n}\left\{E_{\mathcal{P}_n}^{X_n^+,t_n^+}\left[W^*+\frac{1}{\beta}\min((W_T-W^*),0) \;+\; \kappa W_T\middle|X(t_n^-)=(s,b)\right]\right\}. \tag{C.2}$$

$$\text{subject to } \begin{cases} (S_t,B_t) \text{ follow processes (4.3) and (4.4)}; \quad t\notin\mathcal{T} \\ W_\ell^+ = S_\ell^- + B_\ell^- \;;\; X_\ell^+ = (S_\ell^+,B_\ell^+) \\ S_\ell^+ = p_\ell(\cdot)W_\ell^+ \;;\; B_\ell^+ = (1-p_\ell(\cdot))W_\ell^+ \\ p_\ell(\cdot)\in\mathcal{Z} \\ \ell = n,\ldots,M \;;\; t_\ell\in\mathcal{T} \end{cases}. \tag{C.3}$$

We have now decomposed the original problem (6.2) into two steps

- Given an initial cash value of W_0, and a fixed value of W^*, we solve problem (C.2) to determine $V(0,W_0,W^*,0^-)$.
- Then, we solve the original problem (6.2) by maximizing over W^*

$$J(0,W_0,0^-) = \sup_{W^*} V(0,W_0,W^*,0^-)\;. \tag{C.4}$$

C.1 Solution of Problem C.2

We solve Problem C.2 by dynamic programming. Set

$$V(s,b,W^*,T^+) = W^* + \frac{\min((s+b-W^*),0)}{\beta} \;+\; \kappa(s+b)\;. \tag{C.5}$$

For $t \in (t^+_{M-1}, t^-_M)$, we solve the PIDE

$$V_t + \frac{(\sigma^s)^2 s^2}{2} V_{ss} + (\mu^s - \lambda^s_\xi \gamma^s_\xi) s V_s + \lambda^s_\xi \int_{-\infty}^{+\infty} V(e^y s, b, t) f^s(y)\; dy + \frac{(\sigma^b)^2 b^2}{2} V_{bb}$$
$$+ (\mu^b + \mu^b_c \mathbf{1}_{\{b<0\}} - \lambda^b_\xi \gamma^b_\xi) b V_b + \lambda^b_\xi \int_{-\infty}^{+\infty} V(s, e^y b, t) f^b(y)\; dy - (\lambda^s_\xi + \lambda^b_\xi) V + \rho_{sb} \sigma^s \sigma^b s b V_{sb} = 0. \tag{C.6}$$

At rebalancing time t_{M-1}, we determine the optimal control $p_{M-1}(w = s+b, W^*)$ from

$$p_{M-1}(w, W^*) = \arg\max_{p' \in \mathcal{Z}} V(wp', w(1-p'), W^*, t^+_{M-1}), \tag{C.7}$$

so that

$$V(s,b,W^*,t^-_{M-1}) = V(wp_{M-1}(w,W^*), w(1-p_{M-1}(w,W^*)), t^+_{M-1})\;. \tag{C.8}$$

Working backwards, we continue in this way until we reach t_0.

C.2 Numerical Techniques

We localize the infinite domain to $(s,b) \in [s_{\min}, s_{\max}] \times [b_{\min}, b_{\max}]$, and discretize $[b_{\min}, b_{\max}]$ using an equally spaced $\log b$ grid with n_b nodes. Similarly, we discretize $[s_{\min}, s_{\max}]$ on an equally spaced $\log s$ grid with n_s nodes. For the case where we allow leverage, we also define a reflected grid with $b < 0$. We use the Fourier method in (Forsyth and Labahn, 2019) to solve PIDE (C.6). Localization errors are minimized using the domain extension method in (Forsyth and Labahn, 2019).

At rebalancing dates, we solve the optimization problem (C.7) by discretizing $p(\cdot)$ and using exhaustive search. Finally, the optimization problem (C.1) is solved using a one-dimensional optimization technique. Note that each evaluation of the objective function requires solution of problem (C.2) with a fixed value of W^*.

We compute and store the optimal controls from solving Problem 6.2 using the parametric model of the stock and bond processes. We then use the stored controls in Monte Carlo simulations to generate statistical results. As a robustness check, we also use the stored controls and simulate results using bootstrap resampling of historical data.

D Convergence Test

Table D.1 shows a detailed convergence test for the base case problem given in Table 7.2. As expected, we can see that the value function converges at a rate between first and second order. The ES and EW values, which are derived quantities, converge a bit more erratically. Note that there is good agreement between the algorithm in Section C and the Monte Carlo validation. We remind the reader that the controls (the fraction in stocks) are determined using the method in Section C. These controls are then used in the Monte Carlo simulations.

Algorithm in Section C				Monte Carlo	
Grid	ES (5%)	$E[W_T]$	Value Function	ES (5%)	$E[W_T]$
512×512	672.62	1457.68	2130.3	682.06	1450.78
1024×1024	695.15	1437.64	2132.8	698.66	1436.31
2048×2048	696.59	1437.95	2134.5	697.56	1437.73
4096×4096	698.41	1436.72	2135.1	698.68	1436.65

TABLE D.1: *Convergence test for scenario in Table 7.2 with parameters in Table 7.1. The Monte Carlo method used* 2.56×10^6 *simulations.* $\kappa = 1.0, \alpha = .05$. *Grid refers to the grid used in the Algorithm in Section C:* $n_x \times n_b$, *where* n_x *is the number of nodes in the* $\log s$ *direction and* n_b *is the number of nodes in the* $\log b$ *direction.*

E Detailed Efficient Frontiers: Synthetic Market

Tables E.1 and E.2 give the detailed results used to construct Figure 9.1(a).

κ	ES (5%)	$E[W_T]$	$Median[W_T]$
0.1	940.60	1069.19	1051.51
0.25	936.23	1090.89	1056.85
0.4	925.85	1120.77	1063.53
0.6	883.094	1202.04	1083.39
0.7	838.84	1268.85	1107.31
0.8	781.29	1344.45	1146.52
0.9	739.88	1392.98	1179.12
1.0	697.56	1437.73	1222.12
1.2	646.41	1484.39	1291.43
1.5	614.92	1508.10	1347.72
2.0	586.16	1524.96	1381.78
3.0	553.23	1538.65	1399.09
10.0	500.08	1550.05	1405.02
∞	489.00	1550.71	1405.15

TABLE E.1: *Synthetic market results for optimal strategies, assuming the scenario given in Table 7.2. Control determined by solving Problem 6.3. Stock index: real capitalization weighted CRSP stocks; bond index: 30-day T-bills. Parameters from Table 7.1. Units: thousands of dollars. Statistics based on* 2.56×10^6 *Monte Carlo simulation runs.*

F Detailed Efficient Frontiers: Historical Market

Tables F.1-F.2 give the detailed results used to generate Figure 10.1. We show only the case where the expected blocksize is 5 years.

Constant weight p	ES (5%)	$E[W_T]$	$Median[W_T]$
0.0	917.26	1022.76	1023.45
0.1	929.03	1066.65	1063.48
.20	895.80	1112.34	1103.31
.30	849.32	1159.87	1143.32
.35	824.46	1184.35	1163.24
.40	799.09	1209.33	1183.09
.45	773.44	1234.81	1202.82
.50	747.62	1260.79	1222.33
.55	721.71	1287.29	1241.68
.60	695.77	1314.33	1260.89
.65	669.84	1341.90	1279.88
.70	643.93	1370.02	1298.64
.80	592.19	1427.94	1335.30
.90	540.58	1488.19	1370.95
1.0	489.00	1550.71	1405.15

TABLE E.2: *Synthetic market results for constant weight strategies. Stock index: real capitalization weighted CRSP stocks; bond index: 30-day T-bills. Parameters from Table 7.1. Units: thousands of dollars. Statistics based on* 2.56×10^6 *Monte Carlo simulation runs.*

κ	ES (5%)	$E[W_T]$	$Median[W_T]$
0.1	846.10	1085.38	1052.49
0.25	855.16	1101.88	1062.07
0.4	866.53	1124.61	1075.61
0.6	865.83	1189.97	1115.54
0.7	848.92	1247.14	1150.09
0.8	816.59	1313.73	1198.20
0.9	791.41	1356.67	1237.18
1.0	762.13	1394.46	1284.53
1.2	722.44	1432.22	1357.37
1.5	695.87	1450.28	1410.98
2.0	669.71	1463.51	1437.64
3.0	631.84	1476.97	1448.58
10.0	571.11	1494.62	1452.69

TABLE F.1: *Historical market results for optimal strategies, assuming the scenario given in Table 7.2. Control determined by solving Problem 6.3. Stock index: real capitalization weighted CRSP stocks; bond index: 30-day T-bills. Scenario Table 7.1. Units: thousands of dollars. Statistics based on* 10^6 *bootstrapped simulations. Expected blocksize of 5 years.*

References

Bauerle, N. and S. Grether (2015). Complete markets do not allow free cash flow streams. *Mathematical Methods of Operations Research 81*, 137–146.

Bergman, Y. Z., B. D. Grundy, and Z. Wiener (1996). General properties of option prices. *Journal of Finance 51*, 1573–1610.

Bernard, C. and S. Vanduffel (2014). Mean-variance optimal portfolios in the presence of a bench-

Constant weight p	ES (5%)	$E[W_T]$	$Median[W_T]$
0.0	763.98	1029.82	1023.02
0.1	804.41	1070.73	1061.82
0.2	833.61	1112.67	1109.03
0.3	841.16	1155.72	1152.42
0.4	826.08	1199.95	1197.04
0.45	811.92	1222.55	1217.27
0.5	795.27	1245.46	1238.77
0.55	776.71	1268.71	1261.33
.60	756.67	1292.31	1284.25
.65	735.46	1316.27	1306.29
.70	712.98	1340.60	1326.66
.80	664.99	1390.41	1371.08
.90	614.04	1441.81	1411.06
1.0	560.63	1494.91	1452.69

TABLE F.2: *Historical market results for constant weight strategies, assuming the scenario given in Table 7.2. Stock index: real capitalization weighted CRSP stocks; bond index: 30-day T-bills. Scenario in Table 7.1. Units: thousands of dollars. Statistics based on* 10^6 *bootstrapped simulations. Expected blocksize of 5 years.*

mark with applications to fraud detection. *European Journal of Operational Research 234*, 469–480.

Black, F. and R. W. Jones (1987). Simplifying portfolio insurance. *The Journal of Portfolio Management 14:1*, 48–51.

Cont, R. and C. Mancini (2011). Nonparametric tests for pathwise properties of semimartingales. *Bernoulli 17*, 781–813.

Cui, X., D. Li, S. Wang, and S. Zhu (2012). Better than dynamic mean-variance: time inconsistency and free cash flow stream. *Mathematical Finance 22*, 346–378.

Dang, D.-M. and P. A. Forsyth (2016). Better than pre-commitment mean-variance portfolio allocation strategies: a semi-self-financing Hamilton-Jacobi-Bellman equation approach. *European Journal of Operational Research 250*, 827–841.

Dichtl, H., W. Drobetz, and M. Wambach (2016). Testing rebalancing strategies for stock-bond portfolos across different asset allocations. *Applied Economics 48*, 772–788.

Dybvig, P. H. and J. E. Ingersoll (1982). Mean-variance theory in complete markets. *Journal of Business 55:2*, 233–251.

Forsyth, P. and G. Labahn (2019). $\epsilon-$Monotone Fourier methods for optimal stochastic control in finance. *Journal of Computational Finance 22:4*, 25–71.

Forsyth, P. A. (2020a). Multi-period mean CVAR asset allocation: Is it advantageous to be time consistent? *SIAM Journal on Financial Mathematics 11:2*, 358–384.

Forsyth, P. A. (2020b). Optimal dynamic asset allocation for DC plan accumulation/decumulation: Ambition-CVAR. *Insurance: Mathematics and Economics 93*, 230–245.

Forsyth, P. A. and K. R. Vetzal (2017a). Dynamic mean variance asset allocation: Tests for robustness. *International Journal of Financial Engineering 4*, 1750021:1–1750021:37. DOI: 10.1142/S2424786317500219.

Forsyth, P. A. and K. R. Vetzal (2017b). Robust asset allocation for long-term target-based investing. *International Journal of Theoretical and Applied Finance 20:3*. Article 1750017/1-32.

Forsyth, P. A. and K. R. Vetzal (2019). Optimal asset allocation for retirement savings: deterministic vs. time consistent adaptive strategies. *Applied Mathematical Finance 26:1*, 1–37.

Goetzmann, W., J. E. Ingersoll, M. Spiegel, and I. Welch (2002). Sharpening Sharpe ratios. NBER working paper, 9116.

Kou, S. G. (2002). A jump-diffusion model for option pricing. *Management Science 48*, 1086–1101.

Kou, S. G. and H. Wang (2004). Option pricing under a double exponential jump diffusion model. *Management Science 50*, 1178–1192.

Lhabitant, F.-S. (2000). Derivatives in portfolio management: why beating the market is easy. Working paper, EDHEC Business School.

Li, D. and W.-L. Ng (2000). Optimal dynamic portfolio selection: multiperiod mean-variance formulation. *Mathematical Finance 10*, 387–406.

Lin, Y., R. MacMinn, and R. Tian (2015). De-risking defined benefit plans. *Insurance: Mathematics and Economics 63*, 52–65.

Ma, K. and P. A. Forsyth (2016). Numerical solution of the Hamilton-Jacobi-Bellman formulation for continuous time mean variance asset allocation under stochastic volatility. *Journal of Computational Finance 20*(1), 1–37.

MacMinn, R., P. Brockett, J. Wang, Y. Lin, and R. Tian (2014). The securitization of longevity risk and its implications for retirement security. In O. S. Mitchell, R. Maurer, and P. B. Hammond (Eds.), *Recreating Sustainable Retirement*, pp. 134–160. Oxford: Oxford University Press.

Mancini, C. (2009). Non-parametric threshold estimation models with stochastic diffusion coefficient and jumps. *Scandinavian Journal of Statistics 36*, 270–296.

Menoncin, F. and E. Vigna (2017). Mean-variance target based optimisation for defined contribution pension schemes in a stochastic framework. *Insurance: Mathematics and Economics 76*, 172–184.

Moussawi, R., K. Shen, and R. Velthuis (2022). The role of taxes in the rise of ETFs. SSRN 3744519.

Patton, A., D. Politis, and H. White (2009). Correction to: automatic block-length selection for the dependent bootstrap. *Econometric Reviews 28*, 372–375.

Politis, D. and J. Romano (1994). The stationary bootstrap. *Journal of the American Statistical Association 89*, 1303–1313.

Politis, D. and H. White (2004). Automatic block-length selection for the dependent bootstrap. *Econometric Reviews 23*, 53–70.

Rockafellar, R. T. and S. Uryasev (2000). Optimization of conditional value-at-risk. *Journal of Risk 2*, 21–42.

Spurgon, R. B. (2001). How to game your Sharpe ratio. *Journal of Alternative Investments 4:3 (Winter)*, 38–36.

Strub, M., D. Li, and X. Cui (2019). An enhanced mean-variance framework for robo-advising applications. SSRN 3302111.

Ungar, J. and M. T. Moran (2009). The cash-secured putwrite strategy and performance of related benchmark indexes. *Journal of Alternative Investments 11:4 (Spring)*, 43–56.

van Staden, P. M., D.-M. Dang, and P. Forsyth (2018). Time-consistent mean-variance portfolio optimization: a numerical impulse control approach. *Insurance: Mathematics and Economics 83*, 9–28.

van Staden, P. M., D.-M. Dang, and P. A. Forsyth (2021). On the distribution of terminal wealth under dynamic mean-variance optimal investment strategies. *SIAM Journal on Financial Mathematics 12*, 566–601.

Vigna, E. (2014). On efficiency of mean-variance based portfolio selection in defined contribution pension schemes. *Quantitative Finance 14*, 237–258.

Vigna, E. (2022). Tail optimality and preferences consistency for intertemporal optimization problems. *SIAM Journal on Financial Mathematics 13:1*, 295–320.

Wang, J. and P. A. Forsyth (2010). Numerical solution of the Hamilton-Jacobi-Bellman formulation for continuous time mean variance asset allocation. *Journal of Economic Dynamics and Control 34*, 207–230.

Zhou, X. Y. and D. Li (2000). Continuous-time mean-variance portfolio selection: a stochastic LQ framework. *Applied Mathematics and Optimization 42*, 19–33.