Smart leverage? Rethinking the role of Leveraged Exchange Traded Funds in constructing portfolios to beat a benchmark

Pieter M. van Staden* Peter A. Forsyth† Yuying Li‡

December 6, 2024

Abstract

Leveraged Exchange Traded Funds (LETFs), while extremely controversial in the literature, remain stubbornly popular with both institutional and retail investors in practice. While the criticisms of LETFs are certainly valid, we argue that their potential has been underestimated in the literature due to the use of very simple investment strategies involving LETFs. In this paper, we systematically investigate the potential of including a broad stock market index LETF in long-term, dynamically-optimal investment strategies designed to maximize the outperformance over standard investment benchmarks in the sense of the information ratio (IR). Our results exploit the observation that positions in a LETF deliver call-like payoffs, so that the addition of a LETF to a portfolio can be a convenient way to add inexpensive leverage while providing downside protection. Under stylized assumptions, we present and analyze closed-form IR-optimal investment strategies using either a LETF or standard/vanilla ETF (VETF) on the same equity index, which provides the necessary intuition for the potential and benefits of LETFs. In more realistic settings of infrequent trading, leverage restrictions and additional constraints, we use a neural network-based approach to determine the IR-optimal strategies, trained on bootstrapped historical data. We find that IR-optimal strategies with a broad stock market LETF are not only more likely to outperform the benchmark than IR-optimal strategies derived using the corresponding VETF, but are able to achieve partial stochastic dominance over the benchmark and VETF-based strategies in terms of terminal wealth, even if investment in the VETF can be leveraged. Our results help to explain the empirical appeal of LETFs to investors, and encourage the reconsideration in academic research of the role of broad stock market LETFs within the context of more sophisticated investment strategies.

Keywords: Asset allocation, leveraged investing, portfolio optimization, neural network

JEL classification: G11, C61

1 Introduction

Leveraged Exchange Traded Funds (LETFs) are exchange traded funds (ETFs) with the stated objective of replicating some multiple β of the daily returns of their underlying reference assets/indices before costs, where values of β of $+2$, $+3$, -2 and -3 are commonly used. In contrast, standard/vanilla ETFs (VETFs) aim simply to replicate the returns of their underlying assets/indices before costs (i.e. $\beta = 1$).

A review of the academic literature suggests that incorporating LETFs in investment strategies are commonly regarded with at least some suspicion, if not outright distrust. "*Just say no to leveraged ETFs*", the title of a recent article (Bednarek and Patel (2022)), provides perhaps the most succinct summary of the broadly negative view of LETFs that permeates the literature. There are certainly good reasons for these negative perceptions of LETFs. A common criticism in the literature focuses on the "compounding" effect of LETF returns, which arises since a LETF returning β times the *daily* returns of the underlying index obviously does not imply that the LETF also returns β times the *quarterly* returns of the underlying index. This observation, along with the time decay and volatility decay of LETF positions, results in the potential wealth-destroying effects of LETF investments which increase with the holding time horizon (Mackintosh (2008),Carver (2009), Sullivan

*National Australia Bank, Melbourne, Victoria, Australia 3000. The research results and opinions expressed in this paper are solely those of the authors, are not investment recommendations, and do not reflect the views or policies of the NAB Group. pieter.vanstaden@gmail.com

†Cheriton School of Computer Science, University of Waterloo, Waterloo ON, Canada, N2L 3G1, paforsyt@uwaterloo.ca

‡Cheriton School of Computer Science, University of Waterloo, Waterloo ON, Canada, N2L 3G1, yuying@uwaterloo.ca

(2009), Charupat and Miu (2011)). The poor investment outcomes using LETFs reported in the literature should therefore not come as a surprise if an investor uses a basic buy-and-hold investment strategy (Charupat and Miu (2011), Bednarek and Patel (2022), Leung and Santoli (2012)), or very simple (if not outright naive) portfolio rebalancing rules typically considered in conjunction with unrealistically frequent rebalancing[1] from the perspective of long-term investors (Cheng and Madhavan (2009), Avellaneda and Zhang (2010), Bansal and Marshall (2015), DeVault et al. (2021)). While there are a few studies observing that LETFs could have a role within diversified portfolios, especially where the investor might have relatively aggressive performance targets (Bansal and Marshall (2015), Hill and Foster (2009)), wish to circumvent onerous leverage restrictions or large margin rates on borrowing (DeVault et al. (2021)), or want to outperform broad market indices using daily rebalancing (Knapp (2023)), these perspectives remain the exception to the mainstream academic view and tend to leave questions regarding the formulation of practical investment strategies unanswered.

However, the contrast between the general perceptions in the academic literature and investment practice could not be more profound. LETFs have consistently remained incredibly popular financial products since their introduction in 2006, as emphasized by recent headlines such as "*Investors pump record sums into leveraged ETFs*" (Financial Times, November 2022[2]) and "*Retail investors snap up triple-leveraged US equity ETFs*" (Financial Times, May 2024[3]). LETFs consistently dominate the top 10 most popular ETFs listed on US exchanges[4].

Perhaps more significantly, LETFs also enjoy substantial popularity among institutional investors. Analyzing the quarterly reports by institutional investment managers with at least US$100 million in assets under management that were filed with the SEC from September 2006 to December 2016, DeVault et al. (2021) finds that more than 20% of the reports reference at least one LETF in the end-of-quarter portfolio allocation.

Leaving speculative trading aside, what could explain the appeal of LETFs for institutional investors? Suppose an investor wants to leverage returns in a cost-effective way which also offers some downside protection. Since the requirement of downside protection rules out simple leverage, such an institutional investor has broadly speaking at least two options, namely (a) engage with a hedge fund or fund manager to use for example managed futures strategies, or (b) follow a dynamic trading strategy using for example LETFs as discussed in this paper. Since the expense ratios of LETFs range typically between 80 and 150 basis points, whereas standard leveraged positions are subject to substantial margin rates of borrowing which can easily exceed 5% for smaller institutional investors even during periods of low interest rates (DeVault et al. (2021)) while hedge funds charge hefty management fees, LETFs are certainly cost effective. In addition, LETFs can offer great upside returns in combination with limited liability without the need to manage short positions, and as discussed in this paper, positions in LETFs can be *infrequently* rebalanced while still obtaining competitive investment outcomes relative to standard investment benchmarks.

1.1 Main Contributions

While the academic literature offers a sophisticated and careful treatment of optimal rebalancing for hedging purposes in the case of LETFs (Dai et al. (2023)), or optimal replication policies for LETF construction (Guasoni and Mayerhofer (2023)), in this paper we therefore aim to make progress towards closing the observed gap between the literature and investment practice by showing that there is a role for LETFs in *infrequently rebalanced* (e.g. quarterly rebalanced) portfolios designed for *long-term* institutional or retail investors wishing to outperform some investment benchmark.

In more detail, the main contributions of this paper are as follows:

(i) We construct dynamic (multi-period) investment strategies which maximize the information ratio (IR) of the active portfolio manager (or simply "investor") relative to standard investment benchmarks using a LETF or a VETF on the same underlying equity index as well as bonds. The IR is chosen due to its popularity in investment practice (Bajeux-Besnainou et al. (2013); Bolshakov and Chincarini (2020); Hassine and Roncalli (2013); Israelsen and Cogswell (2007)), so that the results are not just of academic interest but also of practical relevance to institutional investors.

[1] For example, daily rebalancing. Note that in some cases, the general investment literature actually advises investors not to hold LETFs for longer than a single trading session. See for example Forbes Advisor, accessed 10 March 2024. Michael Adams. *Eight best leveraged ETFs of March 2024*. https://www.forbes.com/advisor/investing/best-leveraged-etfs.

[2] Financial Times, November 14, 2022. Steven Johnson. *Investors pump record sums into leveraged ETFs*. https://www.ft.com/content/b98ab360-2506-44f2-8e08-9d434df5f15d

[3] Financial Times, May 4, 2024. George Steer and Will Schmitt. *Retail investors snap up triple-leveraged US equity ETFs*. https://on.ft.com/3WsWTom

[4] For example, as at 9 October 2023, four out of the five most popular ETFs as measured by the average daily trading volume over the preceding three months were LETFs. https://etfdb.com/compare/volume. Accessed 9 October 2023

(ii) Under stylized assumptions including parametric dynamics for the underlying assets, we present closed-form IR-optimal dynamic investment strategies which enable us to obtain intuition regarding the expected behavior of IR-optimal investment strategies in more general settings. Note that the closed-form solutions allow for jump-diffusion dynamics of the equity index, which as noted above is crucial to consider in the case of LETFs, so that our results contributes to the existing literature which is almost exclusively based on diffusion dynamics (see for example Giese (2010), Jarrow (2010) Leung and Santoli (2012), Leung et al. (2017), Leung and Sircar (2015), Wagalath (2014), Guasoni and Mayerhofer (2023)). However, in the context of $\mathbb{Q}$ measure option pricing, Ahn et al. (2015) consider jump processes for LETFs.

(iii) Relaxing the stylized assumptions to allow for multi-asset portfolios, infrequent rebalancing, multiple investment constraints including leverage restrictions and no need to specify parametric asset dynamics, we use a neural network-based approach to obtain and analyze the resulting dynamic IR-optimal investment strategies for different scenarios, including: (i) investing in a VETF and bonds but with no leverage allowed, (ii) investing in a VETF and bonds with different levels of maximum leverage allowed and different levels of borrowing premiums being applicable, and (iii) investing in a LETF on the same equity market index as well as bonds with no leverage being allowed. We use a data-driven approach based on stationary block bootstrap resampling of historical data.

(iv) We find that IR-optimal investment strategies involving LETFs are fundamentally *contrarian*. This finding aligns to the empirical asset allocation behavior observed by DeVault et al. (2021) in their analysis of the SEC filings by institutional fund managers, whereby managers seem to decrease their holdings in LETFs after observing strong recent investment performance. In terms of investment performance, we find that IR-optimal strategies including the LETF are not only more likely to outperform the benchmark than IR-optimal strategies derived using the corresponding VETF, but are able to achieve partial stochastic dominance over the investment benchmark in terms of portfolio value (wealth).[5] This result holds even if the VETF can be leveraged with zero borrowing premium over the risk-free rate. Our results therefore encourage the reconsideration of the role of broad equity market LETFs within more sophisticated dynamic investment strategies, and provide a potential additional motivation regarding the enduring popularity of LETFs observed in practice.

To ensure the practical relevance of our conclusions, illustrative investment results are based on data sets generated using (i) stochastic differential equations calibrated to historical data for closed-form solutions and (ii) block bootstrap resampling of historical data (Anarkulova et al. (2022); Cogneau and Zakalmouline (2013); Politis and Romano (1994)) for neural network-based numerical solutions. Due to the relative paucity of LETF data, we construct a proxy LETF replicating the $\beta = 2$ times the daily returns of a broad US equity market index since 1929, deflating the returns by a typical LETF expense ratio and interest rates (see Appendix B). We make the standard assumption (see e.g. Bansal and Marshall (2015), Leung and Sircar (2015)) that the LETF managers do not have challenges in replicating the underlying index. Note that given improvements in designing replication strategies for LETFs that remain robust even during periods of market volatility (see for example Guasoni and Mayerhofer (2023)), this appears to be a reasonable assumption for ETFs (VETFs and LETFs) written on major equity market indices as considered in this paper. All returns time series are inflation-adjusted to reflect real returns.

1.2 Intuition

In this section, we provide some insight into the potential advantages of including LETFs in optimal dynamic asset allocation. We will give an overview here, leaving the technical details to Section 3.1.

Suppose an investor allocates their initial wealth $W(0)$ to US 30-day T-bills and either a LETF ($\beta = 2$) or a VETF on a broad US equity market index S at time $t_0 = 0$, and does not rebalance the portfolio over the holding time horizon $\Delta t > 0$. We discuss in more detail in Section 3 how the LETF and VETF can be viewed as derivative contracts on underlying S costing $F_\ell(0)$ (LETF) and $F_v(0)$ (VETF) to purchase, with payoffs of $F_\ell(\Delta t)$ and $F_v(\Delta t)$, respectively.

Assume that the underlying index price follows a geometric Brownian motion, which implies that the price is always positive. In addition, assume a constant risk-free rate for short term bonds. It can be easily shown

[5] For a definition of partial stochastic dominance, see Atkinson (1987); Van Staden et al. (2021)

that (see e.g. (Avellaneda and Zhang, 2010))

$$\frac{F_\ell(\Delta t)}{F_\ell(0)} = \exp\{-c_\ell \cdot \Delta t\} \cdot f_\ell(\Delta t; \beta) \cdot \left(\frac{S(\Delta t)}{S(0)}\right)^\beta, \tag{1.1}$$

where

$$f_\ell(\Delta t; \beta) = \exp\left\{-\left[(\beta - 1)r + \frac{1}{2}(\beta - 1)\beta\sigma^2\right] \cdot \Delta t\right\} \tag{1.2}$$

$c_\ell > 0$ is the expense ratio, r is the risk-free interest and σ is the volatility of return. In other words, payoff $F_\ell(\Delta t)$ is a deterministic function of the terminal value of the underlying index. Of course, $S(\Delta t)$ itself is stochastic.

Since small values of maturity Δt can be undesirable due to frequent trading, and large values of Δt emphasize the time- and volatility decay of simply holding the LETF F_ℓ, suppose the investor chooses a convenient maturity of $\Delta t = 0.25$ years (one quarter). Figure 1.1 illustrates the payoff diagrams for the investor's wealth at maturity $W(\Delta t)$ under different combinations of T-bills and an ETF, as a function of the value at maturity $S(\Delta t)$ of the underlying equity index. Similar payoff diagrams can be seen in Knapp (2023) and for the unleveraged case in Bertrand and Prigent (2022).

For simplicity, we assume parametric asset dynamics for the T-bills and index S (S follows geometric Brownian motion) in Figure 1.1 calibrated to US market data over 1926 to 2023. We impose realistic ETF expense ratios, and a borrowing premium of 3% over the T-bill rate for short positions. Note that the assumption of parametric asset dynamics is for purposes of intuition only, since the investment results in Section 5 do not rely on any parametric assumptions. Leaving rigorous derivation for subsequent sections, we make the following qualitative observations regarding the payoff diagrams for purposes of intuition:

- Figures 1.1 illustrate that we can characterize the payoffs of LETF investments as *call-like*. This suggests that the addition of an LETF to a portfolio can be a useful way to add inexpensive leverage while preserving downside protection, much like a usual call option. Provided that the investment in the LETF is itself not funded by borrowing (i.e. the LETF position itself is not leveraged), the LETF payoff is always non-negative due to limited liability even in the case of significant drops in the value of the underlying equity index, in contrast with leveraged VETF positions.

- For calibrated geometric Brownian motion (GBM) dynamics for the equity index S, Figure 1.1 illustrates that investing all wealth in the LETF with $\beta = 2$ (Figure 1.1(a)) dominates the 2x leveraged investment in the VETF (200% of wealth in VETF funded by borrowing and amount equal to 100% of wealth) almost everywhere, but underperforms a 100% investment in the VETF for negative underlying index returns (i.e. when $S(\Delta t)/S(0) < 1$). By contrast, investing 50% of wealth in the LETF (Figure 1.1(a)) and the remaining 50% in T-bills dominates the payoff of investing 100% of wealth in the VETF almost everywhere, but underperforms a 2x leveraged investment in the VETF for positive underlying index returns ($S(\Delta t)/S(0) > 1$). Note that under GBM dynamics for S, the terminal wealth $W(\Delta t)$ conditional on the terminal value $S(\Delta t)$ is deterministic for both VETF and LETF investments (see Section 3.1).

When the underlying index is modelled by a jump process, due to limited liability, a correction to the index price is needed so that LETF remains positive, see Section 3.1 for a more detailed discussion. This makes the payoff relation between LETF and the underlying index stochastic. While qualitatively similar observations as in the case of no jumps (Figure 1.1) apply to the *median* payoffs of the LETF investments, allowing for jumps in the underlying asset dynamics can affect the LETF payoff significantly, and jumps are therefore critical to incorporate in the investor's strategy. However, most of the existing literature on investment strategies with LETFs only allows for pure diffusion processes for the equity index underlying the LETF (Giese (2010) Jarrow (2010), Leung and Santoli (2012), Leung et al. (2017), Leung and Sircar (2015), Wagalath (2014), Guasoni and Mayerhofer (2023)).

We emphasize that while Figure 1.1 is for the purposes of intuition only, it is nevertheless based on asset dynamics calibrated to empirical US market data. Since even long-term investments (e.g. 10 years) can be managed effectively using a *dynamic* investment strategy with for example quarterly rebalancing (i.e. at the beginning of each quarter, the investor faces investment choices and associated outcomes such as those in Figure 1.1), this suggests that the benefits of LETFs could potentially be harnessed without being unduly affected by the compounding effects as well as time- and volatility-decay. Our results show that this is indeed the case, even if no parametric form of the underlying dynamics is assumed.

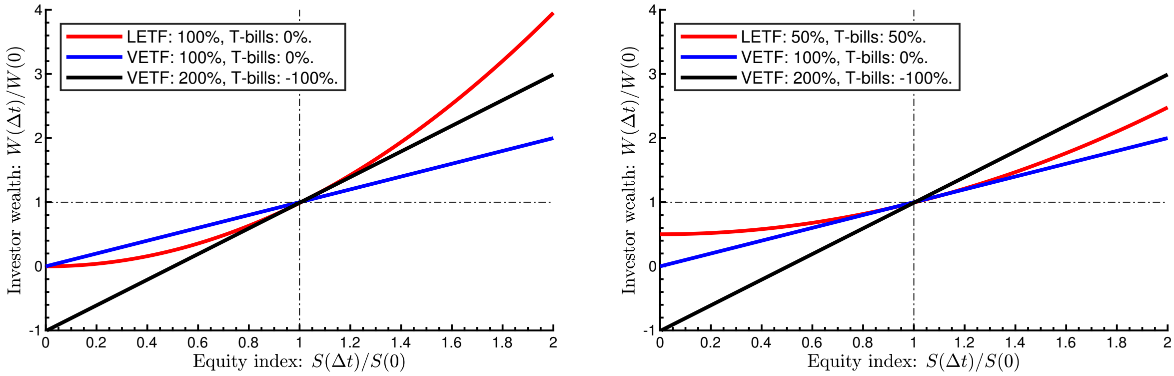

(a) LETF: 100%, T-bills: 0% vs. VETF combinations

(b) LETF: 50%, T-bills: 50% vs. VETF combinations

Figure 1.1: Payoffs when equity market index S follows calibrated GBM dynamics: Investor wealth gross return $W(\Delta t)/W(0)$ as a function of underlying equity index gross return $S(\Delta t)/S(0)$, $\Delta t = 0.25$ (1 quarter), for different proportions of initial wealth $W(0)$ invested in the LETF, VETF and T-bills at time $t_0 = 0$. Asset parameters are calibrated to US equity and bond market data over the period 1926:01 to 2023:12 (Appendix B), LETF and VETF expense ratios are assumed to be 0.89% and 0.06% respectively, and a borrowing premium of 3% over the T-bill rate is applicable to short positions. See Section 3.1 for a rigorous treatment of the illustrated relationships.

1.3 Organization

The remainder of the paper is organized as follows. Section 2 provides the general problem formulation, with Section 3 presenting closed-form results obtained under stylized assumptions. Section 4 discusses a neural network-based solution approach to obtain the optimal investment strategies numerically under multiple investment constraints. Finally, Section 5 presents indicative investment results and Section 6 concludes the paper, with additional analytical and numerical results presented in Appendices A, B and C.

2 General problem formulation

In this section we formulate, in general terms, the dynamic portfolio optimization problem to be solved by an active portfolio manager (simply "investor") over a given time horizon $[t_0 = 0, T]$, where $T > 0$ can be large (e.g. 10 years). We assume that the investment performance of the investor is measured relative to that of a given benchmark portfolio, as is typically the case for professional asset managers (see for example Alekseev and Sokolov (2016); Kashyap et al. (2021); Korn and Lindberg (2014); Lehalle and Simon (2021); Zhao (2007)). To this end, for any $t \in [t_0 = 0, T]$, let $W(t)$ and $\hat{W}(t)$ denote the portfolio value (or informally, simply the "wealth") of the investor and benchmark portfolios, respectively. The same initial wealth $w_0 := W(t_0) = \hat{W}(t_0) > 0$ is assumed to ensure that the performance comparison remains fair. The investor's strategy is based on investing in any of a set of N_a candidate assets, while the benchmark is defined in terms of $\hat{N}_a$ potentially different underlying assets.

In more detail, if $\hat{\boldsymbol{X}}(t)$ denotes the state (or informally, the information) used in obtaining the benchmark asset allocation strategy at time $t \in [t_0, T]$, let $\hat{p}_j\left(t, \hat{\boldsymbol{X}}(t)\right)$ denotes the proportion of the benchmark wealth $\hat{W}(t)$ invested in asset $j \in \left\{1, .., \hat{N}_a\right\}$. The vector $\hat{\boldsymbol{p}}\left(t, \hat{\boldsymbol{X}}(t)\right) = \left(\hat{p}_j\left(t, \hat{\boldsymbol{X}}(t)\right) : j = 1, .., \hat{N}_a\right) \in \mathbb{R}^{\hat{N}_a}$ then denotes the asset allocation or investment strategy of the benchmark at time $t \in [t_0, T]$.

Similarly, if $\boldsymbol{X}(t)$ denotes the state or information incorporated by the investor in making their asset allocation decision, let $p_i(t, \boldsymbol{X}(t))$ denote the proportion of the investor's wealth $W(t)$ invested in asset $i \in \{1, .., N_a\}$, with $\boldsymbol{p}(t, \boldsymbol{X}(t)) = (p_i(t, \boldsymbol{X}(t)) : i = 1, .., N_a) \in \mathbb{R}^{N_a}$ denoting the investor's asset allocation or investment strategy at time $t \in [t_0, T]$.

The set of portfolio rebalancing events is denoted by $\mathcal{T} \subseteq [t_0, T]$, where we consider $\mathcal{T} = [t_0, T]$ in the case of continuous rebalancing (Section 3), or a discrete subset $\mathcal{T} \subset [t_0, T]$ in the case of discrete rebalancing (Section 4). Given the set $\mathcal{T}$, the investor and benchmark investment strategies are respectively defined as

$$\mathcal{P} = \{\boldsymbol{p}(t, \boldsymbol{X}(t)),\ t \in \mathcal{T}\}, \qquad \text{and} \qquad \hat{\mathcal{P}} = \left\{\hat{\boldsymbol{p}}\left(t, \hat{\boldsymbol{X}}(t)\right),\ t \in \mathcal{T}\right\}. \tag{2.1}$$

It is typical for the investor to be subject to investment constraints, which are encoded by the set $\mathcal{A}$ of admissible controls. In the simplest case, admissible investor strategies $\mathcal{P} \in \mathcal{A}$ are such that $\mathcal{P} = \{\boldsymbol{p}(t, \boldsymbol{X}(t)) \in \mathcal{Z} : t \in \mathcal{T}\}$, with $\mathcal{Z}$ denoting the admissible control space. More complex constraints require a more careful formulation of $\mathcal{A}$ and $\mathcal{Z}$, see for example Section 4.

Since the investor aims to construct $\mathcal{P}$ to *outperform* the benchmark strategy $\hat{\mathcal{P}}$, Assumption 2.1 below outlines some general assumptions regarding the investment benchmark $\hat{\mathcal{P}}$. Note that Assumption 2.1 aligns with investment practice and is important for assessing the relevance of LETFs when constructing portfolios for outperforming a benchmark - see further discussion in Remark 2.1 below.

Assumption 2.1. (General assumptions regarding the benchmark strategy $\hat{\mathcal{P}}$) We make the following general assumptions regarding the benchmark strategies considered in this paper:

(i) The investor can observe the asset allocation $\hat{\boldsymbol{p}}\left(t, \hat{\boldsymbol{X}}(t)\right)$ of the benchmark strategy at each $t \in \mathcal{T}$.

(ii) The sets of investable assets available to the investor and benchmark, respectively, do not necessarily correspond. In particular, the benchmark strategy may invest in assets which the investor is unwilling or unable to invest in, or the investor might consider investing in a much larger universe of investable assets than those included in the benchmark. □

Remark 2.1 highlights some key observations regarding Assumption 2.1.

Remark 2.1. (Clarification of benchmark assumptions) With regards to Assumption 2.1, we observe the following:

(i) Observable benchmarks play a key role in performance reporting for many institutional investors, since active portfolio managers often explicitly pursue the outperformance of a *predetermined* investment benchmark (see for example Alekseev and Sokolov (2016); Kashyap et al. (2021); Korn and Lindberg (2014); Lehalle and Simon (2021); Zhao (2007)). As a result, the benchmark is clearly defined and transparent in the sense of the underlying asset allocation, which often incorporates broad market indices and bonds. In the case of pension funds, the benchmark (or "reference") portfolios are usually constructed using traded assets in fixed proportions. Examples include the Canadian Pension Plan (CPP) with a base reference portfolio of 15% Canadian government bonds and 85% global equity (Canadian Pension Plan (2022)), or the Norwegian government pension plan ("Government Pension Fund Global", or GPFG) using a benchmark portfolio of 70% equities and 30% bonds (Government Pension Fund Global (2022)).

(ii) Active portfolio managers often consider not only different but indeed larger/broader sets of assets than the benchmark. For example, pension funds might include private equity whereas the benchmark might be based on publicly traded assets only (see for example Canadian Pension Plan (2022)). In the assessment of the effect of replacing VETFs with LETFs discussed in Section 3, we consider scenarios where the benchmark strategy is defined in terms of a broad stock market index, but the investor might not be able to invest directly in the index itself, and invests instead in an ETF (VETF or LETF) replicating the index returns. Since the ETFs only replicate (a multiple of) the index returns *before* costs, the existence of a non-zero ETF expense ratios implies that investing in ETFs is not exactly the same as investing in the underlying index, i.e. an ETF and its underlying index can be viewed as different assets. Assumption 2.1(ii) is therefore relevant to an assessment of the role of a VETF or LETF within portfolios designed to beat a broad equity index-based investment benchmark. □

Let $E_{\mathcal{P}}^{t_0,w_0}[\cdot]$ and $Var_{\mathcal{P}}^{t_0,w_0}[\cdot]$ denote the expectation and variance, respectively, given initial wealth $w_0 = W(t_0) = \hat{W}(t_0)$ at time $t_0 = 0$ and using admissible investor strategy $\mathcal{P} \in \mathcal{A}$ over $[t_0, T]$. As a result of Assumption 2.1, the benchmark strategy $\hat{\mathcal{P}}$ remains implicit and fixed for notational simplicity.

As discussed in the Introduction, for an investment objective measuring outperformance, we wish to maximize the information ratio (IR) of the investor relative to the benchmark. In the context of dynamic trading with strategies of the form (2.1), the information ratio (IR) is defined as (Bajeux-Besnainou et al. (2013); Goetzmann et al. (2002))

$$\mathcal{IR}_{\mathcal{P}}^{t_0,w_0} = \frac{E_{\mathcal{P}}^{t_0,w_0}\left[W(T) - \hat{W}(T)\right]}{Stdev_{\mathcal{P}}^{t_0,w_0}\left[W(T) - \hat{W}(T)\right]}. \tag{2.2}$$

Maximizing the IR (2.2) can be achieved by solving a mean-variance (MV) optimization problem (Bajeux-Besnainou et al. (2013)) ,

$$\sup_{\mathcal{P}\in\mathcal{A}} \left\{ E_{\mathcal{P}}^{t_0,w_0}\left[W(T) - \hat{W}(T)\right] - \rho \cdot Var_{\mathcal{P}}^{t_0,w_0}\left[W(T) - \hat{W}(T)\right] \right\}, \quad \rho > 0, \tag{2.3}$$

where ρ denotes a scalarization parameter.

Using the embedding technique of Li and Ng (2000); Zhou and Li (2000), we solve (2.3) by formulating the equivalent problem (Van Staden et al. (2023))

$$(IR(\gamma)): \qquad \inf_{\mathcal{P}\in\mathcal{A}} E_{\mathcal{P}}^{t_0,w_0}\left[\left(W(T) - \left[\hat{W}(T) + \gamma\right]\right)^2\right], \qquad \gamma > 0, \tag{2.4}$$

where $\gamma > 0$ denotes the embedding parameter. As discussed in Van Staden et al. (2023), the parameter γ can be viewed as the investor's (implicit) target for benchmark outperformance formulated in terms of terminal wealth.

Remark 2.2. (Time consistency.) Note that the control for problem (2.3) is formally the pre-commitment policy, i.e. not time consistent. However, the pre-commitment policy solution of Problem (2.3) is identical to the strategy for an induced time consistent policy (Forsyth, 2020; Strub et al., 2019), and hence it is implementable.[6] The induced time consistent strategy in this case is the target based Problem (2.4), with a fixed value of $\gamma, \forall t > 0$. The relationship between pre-commitment and implementable target-based schemes in the mean-variance context is discussed in Menoncin and Vigna (2017); Vigna (2014, 2020, 2022). We consider the policy followed by the investor for $t > 0$ to be the implementable solution of Problem (2.4) with a fixed value of γ. This is identical to the solution of Problem (2.3) as seen at $t = 0$.

3 Closed-form solutions

To obtain the valuable intuition regarding the characteristics of IR-optimal investment strategies incorporating a LETF or VETF on a broad equity market index, we present closed-form solutions to the IR problem (2.4) under stylized assumptions. Remark 3.1 emphasizes that these assumptions are required for the derivation of closed-form solutions in this section only.

Remark 3.1. (Relaxing closed-form assumptions) The closed-form results presented in this section require stylized assumptions (Assumption 3.1 and Assumption 3.2 below), but we will use numerical techniques (Section 4) and present indicative investment results (Section 5) where these assumptions are not required. The investment problem is solved numerically in a setting where the following is allowed: (i) no restrictions on the number of underlying assets, (ii) no parametric assumptions are required for the dynamics of the underlying assets, (iii) discrete portfolio rebalancing is used, (iv) leverage is restricted and in some scenarios not allowed at all, (v) nonzero borrowing premiums over the risk-free rate are applicable when funding leveraged positions, (vi) no trading in the event of insolvency can occur, and (vii) more general benchmark strategies are allowed, though for illustrative purposes we will use constant proportion strategies due to their popularity in practical applications. □

The first set of general assumptions for the derivation of the closed-form solution in this section is outlined in Assumption 3.1.

Assumption 3.1. (Stylized assumptions for closed-form solutions) For the purposes of obtaining closed-form solutions in this section, we assume the following:

(i) The benchmark investment strategy (asset allocation) is a deterministic function of time defined in terms of the 30-day T-bills ("risk-free" asset) denoted by B and a broad equity market index ("risky" asset) denoted by S. Note that any known deterministic benchmark strategy clearly satisfies Assumption 2.1, and includes as a special case the constant proportion strategies which are popular benchmarks used in practice (see for example Canadian Pension Plan (2022),Government Pension Fund Global (2022)).

(ii) We consider two investors, each optimizing their respective portfolios relative to the same benchmark. Both investors are assumed to be unable or unwilling to invest directly the underlying broad equity market index itself (i.e. replicate the index with individual stocks), and instead invests in ETFs referencing the index.

[6] 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, 2020).

The first investor, informally referred to as the "VETF investor", allocates wealth to two underlying assets, namely 30-day T-bills B and a VETF F_v with expense ratio $c_v > 0$, where the VETF simply replicates the instantaneous returns of the index S before costs. The second investor, informally referred to as the "LETF investor", allocates wealth to two underlying assets, namely 30-day T-bills B and a LETF F_ℓ with expense ratio $c_\ell > 0$, where the LETF returns $\beta > 1$ times the instantaneous returns of the index S before costs.

(iii) Parametric dynamics for all underlying assets are assumed, including jump-diffusion dynamics for the broad equity market index S - see (3.1)-(3.2), (3.4) and (3.8) below. □

Table 3.1 provides an example of an investment scenario consistent with Assumption 3.1(i)-(ii), which will be used for the illustration of the closed-form solutions of this section.

Table 3.1: Closed-form solutions - Candidate assets and benchmark: Example of the investment scenario outlined in Assumption 3.1(i)-(ii), which will be used when illustrating the closed-form solutions in this section. The constant proportion benchmark has been chosen to align with typical benchmarks used by pension funds, while the indicative expense (or cost) ratios are chosen from the range of expense ratios of VETFs and LETFs on broad equity market indices typically observed in practice.

Underlying assets			Benchmark	Investor candidate assets	
Label	Value	Asset description		Using VETF	Using LETF
T30	$B(t)$	30-day Treasury bill	30%	✓	✓
Market	$S(t)$	Market portfolio (broad equity market index)	70%	-	-
VETF	$F_v(t)$	Vanilla or standard/unleveraged ETF (VETF) replicating the returns of the market portfolio $S(t)$, with expense ratio $c_v = 0.06\%$	0%	✓	-
LETF	$F_\ell(t)$	Leveraged ETF (LETF) with daily returns replicating $\beta = 2$ times the daily returns of the market portfolio $S(t)$, with expense ratio $c_\ell = 0.89\%$	0%	-	✓

We assume that the underlying index S can follow any of the commonly-encountered jump diffusion processes in finance (see for example Kou (2002); Merton (1976)), resulting in the following assumed dynamics for B and S, respectively,

$$\frac{dB(t)}{B(t)} = r \cdot dt, \tag{3.1}$$

$$\frac{dS(t)}{S(t^-)} = (\mu - \lambda\kappa_1^s)\, dt + \sigma \cdot dZ + d\left(\sum_{i=1}^{\pi(t)} (\xi_i^s - 1)\right). \tag{3.2}$$

In (3.1)-(3.2), r denotes the continuously compounded risk-free rate, $\pi(t)$ denotes a Poisson process with intensity $\lambda \geq 0$, while μ and σ denote the drift and volatility coefficients, respectively, under the objective (real-world) probability measure. ξ_i^s are i.i.d. random variables with the same distribution as ξ^s, which represents the jump multiplier associated with the S-dynamics, and we define

$$\kappa_1^s = \mathbb{E}[\xi^s - 1], \qquad \kappa_2^s = \mathbb{E}\left[(\xi^s - 1)^2\right], \tag{3.3}$$

which can be obtained using a given probability density function (pdf) of ξ^s, denoted by $G(\xi^s)$. Finally, for any functional $\psi(t), t \in [t_0, T]$, we use $\psi(t^-)$ and $\psi(t^+)$ as shorthand notation for the one-sided limits $\psi(t^-) = \lim_{\epsilon \downarrow 0} \psi(t - \epsilon)$ and $\psi(t^+) = \lim_{\epsilon \downarrow 0} \psi(t + \epsilon)$, respectively. Note that we can recover the assumption of geometric Brownian motion (GBM) dynamics for S by simply setting the intensity $\lambda \equiv 0$ in (3.2).

Since the VETF F_v with expense ratio c_v is a vanilla/standard ETF simply replicating the returns of S before costs, it has dynamics given by

$$\begin{aligned} \frac{dF_v(t)}{F_v(t^-)} &= \frac{dS(t)}{S(t^-)} - c_v \cdot dt \\ &= (\mu - \lambda\kappa_1^s - c_v) \cdot dt + \sigma \cdot dZ + d\left(\sum_{i=1}^{\pi(t)} (\xi_i^s - 1)\right). \end{aligned} \tag{3.4}$$

In contrast, the LETF F_ℓ with expense ratio $c_\ell > 0$ aims at replicating $\beta > 1$ times the instantaneous returns of the underlying broad stock market index S before costs, and therefore has dynamics *approximately* given by

$$\frac{dF_\ell(t)}{F_\ell(t^-)} \simeq \beta \frac{dS(t)}{S(t^-)} - [(\beta - 1) r + c_\ell] \, dt. \tag{3.5}$$

It should be emphasized that (3.5) is only an approximation. Since the investor in an LETF has limited liability, exact equality in (3.5) only holds in the case where there are no jumps in the dynamics of S. In the case of pure GBM dynamics, F_ℓ can never become negative, hence limited liability is irrelevant. As a result, (3.5) is indeed used to model LETF dynamics in the literature where the analysis is limited to GBM dynamics for S (see for example Avellaneda and Zhang (2010), Jarrow (2010), Guasoni and Mayerhofer (2023)), with the notable exception of Ahn et al. (2015) in the context of $\mathbb{Q}$ measure dynamics.

However, in the case of jump-diffusion dynamics for S, (3.5) is not quite correct since the LETF investor is protected by limited liability. In more detail, if there is a jump with multiplier ξ^s in the underlying index (3.2) at a specific time t, then the approximation (3.5) suggests that the value of the LETF would jump to $F_\ell(t) = F_\ell(t^-) \cdot [1 + \beta(\xi^s - 1)]$.

Therefore, in the case of a large downward jump in the underlying index, in particular where the jump multiplier satisfies $\xi^s < (\beta - 1)/\beta$, the approximation (3.5) implies that $F_\ell(t) < 0$, which cannot occur due to limited liability of the LETF. Instead, if it is indeed the case that $\xi^s \leq (\beta - 1)/\beta$, the value of the LETF simply drops to zero, i.e. $F_\ell(t) \equiv 0$.

We can therefore model the limited liability of the LETF F_ℓ by observing that F_ℓ therefore experiences jumps which are related to, but not necessarily exactly the same as the jumps experienced by the underlying index S. To this end, we define a jump multiplier ξ^ℓ for the F_ℓ-dynamics in terms of the jump multiplier ξ^s in the S-dynamics as

$$\xi^\ell = \begin{cases} \xi^s & \text{if } \xi^s > (\beta - 1)/\beta, \\ \dfrac{(\beta - 1)}{\beta} & \text{if } \xi^s \leq (\beta - 1)/\beta. \end{cases} \tag{3.6}$$

The second case in (3.6) enforces the limited liability of the LETF investor in the case of large downward jumps, i.e. $F_\ell(t) \equiv 0$ if $\xi^s \leq (\beta - 1)/\beta$. For subsequent use, we also define the following quantities involving the LETF jump multiplier ξ^ℓ,

$$\kappa_1^\ell = \mathbb{E}\left[\xi^\ell - 1\right], \qquad \kappa_2^\ell = \mathbb{E}\left[\left(\xi^\ell - 1\right)^2\right], \qquad \kappa_\chi^{\ell,s} = \mathbb{E}\left[\left(\xi^\ell - 1\right)\left(\xi^s - 1\right)\right]. \tag{3.7}$$

Given (3.2), (3.5) and (3.6), the LETF dynamics correctly incorporating jumps is therefore given by

$$\frac{dF_\ell(t)}{F_\ell(t^-)} = [\beta(\mu - \lambda\kappa_1^s) - (\beta - 1) r - c_\ell] \cdot dt + \beta\sigma \cdot dZ + \beta \cdot d\left(\sum_{i=1}^{\pi(t)} \left(\xi_i^\ell - 1\right)\right), \tag{3.8}$$

where ξ_i^ℓ are i.i.d. random variables with the same distribution as ξ^ℓ, which represents the jump multiplier associated with the F_ℓ-dynamics (3.6). We highlight the following observations regarding the dynamics (3.2), (3.4) and (3.8).

(i) The jumps in the underlying index S, the VETF F_v and LETF F_ℓ occur at the same time, so the Poisson process $\pi(t)$ and intensity λ are the same in (3.2), (3.4) and (3.8). More formally, the processes (3.2), (3.4) and (3.8) have the same Poisson random measures, but the compensated Poisson random measure is different for the F_ℓ-dynamics since the LETF jump sizes are slightly different due to (3.6).

(ii) The dynamics (3.4) and (3.8) implicitly assume that the ETFs have negligible tracking errors, while having non-negligible expense ratios. While ETF expense ratios can indeed be material, especially for LETFs, the assumption that tracking errors are negligible are often employed in the literature (see for example Bansal and Marshall (2015), Leung and Sircar (2015)). Given the recent developments in designing replication strategies for LETFs that remain robust even during periods of market volatility (see for example Guasoni and Mayerhofer (2023)), this appears to be a reasonable assumption especially in the case of the most popular VETFs and LETFs written on the major stock market indices.

(iii) As noted in the Introduction, we limit the analysis to the case of LETFs where $\beta > 1$ (i.e. "bullish"

LETFs). However, the dynamics (3.8) could also be applicable to inverse or "bearish" ETFs where $\beta < 1$ (see for example Jarrow (2010)), but adjustments are usually required to incorporate the time-dependent borrowing cost involved in short-selling the particular components of the underlying replication basket each time t (see for example Avellaneda and Zhang (2010)).

Remark 3.2. (Relationship to Ahn et al. (2015)) In Ahn et al. (2015), the authors model the limited liability of the LETF by taking the point of view of the manager of the LETF. The manager must purchase insurance to handle the cases where the manager's position becomes negative. In this work, we simply take the point of view of the holder the LETF (not the manager), who has no exposure to the possible negative value of the manager's position. The cost of this insurance (to the manager) is assumed to be passed on to the LETF investor as part of the fee c_ℓ charged by the manager, which is easily observable.

3.1 Intuition: lump sum investment scenario

As a simple and intuitive illustration of the potential and risks of using LETFs vs. VETFs, we consider a simple version of the general formulation of the problem as outlined in Section 2. Specifically, we consider a lump sum investment scenario as discussed in the Introduction (see Figures 1.1 and 3.1), where the initial wealth $w_0 = W(t_0) = \hat{W}(t_0) > 0$ is invested at time $t_0 = 0$ with no intermediate intervention/rebalancing until the terminal time $T = \Delta t = 0.25$ years (i.e. one quarter). In the notation of Section 2, we therefore have a trivial set of rebalancing events $\mathcal{T} = [t_0]$.

First, we consider the implications of the underlying asset dynamics without referencing the investment strategy (i.e. wealth allocation to assets). The dynamics (3.1)-(3.2) imply that

$$\frac{B(\Delta t)}{B(0)} = \exp\{r \cdot \Delta t\}, \tag{3.9}$$

$$\frac{S(\Delta t)}{S(0)} = \exp\left\{\left(\mu - \lambda\kappa_1^s - \frac{1}{2}\sigma^2\right)\cdot \Delta t + \sigma \cdot Z(\Delta t) + \sum_{i=1}^{\pi(\Delta t)} \log \xi_i^s\right\}, \tag{3.10}$$

where the values $B(0)$ and $S(0)$ are observable at time $t_0 = 0$. In the case of the VETF, we simply have

$$\frac{F_v(\Delta t)}{F_v(0)} = \exp\{-c_v \cdot \Delta t\} \cdot \left(\frac{S(\Delta t)}{S(0)}\right). \tag{3.11}$$

In the case of the LETF, we have

$$\frac{F_\ell(\Delta t)}{F_\ell(0)} = \exp\{-c_\ell \cdot \Delta t\} \cdot f_\ell(\Delta t; \beta) \cdot \widetilde{Y}_\ell(\Delta t; \beta) \cdot \left(\frac{S(\Delta t)}{S(0)}\right)^\beta, \tag{3.12}$$

where

$$f_\ell(\Delta t; \beta) = \exp\left\{-\left[(\beta - 1)r + \frac{1}{2}(\beta - 1)\beta\sigma^2\right]\cdot \Delta t\right\}, \quad \text{and} \quad \widetilde{Y}_\ell(\Delta t; \beta) = \prod_{i=1}^{\pi(\Delta t)}\left[\frac{1 + \beta\left(\xi_i^\ell - 1\right)}{\left(\xi_i^s\right)^\beta}\right]. \tag{3.13}$$

Expression (3.12) for the case where there are no jumps, is given in Avellaneda and Zhang (2010).

For purposes of intuition, consider the *special case* where the underlying index S experiences *zero growth/decline* over the time horizon Δt. If $S(\Delta t) = S(0)$, it is clear that the LETF will perform worse than the VETF, i.e. assuming $\beta > 1$, $F_\ell(\Delta t) < F_v(\Delta t)$, due to the following:

- Decay due to volatility: The term $f_\ell(\Delta t; \beta) < 1$, which only affects the LETF, is dominated by the diffusive volatility σ in the underlying S-dynamics. All else being equal, the larger volatility of S, the worse the performance of the LETF relative to the VETF, with the limiting case $\lim_{\sigma\to\infty} f_\ell(\Delta t; \beta) = 0$.
- Time decay: Even if there is no change in the value of the underlying index, $S(\Delta t) = S(0)$, the value of the LETF tends to zero if held for a long time, since $\lim_{\Delta t\to\infty} f_\ell(\Delta t; \beta) = 0$.
- Costs and interest rates: Expense ratios for LETFs are typically substantially higher than for VETFs, $0 < c_v \ll c_\ell$. In addition, all else being equal, increasing interest rates $r > 0$ also decreases $f_\ell(\Delta t; \beta)$. However, while these effects further decrease the value of the VETF relative to the LETF, they are expected to be comparatively small compared to the other effects.

- Decay due to jumps: It can be shown that the term $\widetilde{Y}_{\ell}(\Delta t;\beta)$, which only affects the LETF, satisfies $\widetilde{Y}_{\ell}(\Delta t;\beta) \leq 1$, where the maximum value ($\widetilde{Y}_{\ell}(\Delta t;\beta) = 1$) is achieved either when there are no jumps in the value of the underlying over $[0, \Delta t]$ (i.e. $\pi(\Delta t) = 0$), or if there are jumps but they all satisfy $\xi_i^s = 1$ which has probability of almost surely zero. In other words, when $S(\Delta t) = S(0)$, the mere presence of jumps in the underlying S decreases the value of the LETF relative to the VETF via the term $\widetilde{Y}_{\ell}(\Delta t;\beta)$. This is illustrated in Figure 3.1 at the point where $S(\Delta t)/S(0) = 1$.

The preceding observations summarize what are effectively the standard objections to LETFs that can be found in the literature, the only addition being the rigorous treatment of jumps in the LETF dynamics and the associated jump decay component.

Next, we discuss lump sum investment strategies, i.e. wealth allocation to assets at time $t_0 = 0$ with no subsequent rebalancing prior to maturity $T = \Delta t$. For simplicity, we consider a constant proportion benchmark strategy $\hat{\mathcal{P}} = \hat{\boldsymbol{p}}(t_0) \coloneqq (1-\hat{p}_s, \hat{p}_s)$, where $\hat{p}_s$ denotes the proportion of benchmark wealth $\hat{W}(t_0) = w_0$ invested in the broad equity market index S at time t_0, with the remaining proportion $(1-\hat{p}_s)$ invested in 30-day T-bills. To emphasize that the benchmark wealth at the end of the investment time horizon depends on $\hat{p}_s$, we use the notation $\hat{W}(\Delta t;\hat{p}_s)$, and observe that

$$\frac{\hat{W}(\Delta t;\hat{p}_s)}{w_0} = (1-\hat{p}_s)\cdot \exp\{r\cdot\Delta t\} + \hat{p}_s \cdot \frac{S(\Delta t)}{S(0)}. \tag{3.14}$$

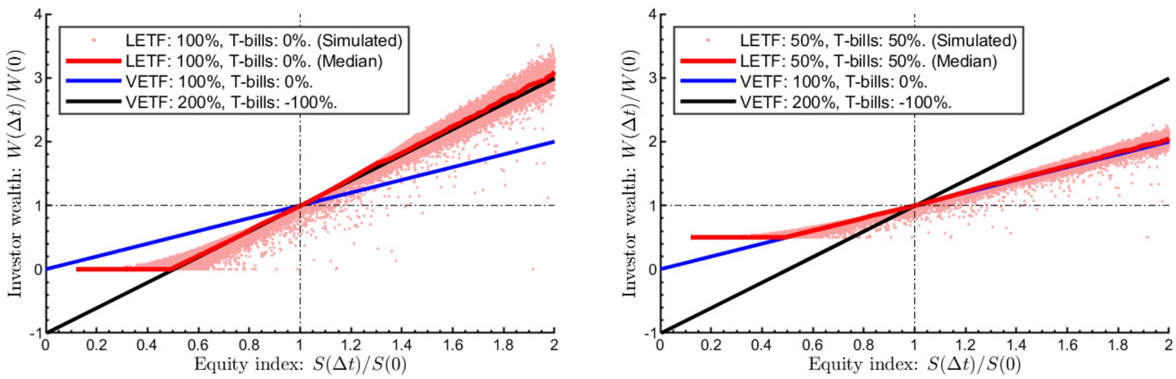

(a) LETF: 100%, T-bills: 0% vs. VETF combinations **(b)** LETF: 50%, T-bills: 50% vs. VETF combinations

Figure 3.1: Payoffs when equity market index S follows calibrated jump-diffusion dynamics (Kou (2002) model): Investor wealth gross return $W(\Delta t)/W(0)$ as a function of underlying equity index gross return $S(\Delta t)/S(0)$, $\Delta t = 0.25$ (1 quarter), for different proportions of initial wealth $W(0)$ invested in the LETF, VETF and T-bills at time $t_0 = 0$. Asset parameters are calibrated to US equity and bond market data over the period 1926:01 to 2023:12 (Appendix B), LETF and VETF expense ratios are assumed to be 0.89% and 0.06% respectively, and a borrowing premium of 3% over the T-bill rate is applicable to short positions.

The investor, being unable to invest directly in S, can combine an ETF investment with T-bills. We will assume that the investor does not short-sell the LETF or VETF[7], but might short-sell the T-bills (i.e. borrow funds) to leverage their investment in the ETF, in which case a constant borrowing premium $b \geq 0$ is added to the T-bill returns. In more detail, if p denotes the fraction of wealth $W(t_0) = w_0$ that the LETF or VETF investors invest in their respective ETFs, an investment fraction $p > 1$ in the ETF is funded by borrowing the amount $(1-p)\cdot w_0$ at an interest rate of $(r+b)$, so the T-bill dynamics applicable to the investors can be modified as

$$\frac{\overline{B}(\Delta t)}{\overline{B}(0)} = \exp\{\overline{r}(p)\cdot\Delta t\}, \qquad \text{where} \qquad \overline{r}(p) = r + b\cdot\mathbb{I}_{[p>1]}, \tag{3.15}$$

with $\mathbb{I}_{[A]}$ denoting the indicator of the event A.

[7] As can be seen from the relationship between the objective functions (2.3) and (2.4), optimizing the IR essentially places us within a variant of the Mean-Variance (MV) framework with a constant risk aversion parameter. In MV optimization (see for example Bensoussan et al. (2014); Van Staden et al. (2018)) with a constant risk aversion parameter, it is never optimal to short-sell the risky asset. The subsequent results of Section 3 and the results of Ni et al. (2024) suggest that this observation also holds our investment scenario, i.e. it is never expected to be IR-optimal to short-sell high return/high volatility assets given the existence of low return/low volatility assets.

The VETF investor (see Assumption 3.1(ii)) specifies an investment strategy $\mathcal{P}_v = \boldsymbol{p}_v(t_0) = (1-p_v, p_v)$, where p_v denotes the fraction of wealth $W(t_0) = w_0$ invested in the VETF F_v at time $t_0 = 0$, and the remaining fraction of wealth $(1-p_v)$ invested in 30-day T-bills. The VETF investor's wealth at the end of the investment time horizon,$W_v(\Delta t; p_v)$ therefore satisfies

$$\frac{W_v(\Delta t; p_v)}{w_0} = (1-p_v)\cdot\exp\{\bar{r}(p_v)\cdot\Delta t\} + p_v\cdot\exp\{-c_v\cdot\Delta t\}\cdot\left(\frac{S(\Delta t)}{S(0)}\right). \tag{3.16}$$

Similarly, the LETF investor specifies investment strategy $\mathcal{P}_\ell = \boldsymbol{p}_\ell(t_0) = (1-p_\ell, p_\ell)$, where p_ℓ is the fraction of wealth $W(t_0) = w_0$ invested in the LETF F_ℓ at time $t_0 = 0$, and the remaining fraction of wealth $(1-p_\ell)$ invested in 30-day T-bills. Using (3.12), the LETF investor's wealth at the end of the investment time horizon,$W_\ell(\Delta t; p_\ell)$ therefore satisfies

$$\frac{W_\ell(\Delta t; p_\ell)}{w_0} = (1-p_\ell)\cdot\exp\{\bar{r}(p_\ell)\cdot\Delta t\} + p_\ell\cdot\exp\{-c_\ell\cdot\Delta t\}\cdot f_\ell(\Delta t;\beta)\cdot\widetilde{Y}_\ell(\Delta t;\beta)\cdot\left(\frac{S(\Delta t)}{S(0)}\right)^\beta. \tag{3.17}$$

Varying p_v and p_ℓ in (3.16) and (3.17) therefore trace out different payoffs for $W_v(\Delta t; p_v)$ and $W_\ell(\Delta t; p_\ell)$ as a function of the (random) underlying index outcome $S(\Delta t)$, with specific choices of p_v and p_ℓ illustrated in Figures 1.1 and 3.1. Note however that the wealth of the VETF investor $W_v(\Delta t; p_v)$ is linear in $S(\Delta t)$, and conditional on $S(\Delta t)$ the outcome $W_v(\Delta t; p_v)$ is deterministic. However, this is not the case for the wealth $W_\ell(\Delta t; p_\ell)$ of the LETF investor, which has a power call-like payoff due to the $[S(\Delta t)]^\beta$ term of (3.17) in conjunction with limited liability. Note that even if we condition on the value of $S(\Delta t)$, the wealth outcome $W_\ell(\Delta t; p_\ell)$ is *not* deterministic due to the presence of the jump term $\widetilde{Y}_\ell(\Delta t;\beta)$in (3.17). However, if no jumps are present then $W_\ell(\Delta t; p_\ell)$ conditional on $S(\Delta t)$ is linear in $[S(\Delta t)]^\beta$, compare Figures 1.1 and 3.1.

Suppose the LETF and VETF investors want to choose values p_v^* and p_ℓ^*, respectively, to maximize the IR (2.2) subject to an implicit target $\gamma > 0$ in (2.4). In this setting of parametric asset dynamics, we can simulate N_d paths of the underlying equity index using (3.10) use each path's information together.

First, we discretize possible values of the fractions p_v and p_ℓ using a fine grid, so that using each discretized value of p_v and p_ℓ, we can obtain the corresponding values of $W_v^{(j)}(\Delta t; p_v)$ and $W_\ell^{(j)}(\Delta t; p_\ell)$ respectively, along each path $j = 1, ..., N_d$. Next, using a discretization of the objective (2.4) in this setting, we can find the approximate IR-optimal values p_v^* and p_ℓ^* by exhaustive search over the grid by solving

$$p_k^* = \arg\min_{p_k}\left\{\frac{1}{N_d}\sum_{j=1}^{N_d}\left(W_k^{(j)}(\Delta t; p_k) - \left[\hat{W}^{(j)}(\Delta t;\hat{p}_s) + \gamma\right]\right)^2\right\}, \qquad k\in\{v,\ell\}. \tag{3.18}$$

Figure 3.2 illustrates the results of this procedure for two different values of the implicit target, $\gamma = 20$ and $\gamma = 50$, where we observe the following:

- In the case of $\gamma = 20$ in Figure 3.2(a), the VETF investor simply invests all wealth in the VETF ($p_v^* =$ 100%), whereas the LETF investor invests slightly less than half of total wealth in the LETF ($p_\ell^* = 48.3\%$). Note that the IR-optimal strategies in this case satisfy $p_v^*/p_\ell^* = 2.070$.
- With a significantly more aggressive benchmark outperformance target of $\gamma = 50$ in Figure 3.2(b), the LETF investor now invests $p_\ell^* = 70.1\%$ in the LETF, i.e. there is no need to leverage the LETF investment itself, whereas the VETF investor borrows 20% of wealth to invest $p_v^* = 120\%$ in the VETF. This leverage is costly for the VETF investor due to the lack of downside protection and borrowing premiums, which can be seen in both the upside and extreme downside outcomes of Figure 3.2(b). In this case, we have $p_v^*/p_\ell^* = 1.712$.

Comparing IR-optimal investment strategies implemented using a LETF or VETF with the same benchmark outperformance target $\gamma > 0$, the observation that $p_v^*/p_\ell^* \approx \beta = 2$ holds in the cases illustrated in Figure 3.2 is not an accident. This relationship is more rigorously discussed in the subsequent results of this section, but for now we note that exact equality $p_v^*/p_\ell^* \equiv \beta$, where β is the returns multiplier of the LETF, only holds for the IR-optimal investment strategies in the case of continuous rebalancing ($\Delta t \downarrow 0$), zero expense ratios ($c_v = c_\ell = 0$), zero borrowing premium over the risk-free rate r ($b = 0$), and when no leverage restrictions are applicable. While this is a very restrictive set of assumptions, $p_v^*/p_\ell^* \approx \beta$ is nevertheless a useful rule-of-thumb to keep in mind when comparing IR-optimal investment strategies in more general cases, a simple example being Figure 3.2. However, it should be emphasized that while the *strategies* might satisfy $p_v^*/p_\ell^* \approx \beta$, this does *not* mean that the ultimate investment *outcomes* for the LETF and VETF investors (such as investor wealth, benchmark outperformance etc.) have straightforward relationships except under very restrictive conditions, since the outcomes are affected by limited liability and the path-dependent role of jumps (see Figure 3.2).

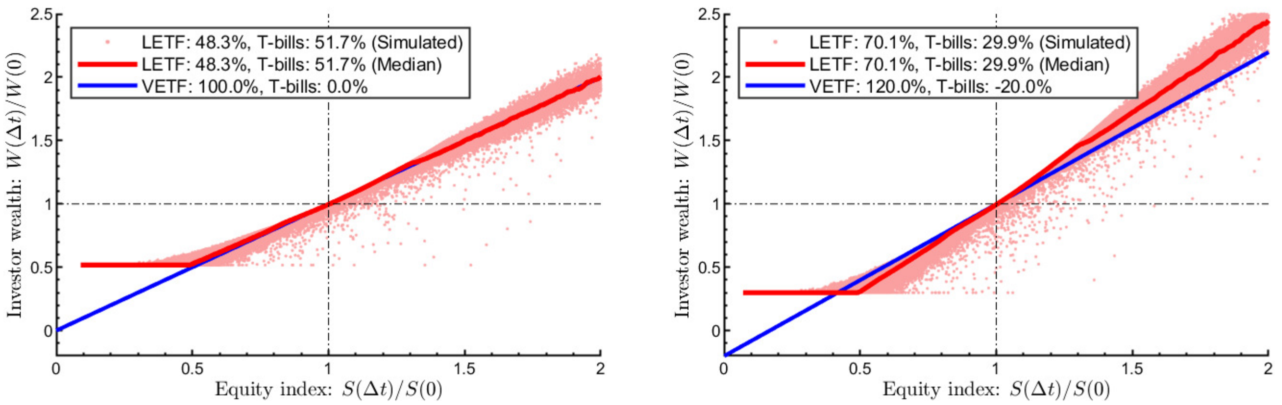

(a) $\gamma = 20$: $p_\ell^* = 48.3\%$, $p_v^* = 100\%$ **(b)** $\gamma = 50$: $p_\ell^* = 70.1\%$, $p_v^* = 120\%$

Figure 3.2: Payoffs when equity market index S follows calibrated jump-diffusion dynamics (Kou (2002) model): Investor wealth gross return $W(\Delta t)/W(0)$ as a function of underlying equity index gross return $S(\Delta t)/S(0)$, $\Delta t = 0.25$ (1 quarter), for different proportions of initial wealth $W(0)$ invested in the LETF, VETF and T-bills at time $t_0 = 0$. Asset parameters are calibrated to US equity and bond market data over the period 1926:01 to 2023:12 (Appendix B), LETF and VETF expense ratios are assumed to be 0.89% and 0.06% respectively, and a borrowing premium of 3% over the T-bill rate is applicable to short positions.

3.2 Dynamically-optimal strategies under continuous rebalancing

Section 3.1 treated the lump-sum investment scenario with no subsequent intervention over $(t_0, T]$, i.e. the set of rebalancing times being simply $\mathcal{T} = [t_0]$. We now consider the other extreme, namely that of continuous rebalancing, where the set of rebalancing times is given by $\mathcal{T} = [t_0, T]$.

Derivation of closed-form optimal strategies necessarily requires stylized assumptions, in this case outlined in Assumption 3.2. As per Remark 3.1, we emphasize that these assumptions are not required for the subsequent results discussed in Section 5.

Assumption 3.2. (Stylized assumptions - continuous rebalancing) In the case of continuous rebalancing, for the purposes of obtaining closed-form solutions in this section, we assume the following:

(i) Assumption 3.1 holds, including parametric dynamics (3.1)-(3.2), (3.4) and (3.8) for the underlying assets.

(ii) The investor injects cash into the portfolio at a constant rate of $q \geq 0$ per year. To ensure the wealth processes remain comparable, the identical rate of cash injection is assumed for the benchmark portfolio.

(iii) We assume continuous portfolio rebalancing ($\mathcal{T} = [t_0, T]$), no investment constraints (i.e. no leverage limits or short-selling constraints), zero borrowing premium so that both borrowing and lending occurs at the risk-free rate r, and trading is allowed to continue in the event of insolvency. Note that these assumptions are standard in the derivation of closed-form solutions of multi-period portfolio optimization problems (see for example Zhou and Li (2000)). This implies that the investment in the LETF can itself be leveraged, which is plausible since even retail investors can borrow and invest in LETFs. However, the degree to which leverage is required by either the LETF or VETF investors depends on the aggressiveness of the outperformance target γ, as shown by the subsequent results. $\square$

Since Assumption 3.1 remains applicable (see Assumption 3.2(i)), the deterministic benchmark strategy allocates wealth to two assets, namely the T-bills B and the broad equity market index S. For the special case of continuous rebalancing, let $t \rightarrow \hat{\varrho}_s(t)$ be a deterministic function of time denoting the proportional allocation to S at time $t \in \mathcal{T} = [t_0, T]$, with $(1 - \hat{\varrho}_s(t))$ denoting the corresponding allocation to T-bills. The benchmark strategy in this section is therefore given by

$$\hat{\mathcal{P}} \quad = \quad \left\{ \hat{\boldsymbol{p}}\left(t, \hat{W}(t)\right) = \left(1 - \hat{\varrho}_s(t),\ \hat{\varrho}_s(t)\right) : t \in [t_0, T] \right\}. \tag{3.19}$$

By Assumption 3.1(ii), in the case of the VETF investor, let $\varrho_v(t, \boldsymbol{X}_v(t))$ denote the proportional allocation of wealth $W_v(t)$ at time t to the VETF F_v in the case of continuous rebalancing, with $\boldsymbol{X}_v(t) = \left(W_v(t), \hat{W}(t), \hat{\varrho}_s(t)\right)$. The VETF investor strategy is therefore of the form

$$\mathcal{P}_v = \left\{\boldsymbol{p}_v(t, \boldsymbol{X}_v(t)) = \left(1 - \varrho_v(t, \boldsymbol{X}_v(t)), \ \varrho_v(t, \boldsymbol{X}_v(t))\right) : t \in [t_0, T]\right\}. \tag{3.20}$$

In the case of the LETF investor, let $\varrho_\ell(t, \boldsymbol{X}_\ell(t))$ denote the proportional allocation of wealth $W_\ell(t)$ at time t to the LETF F_ℓ in the case of continuous rebalancing, with $\boldsymbol{X}_\ell(t) = \left(W_\ell(t), \hat{W}(t), \hat{\varrho}_s(t)\right)$, to obtain the LETF investor strategy as

$$\mathcal{P}_\ell = \left\{\boldsymbol{p}_\ell(t, \boldsymbol{X}_\ell(t)) = \left(1 - \varrho_\ell(t, \boldsymbol{X}_\ell(t)), \ \varrho_\ell(t, \boldsymbol{X}_\ell(t))\right) : t \in [t_0, T]\right\}. \tag{3.21}$$

Given investment strategies of the form (3.19) and (3.21), as well as dynamics (3.1), (3.2), (3.4) and (3.8), we therefore have the following wealth dynamics in the case of continuous rebalancing:

$$\begin{aligned} d\hat{W}(t) &= \left\{\hat{W}(t^-) \cdot \left[r + \hat{\varrho}_s(t)(\mu - \lambda\kappa_1^s - r)\right] + q\right\} \cdot dt \\ &\quad + \hat{W}(t^-)\,\hat{\varrho}_s(t)\,\sigma \cdot dZ(t) + \hat{W}(t^-)\,\hat{\varrho}_s(t) \cdot d\left(\sum_{i=1}^{\pi(t)}(\xi_i^s - 1)\right), \end{aligned} \tag{3.22}$$

$$\begin{aligned} dW_v(t) &= \left\{W_v(t^-) \cdot \left[r + \varrho_v\left(t, \boldsymbol{X}_v(t^-)\right)\left\{(\mu - \lambda\kappa_1^s - r) - c_v\right\}\right] + q\right\} \cdot dt \\ &\quad + W_v(t^-)\,\varrho_v\left(t, \boldsymbol{X}_v(t^-)\right)\sigma \cdot dZ(t) + W_v(t^-)\,\varrho_v\left(t, \boldsymbol{X}_v(t^-)\right) \cdot d\left(\sum_{i=1}^{\pi(t)}(\xi_i^s - 1)\right), \end{aligned} \tag{3.23}$$

$$\begin{aligned} dW_\ell(t) &= \left\{W_\ell(t^-) \cdot \left[r + \varrho_\ell\left(t, \boldsymbol{X}_\ell(t^-)\right)\left\{\beta(\mu - \lambda\kappa_1^s - r) - c_\ell\right\}\right] + q\right\} \cdot dt \\ &\quad + W_\ell(t^-)\,\varrho_\ell\left(t, \boldsymbol{X}_\ell(t^-)\right)\beta\sigma \cdot dZ(t) + W_\ell(t^-)\,\varrho_\ell\left(t, \boldsymbol{X}_\ell(t^-)\right)\beta \cdot d\left(\sum_{i=1}^{\pi(t)}(\xi_i^\ell - 1)\right), \end{aligned} \tag{3.24}$$

for $t \in (t_0, T]$, with initial wealth $W_v(t_0) = W_\ell(t_0) = \hat{W}(t_0) = w_0$. As a reminder, $q \geq 0$ denotes the constant rate per year at which cash is contributed to each portfolio (see Assumption 3.2(ii)), and $\beta > 1$ in (3.24) denotes the multiplier of the LETF.

Due to Assumption 3.2(iii), the set of admissible investor strategies is given by $\varrho_k(t, \boldsymbol{X}_k(t)) \in \mathcal{A}_0$ for $k \in \{v, \ell\}$, where

$$\mathcal{A}_0 = \left\{\varrho_k(t, w, \hat{w}, \hat{\varrho}_s(t)) \,|\, \varrho_k : [t_0, T] \times \mathbb{R}^3 \to \mathbb{R}\right\}, \qquad k \in \{v, \ell\}. \tag{3.25}$$

The IR optimization problem (2.4) in this setting can therefore be written as

$$(IR(\gamma)) : \quad \inf_{\varrho_k \in \mathcal{A}_0} E_{\varrho_k}^{t_0, w_0}\left[\left(W_k(T) - \left[\hat{W}(T) + \gamma\right]\right)^2\right], \qquad \gamma > 0, \quad \text{for } k \in \{v, \ell\}, \tag{3.26}$$

with wealth dynamics (3.22),(3.23) and (3.24) respectively.

The following theorem describes the HJB partial integro-differential equation (PIDE) satisfied by the value function of (3.26) for the LETF investor.

Theorem 3.1. *(IR optimization for the LETF investor: Verification theorem) Let $\gamma > 0$, and assume a given benchmark strategy $t \to \hat{\varrho}_s(t)$ that is deterministic and integrable. Suppose that for all $(t, w, \hat{w}, \hat{\varrho}_s) \in [t_0, T] \times \mathbb{R}^3$, there exist functions $V(t, w, \hat{w}, \hat{\varrho}_s) : [t_0, T] \times \mathbb{R}^3 \to \mathbb{R}$ and $\varrho_\ell^*(t, w, \hat{w}, \hat{\varrho}_s; \gamma) : [t_0, T] \times \mathbb{R}^3 \to \mathbb{R}$ such that: (i) V and ϱ_ℓ^* are sufficiently smooth and solve the HJB PIDE (3.27)-(3.28), and (ii) the pointwise supremum in (3.27) is attained by the function $\varrho_\ell^*(t, w, \hat{w}, \hat{\varrho}_s; \gamma)$.*

$$\frac{\partial V}{\partial t} + \inf_{\varrho_\ell \in \mathbb{R}}\left\{ \mathcal{H}(\varrho_\ell; t, w, \hat{w}, \hat{\varrho}_s) \right\} = 0, \tag{3.27}$$

$$V(T, w, \hat{w}, \hat{\varrho}_s) = (w - \hat{w} - \gamma)^2, \tag{3.28}$$

where

$$\begin{aligned}
\mathcal{H}\left(\varrho_\ell; t, w, \hat{w}, \hat{\varrho}_s\right) &= \left(w \cdot \left[r + \left\{\beta\left(\mu - \lambda\kappa_1^s - r\right) - c_\ell\right\} \cdot \varrho_\ell\right] + q\right) \cdot \frac{\partial V}{\partial w} \\
&\quad + \left(\hat{w} \cdot \left[r + \left(\mu - \lambda\kappa_1^s - r\right) \cdot \hat{\varrho}_s\right] + q\right) \cdot \frac{\partial V}{\partial \hat{w}} \\
&\quad + \frac{1}{2}\left(\varrho_\ell \cdot w\beta\sigma\right)^2 \cdot \frac{\partial^2 V}{\partial w^2} + \frac{1}{2}\left(\hat{\varrho}_s \hat{w}\sigma\right)^2 \cdot \frac{\partial^2 V}{\partial \hat{w}^2} + \left(\varrho_\ell \cdot w\beta\sigma\right)\left(\hat{\varrho}_s \hat{w}\sigma\right) \cdot \frac{\partial^2 V}{\partial w \partial \hat{w}} \\
&\quad - \lambda \cdot V + \lambda \cdot \int_0^\infty V\left(w + \varrho_\ell \cdot w\beta\left(\xi^\ell - 1\right), \hat{w} + \hat{\varrho}_s \hat{w}\left(\xi^s - 1\right), t\right) G\left(\xi^s\right) d\xi^s. \qquad (3.29)
\end{aligned}$$

Then given Assumption 3.2 and wealth dynamics (3.22) and (3.24), V is the value function and ϱ_ℓ^* is the optimal control (i.e. optimal proportion of the investor's wealth to be invested in the LETF with $\beta > 1$) for the $IR(\gamma)$ problem (3.26) for the LETF investor.

Proof. See Appendix A.1. Note that since ξ^ℓ is a function of ξ^s (see (3.6)), the integral in (3.29) is only written with respect to values of ξ^s with associated PDF $G(\xi^s)$. □

Solving the HJB PIDE (3.27)-(3.28), we obtain the IR-optimal investment strategy for the LETF investor as per Proposition 3.2.

Proposition 3.2. *(IR-optimal investment strategy using the LETF) Let $\gamma > 0$ be fixed. Suppose that Assumption 3.2 and wealth dynamics (3.22) and (3.24) apply. Let $W_\ell^*(t)$ denote the LETF investor's wealth process (3.24) under the optimal strategy ϱ_ℓ^*, and let $\boldsymbol{X}_\ell^*(t) = \left(W_\ell^*(t), \hat{W}(t), \hat{\varrho}_s(t)\right)$. Then the IR-optimal fraction of the investor's wealth invested in the LETF, ϱ_ℓ^*, satisfies*

$$\begin{aligned}
&\varrho_\ell^*\left(t, \boldsymbol{X}_\ell^*\left(t^-\right)\right) \cdot W_\ell^*\left(t^-\right) \\
&= \left(\frac{\beta\left[\mu + \lambda\left(\kappa_1^\ell - \kappa_1^s\right) - r\right] - c_\ell}{\beta^2\left(\sigma^2 + \lambda\kappa_2^\ell\right)}\right) \cdot \left[h_\ell(t) + \gamma e^{-r(T-t)} - \left(W_\ell^*\left(t^-\right) - g_\ell(t) \cdot \hat{W}\left(t^-\right)\right)\right] \\
&\quad + \frac{1}{\beta} g_\ell(t) \left(\frac{\sigma^2 + \lambda\kappa_\chi^{\ell,s}}{\sigma^2 + \lambda\kappa_2^\ell}\right) \cdot \hat{\varrho}_s(t) \hat{W}\left(t^-\right), \qquad (3.30)
\end{aligned}$$

where g_ℓ and h_ℓ are the following deterministic functions,

$$g_\ell(t) = \exp\left\{K_\beta^{\ell,s} \cdot \int_t^T \hat{\varrho}_s(u)\, du\right\}, \qquad (3.31)$$

$$h_\ell(t) = -\frac{q}{r}\left(1 - e^{-r(T-t)}\right) + q e^{-r(T-t)} \cdot \int_t^T \exp\left\{r(T-y) + K_\beta^{\ell,s} \cdot \int_y^T \hat{\varrho}_s(u)\, du\right\} dy, \qquad (3.32)$$

with constant $K_\beta^{\ell,s}$ given by

$$K_\beta^{\ell,s} = \mu - r - \frac{\left(\beta\left[\mu + \lambda\left(\kappa_1^\ell - \kappa_1^s\right) - r\right] - c_\ell\right)\left(\sigma^2 + \lambda\kappa_\chi^{\ell,s}\right)}{\beta\left(\sigma^2 + \lambda\kappa_2^\ell\right)}. \qquad (3.33)$$

Proof. See Appendix A.2. □

Note that values of κ_1^ℓ, κ_2^ℓ and $\kappa_\chi^{\ell,s}$ (see (3.7)), which depend on the multiplier β through the LETF jumps (3.6), are required by the optimal strategy (3.30). Expressions for these quantities can be derived in terms of the calibrated parameters of the underlying asset dynamics (3.1)-(3.2) without difficulty. As an illustration, Lemma A.1 in Appendix A.3 presents expressions for κ_1^ℓ, κ_2^ℓ and $\kappa_\chi^{\ell,s}$ in the case of the double-exponential Kou model (Kou (2002)) used for illustrating the results of this section, but we note that this can also be done similarly for other jump diffusion models (e.g. Merton (1976)).

The IR-optimal investment strategy for the VETF investor is given in Corollary 3.3.

Corollary 3.3. *(IR-optimal investment strategy using the VETF) Let $\gamma > 0$ be fixed. Suppose that Assumption 3.2 and wealth dynamics (3.22) and (3.23) apply. Let $W_v^*(t)$ denote the VETF investor's wealth process (3.23) under the optimal strategy ϱ_v^*, and let $\boldsymbol{X}_v^*(t) = \left(W_v^*(t), \hat{W}(t), \hat{\varrho}_s(t)\right)$. Then the IR-optimal fraction of the* *investor's wealth invested in the VETF, ϱ_v^*, satisfies*

$$\begin{aligned}\varrho_v^*\left(t, \boldsymbol{X}_v^*\left(t^-\right)\right) \cdot W_v^*\left(t^-\right) &= \left(\frac{\mu - r - c_v}{\sigma^2 + \lambda \kappa_2^s}\right) \cdot \left[h_v\left(t\right) + \gamma e^{-r(T-t)} - \left(W_v^*\left(t^-\right) - g_v\left(t\right) \cdot \hat{W}\left(t^-\right)\right)\right] \\ &\quad + g_v\left(t\right) \cdot \hat{\varrho}_s\left(t\right) \hat{W}\left(t^-\right),\end{aligned} \tag{3.34}$$

where g_v and h_v are the following deterministic functions,

$$g_v\left(t\right) = \exp\left\{c_v \cdot \int_t^T \hat{\varrho}_s\left(u\right) du\right\}, \tag{3.35}$$

$$h_v\left(t\right) = -\frac{q}{r}\left(1 - e^{-r(T-t)}\right) + q e^{-r(T-t)} \cdot \int_t^T \exp\left\{r\left(T - y\right) + c_v \cdot \int_y^T \hat{\varrho}_s\left(u\right) du\right\} dy. \tag{3.36}$$

Proof. See Appendix A.4. □

The following remark relates the results of Proposition 3.2 and Corollary 3.3 to the available results in the literature.

Remark 3.3. (Relationship of Proposition 3.2 and Corollary 3.3 to results in the literature). In the special case of a VETF with zero expense ratio $c_v = 0$, the results of Corollary 3.3 imply that $g_v\left(t\right) = 1$ and $h_v\left(t\right) = 0$ for all $t \in [t_0 = 0, T]$, so that (3.34) simplifies considerably to

$$\varrho_v^*\left(t, \boldsymbol{X}_v^*\left(t^-\right)\right) \cdot W_v^*\left(t^-\right) = \left(\frac{\mu - r}{\sigma^2 + \lambda \kappa_2^s}\right) \cdot \left[\gamma e^{-r(T-t)} - \left(W_v^*\left(t^-\right) - \hat{W}\left(t^-\right)\right)\right] + \hat{\varrho}_s\left(t\right) \hat{W}\left(t^-\right), \tag{3.37}$$

which corresponds to the IR-optimal investment strategy where direct investment in the underlying equity market index S is possible. This special case (3.37) can be found in Van Staden et al. (2023), where the results of Goetzmann et al. (2002, 2007) are extended to the case of jumps in the risky asset processes. Corollary 3.3 therefore extends this to the case of investing in the equity index *indirectly* via a VETF with a non-negligible expense ratio, whereas Proposition 3.2 extends these results further to the case a LETF with multiplier $\beta > 1$ and expense ratio c_ℓ referencing an equity index S with jump-diffusion dynamics. □

To gain some intuition regarding the behavior of the IR-optimal investment strategies in Proposition 3.2 and Corollary 3.3, we compare the illustrative investment results from implementing these strategies over a 10-year time horizon. We use a benchmark and assets as in Table 3.1, illustrative investment parameters as in Table 3.2, and a Kou model (Kou (2002)) is assumed for the jump diffusion dynamics with calibrated parameters as in Appendix B (Table B.1). Since LETFs are a relatively recent invention, we follow the example of Bansal and Marshall (2015) in constructing a proxy LETF replicating $\beta = 2$ times the daily returns of a broad stock market index, in this case using the CRSP VWD index, which is a capitalization-weighted index consisting of all domestic stocks trading on major US exchanges, with historical data available since January 1926. As in for example Bansal and Marshall (2015) and Leung and Sircar (2015), we assume that the managers of the LETF do not have challenges in replicating the underlying index, which is reasonable given the possibility of designing replication strategies for LETFs that remain robust even during periods of market volatility (see for example Guasoni and Mayerhofer (2023)). For more information on the source data and calibrated, inflation-adjusted parameters, please refer to Appendix B.

Table 3.2: Closed-form solutions - Investment parameters for illustrating the results. Note that the calibrated parameters for the jump-diffusion process are given in Appendix B (Table B.1), while the underlying assets, benchmark and ETF expense ratios are given in Table 3.1.

Parameter	T	w_0	q	γ
Value	10 years	\$ 100	\$ 5 per year	125

Figure 3.3 illustrates the IR-optimal proportion of wealth invested in the ETF as a function of time t and the wealth difference $W_k^*\left(t\right) - \hat{W}\left(t\right)$, $k \in \{v, \ell\}$. In the case of the LETF investor, Figure 3.3(a) illustrates ϱ_ℓ^*, whereas for the VETF investor, Figure 3.3(b) illustrates ϱ_v^*, where both strategies use the same target $\gamma = 125$.

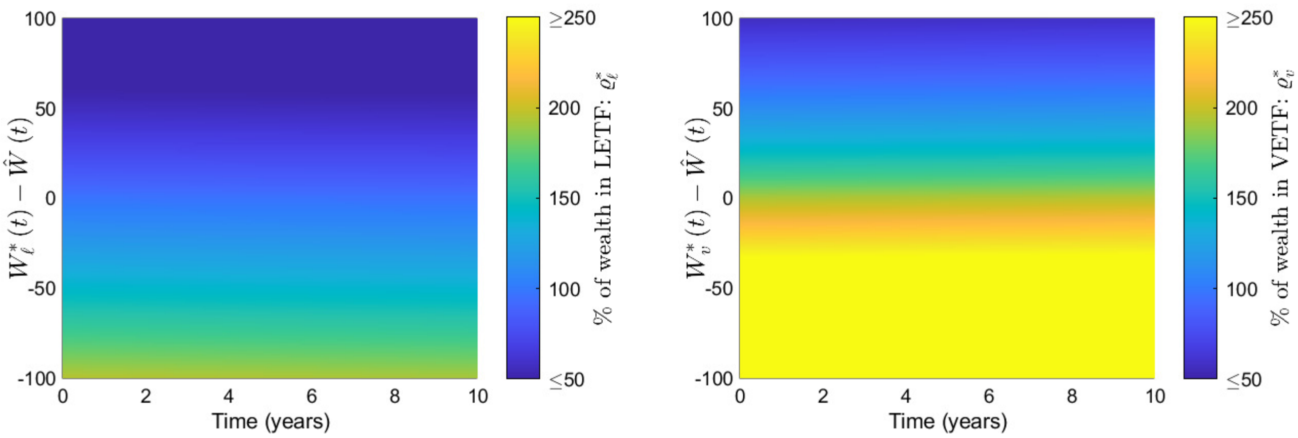

(a) IR-optimal investment in LETF ($\beta = 2$), ϱ_ℓ^* **(b)** IR-optimal investment in VETF, ϱ_v^*

Figure 3.3: Closed-form IR-optimal investment strategies using the LETF (ϱ_ℓ^* as per (3.30)) or the VETF (ϱ_v^* as per (3.34)) as a function of time t and the wealth difference $W_k^*(t) - \hat{W}(t)$, $k \in \{v, \ell\}$, given the same implicit benchmark outperformance target γ. The underlying assets, benchmark, investment parameters and calibrated process parameters are as in Table 3.1, Table 3.2 and Table B.1, respectively. Note that the same color scale is used in both figures for comparison purposes.

Using the same strategies as illustrated in Figure 3.3, Figure 3.4 shows the *ratio* of IR-optimal proportions of wealth invested in the VETF relative to the investment in the LETF, $\varrho_v^*/\varrho_\ell^*$, given an otherwise identical wealth difference $W_k^*(t) - \hat{W}(t), k \in \{v, \ell\}$ at time t. We emphasize that both Figure 3.3 and Figure 3.4 treat the IR-optimal strategies from Proposition 3.2 and Corollary 3.3 simply as functions of time and the wealth difference relative to the benchmark.

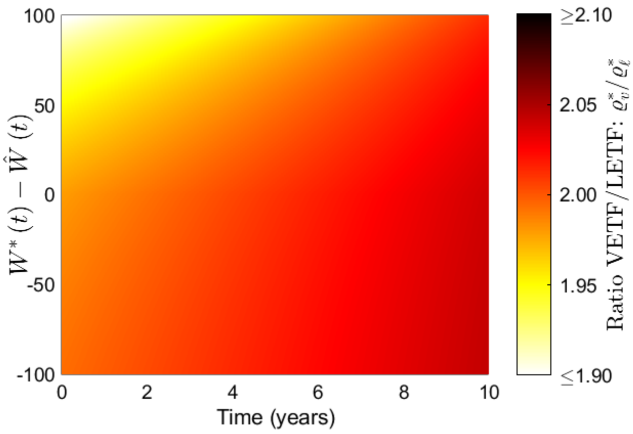

Figure 3.4: Ratio of IR-optimal proportions of wealth in the VETF vs. the LETF, $\varrho_v^*/\varrho_\ell^*$, given identical wealth differences relative to the benchmark $W^*(t) - \hat{W}(t) \equiv W_v^*(t) - \hat{W}(t) = W_\ell^*(t) - \hat{W}(t)$ at each time t and same target γ. The underlying assets, benchmark, investment parameters and calibrated process parameters are as in Table 3.1, Table 3.2 and Table B.1, respectively.

With regards to Figure 3.3 and Figure 3.4, we make the following observations regarding the IR-optimal investment strategies of the LETF investor vs. the VETF investor:

(i) The IR-optimal investment strategy using a LETF (Figure 3.3(a)) is, like the strategy using a VETF (Figure 3.3(b)), fundamentally *contrarian*. Specifically, in the case of the LETF, we observe that the IR-optimal proportion of wealth ϱ_ℓ^* in the LETF decreases as the wealth difference $W_\ell^*(t) - \hat{W}(t)$ increases, which happens after a period of strong LETF return performance. In their analysis of reports to the SEC by institutional fund managers, DeVault et al. (2021) show that institutional investors indeed empirically tend to decrease their holdings in LETFs following periods of strong investment performance. While DeVault et al. (2021) concludes that this behavior might be explained as being a result of compensationbased incentives, our results show that strategies based on maximizing the IR (a widely-used investment metric) relative to a standard investment benchmark could also be related to this empirical investment behavior.

(ii) Figure 3.4 shows that the IR-optimal strategies under stylized assumptions (Assumption 3.1) and identical implicit benchmark outperformance target γ satisfy $\varrho_v^*/\varrho_\ell^* \approx \beta = 2$. Informally, the LETF and VETF investors therefore take on nearly identical "risk" exposure to the movements of the underlying index, which provides valuable intuition when interpreting the results Section 5 where Assumption 3.2 is relaxed.

While Figure 3.3 and Figure 3.4 illustrate the IR-optimal strategy as a function of time and the wealth difference relative to the benchmark, implementing this strategy in a Monte Carlo simulation provides an additional perspective. Figure 3.5 shows the median and 95th percentiles of the IR-optimal proportion of wealth invested in the LETF and VETF, given the same benchmark outperformance target γ. Note that under the stylized assumptions (Assumption 3.2), the LETF and VETF positions can be leveraged without restriction, with leverage constraints only subsequently introduced in Section 4. We observe that the corresponding ETF exposure percentiles tend to approximately satisfy $\varrho_v^*/\varrho_\ell^* \approx \beta = 2$, with decreasing exposure over time due to the contrarian nature of both strategies.

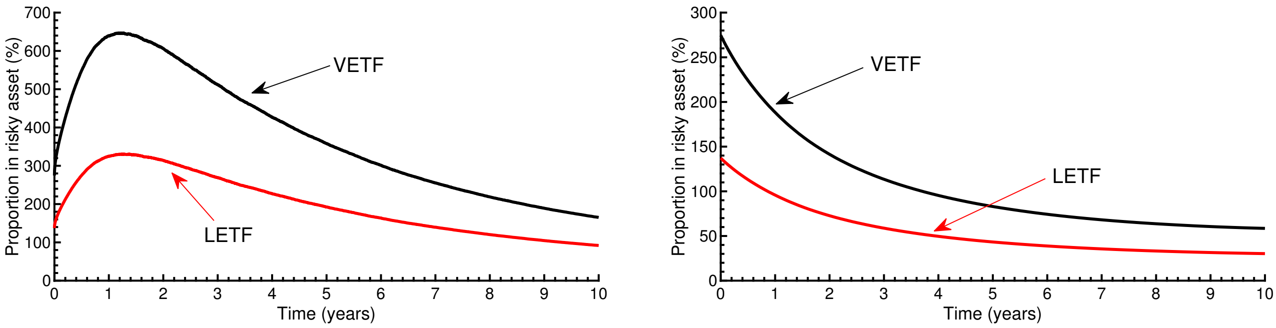

(a) ϱ_ℓ^* and ϱ_v^*: 95th percentiles over time

(b) ϱ_ℓ^* and ϱ_v^*: 50th percentiles over time

Figure 3.5: Closed-form IR-optimal investment strategies: 95th and 50th percentiles over time of the IR-optimal proportion of wealth invested in the LETF (ϱ_ℓ^* as per (3.30)) or the VETF (ϱ_v^* as per (3.34)) obtained using Monte Carlo simulation of the underlying dynamics (3.22)-(3.24), and same target γ. The underlying assets, benchmark, investment parameters and calibrated process parameters are as in Table 3.1, Table 3.2 and Table B.1, respectively.

Figure 3.6 compares the simulated CDFs of IR-optimal terminal wealth $W_k^*(T), k \in \{v, \ell\}$ and CDFs of the terminal wealth ratio relative to the benchmark $W_k^*(T)/\hat{W}(T), k \in \{v, \ell\}$ for the LETF and VETF investors, respectively.

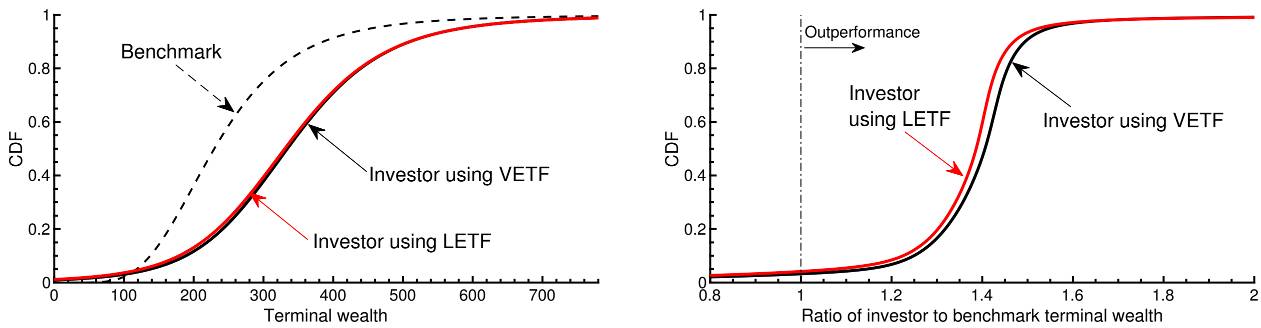

(a) CDFs of $\hat{W}(T), W_k^*(T)$, $k \in \{v, \ell\}$

(b) CDFs of ratio $W_k^*(T)/\hat{W}(T)$, $k \in \{v, \ell\}$

Figure 3.6: Closed-form IR-optimal investment strategies: CDFs of the IR-optimal terminal wealth for the same target γ obtained using Monte Carlo simulation of the underlying dynamics and investing according to optimal strategies (3.30) and (3.34). The underlying assets, benchmark, investment parameters and calibrated process parameters are as in Table 3.1, Table 3.2 and Table B.1, respectively.

Figure 3.6 shows that due to implicitly similar exposure levels to movements in the underlying equity index (since $\varrho_v^* \simeq \beta \cdot \varrho_\ell^*$) given identical targets γ for the LETF and VETF investors, implementing the strategies illustrated in Figure 3.3 result in nearly identical terminal wealth and outperformance outcomes. In fact, under the stylized assumptions of this section, Proposition 3.4 shows that in the special case of (i) zero expense ratios and (ii) no jumps in the S-dynamics, we have $\varrho_v^* \equiv \beta \cdot \varrho_\ell^*$, the IR-optimal investor should be entirely indifferent as to whether the LETF-based strategy (3.30) or VETF-based strategy (3.34) is used for investment purposes.

Proposition 3.4. *(Special case: Zero expense ratios, no jumps) Let $\gamma > 0$ be fixed, and let Assumption 3.2 and wealth dynamics (3.22)-(3.24) apply. If (i) both LETF ($\beta > 1$) and the VETF have zero expense ratios, i.e. $c_v = c_\ell = 0$, and (ii) there are no jumps in the underlying S-dynamics (i.e. $\lambda = 0$ in (3.2)), then following results hold:*

(i) The IR-optimal proportion of wealth invested in the VETF (3.34) is equal to β times the IR-optimal proportion of wealth invested in the LETF (3.30),

$$\varrho_v^*\left(t, \boldsymbol{X}_v^*(t)\right) \quad = \quad \beta \cdot \varrho_\ell^*\left(t, \boldsymbol{X}_\ell^*(t)\right), \qquad \forall t \in [t_0, T], \tag{3.38}$$

where $\boldsymbol{X}_v^*(t) = \left(W_v^*(t), \hat{W}(t), \hat{\varrho}_s(t)\right)$ and $\boldsymbol{X}_\ell^*(t) = \left(W_\ell^*(t), \hat{W}(t), \hat{\varrho}_s(t)\right)$.

(ii) An IR-optimal investor with given target $\gamma > 0$ would be indifferent as to whether the optimal strategy is executed with a LETF (3.30) or VETF (3.34), since wealth outcomes are identical,

$$W_\ell^*(t) \quad = \quad W_v^*(t), \qquad \forall t \in [t_0, T], \tag{3.39}$$

and the same Information Ratio is obtained,

$$\frac{E_{\varrho_\ell^*}^{t_0, w_0}\left[W_\ell^*(T) - \hat{W}(T)\right]}{Stdev_{\varrho_\ell^*}^{t_0, w_0}\left[W_\ell^*(T) - \hat{W}(T)\right]} = \frac{E_{\varrho_v^*}^{t_0, w_0}\left[W_v^*(T) - \hat{W}(T)\right]}{Stdev_{\varrho_v^*}^{t_0, w_0}\left[W_v^*(T) - \hat{W}(T)\right]} = \left(\exp\left\{\left(\frac{\mu - r}{\sigma}\right)^2 \cdot T\right\} - 1\right)^{1/2}. \tag{3.40}$$

Proof. See Appendix A.5. □

We summarize the closed-form solutions results presented in this section as follows. In the case of continuous rebalancing (i.e. rebalancing times $\mathcal{T} = [t_0, T]$) with no investment constraints, the IR-optimal investor should be largely indifferent whether a VETF or LETF is used on the same underlying equity index to execute the IR-optimal investment strategy involving the ETF and T-bills for a given outperformance target γ. Since the IR-optimal investment strategies satisfy the approximate relationship $\varrho_v^* \simeq \beta \cdot \varrho_\ell^*$ (Figure 3.4), when continuously rebalancing the LETF and VETF investors effectively maintain similar implicit risk exposures at each time instant to movements of the underlying index, resulting in broadly similar investment outcomes (Figure 3.6), with differences between outcomes entirely driven by different ETF expense ratios and the presence of jumps (see Proposition 3.4).

At the other extreme, namely the lump-sum investment scenario with no subsequent intervention (i.e. $\mathcal{T} = [t_0]$) and one-quarter time horizon ($\Delta t = 0.25$), we still observe the approximate relationship $p_v^* \approx \beta \cdot p_\ell^*$ between corresponding IR-optimal strategies (Figure 3.2). However, in this case the power call-like payoff of the LETF can clearly be observed, whereby the LETF investor benefits from an upside due to comparatively inexpensive leverage while simultaneously enjoying downside protection due to limited liability (Figures 1.1, 3.1 and 3.2).

Far from ignoring the standard criticisms of LETFs in the literature (see Section 3.1), the closed form results of this section - illustrated using parametric dynamics calibrated to empirical market data - suggest that we should not be entirely surprised that LETFs might have substantial appeal to investors *despite* these shortcomings. However, we emphasize that none of the trading strategies illustrated - not even the lump-sum investment scenario results - advocate for simplistic strategies like buy-and-hold positions in the LETF over indefinite time horizons, so a certain degree of sophistication on the part of the investor is implicitly assumed. In addition, while the closed-form results provide valuable intuition, they were derived under stylized assumptions, which we relax in the subsequent sections to model the potential of LETFs under more realistic conditions.

4 Numerical solutions

To assess the potential value of LETFs in designing investment strategies for benchmark outperformance under more reasonable assumptions than in Section 3, analytical solutions are typically no longer obtainable and a numerical solution technique is instead required.

This section starts by formulating a more realistic investment setting with the following characteristics: (i) Restrictions on the maximum leverage and limitations on short-selling. (ii) An optional borrowing premium applicable when short-selling an asset. (iii) Prohibition of trading in insolvency. (iv) Infrequent (discrete) rebalancing of the portfolio. (v) Direct use of market data without the need to specify parametric dynamics for the underlying assets. This is followed by a brief overview of a neural network-based numerical solution approach to solve the IR problem (2.4) in this setting. Indicative investment results obtained by implementing these techniques on empirical market data are discussed in Section 5.

4.1 Investment constraints and discrete rebalancing

Recall from Section 2 that the investor's strategy is based on investing in a set of N_a candidate assets indexed by $i \in \{1, .., N_a\}$, while the benchmark is defined in terms of $\hat{N}_a$ potentially different underlying assets indexed by $j \in \left\{1, .., \hat{N}_a\right\}$. With investor and benchmark strategies of the form (2.1), and the benchmark strategy satisfying only the general assumptions outlined in Assumption 2.1, we assume that both the investor and benchmark portfolios are rebalanced at each of N_{rb} discrete rebalancing times during the investment time horizon $[t_0 = 0, T]$. As a result, $\mathcal{T}$ is now of the form

$$\mathcal{T} \quad = \quad \{ t_n = n\Delta t | \, n = 0, ..., N_{rb} - 1 \}, \qquad \Delta t = T/N_{rb}. \tag{4.1}$$

The assumption of equally-spaced rebalancing times in (4.1) is only for convenience, and can be relaxed without difficulty. At each rebalancing time $t_n \in \mathcal{T}$, we assume a pre-specified cash contribution $q(t_n)$ is made to the investor portfolio, with the contribution also being added to the benchmark portfolio to ensure the comparison in performance remains appropriate.

There is no need to specify any parametric dynamics for the underlying assets in the numerical solution approach, which only requires the availability of empirical market data for deriving the optimal strategy (see Subsection 4.2 below). Specifically, we assume that at each time $t_{n+1} \in \mathcal{T} \cup T$ we can observe $R_i(t_n)$ and $\hat{R}_j(t_n)$, the returns on investor asset $i \in \{1, .., N_a\}$ and benchmark asset $j \in \left\{1, .., \hat{N}_a\right\}$, respectively, over the time interval $[t_n, t_{n+1}]$. Note that these returns might be inflation-adjusted and might include a borrowing premium applicable to assets that have been shorted (see e.g. Assumption 4.1 below). As a result, for the purposes of numerical solutions, the investor and benchmark wealth dynamics are respectively of the form

$$W\left(t_{n+1}^-\right) \quad = \quad \left[W\left(t_n^-\right) + q\left(t_n\right)\right] \cdot \sum_{i=1}^{N_a} p_i\left(t_n, \boldsymbol{X}\left(t_n^-\right)\right) \cdot \left[1 + R_i\left(t_n\right)\right], \tag{4.2}$$

$$\hat{W}\left(t_{n+1}^-\right) \quad = \quad \left[\hat{W}\left(t_n^-\right) + q\left(t_n\right)\right] \cdot \sum_{j=1}^{\hat{N}_a} \hat{p}_j\left(t_n, \hat{\boldsymbol{X}}\left(t_n^-\right)\right) \cdot \left[1 + \hat{R}_j\left(t_n\right)\right], \tag{4.3}$$

where $W\left(t_0^-\right) = \hat{W}\left(t_0^-\right) \coloneqq w_0 > 0$ and $n = 0, ..., N_{rb} - 1$.

Since active funds often have restrictions on leverage and short-selling (see for example Forsyth et al. (2019); Ni et al. (2024)), these constraints are included in the formulation.

The investor's candidate asset $i = 1$, assumed to be 30-day T-bills in the indicative investment results of Section 5, plays a special role in leveraged portfolios. The investor is assumed to be able to short-sell this asset with a potential borrowing premium payable, i.e. the investor can borrow funds at an approximation of the prevailing short-term interest rate plus a borrowing premium to fund leveraged investments in the other assets. In addition, in the case of insolvency, defined as occurring when the investor wealth is negative, $W(t_n) < 0$, we will assume that the negative wealth (i.e. the outstanding debt) is placed in asset $i = 1$, where it grows at the rate of return of this asset with an addition of a possible borrowing premium. Note that this effectively implies that trading ceases when $W < 0$, either until maturity T or until such a time where the cash injections pay off the debt resulting in $W > 0$, in which case trading can resume. A maximum leverage ratio at a portfolio level of p_{max} is also assumed, where typical values are in the range $p_{max} \in [1.0, 1.5]$. Assumption 4.1 outlines the details more formally.

Assumption 4.1. (Investor strategy: Leverage restrictions, borrowing premium and no trading in insolvency) The following assumptions and restrictions apply to the investor strategy, where the investor considers investment in $N_a \geq 2$ candidate assets. As discussed, the investor's set of candidate assets may not correspond to the assets included in the benchmark strategy.

(i) **Shortable and long-only assets**: Only investor candidate asset $i = 1$ will (potentially) be shorted, with the remaining investor candidate assets $i \in \{2, .., N_a\}$ being long only. In other words, at any rebalancing time $t_n \in \mathcal{T}$, we have

$$
\begin{aligned}
\text{(Shortable asset } i = 1) : \quad & p_1\left(t_n, \boldsymbol{X}\left(t_n^-\right)\right) \in \mathbb{R}, \quad & , t_n \in \mathcal{T}, \quad & (4.4) \\
\text{(Long-only assets } i \in \{2, .., N_a\}) : \quad & p_i\left(t_n, \boldsymbol{X}\left(t_n^-\right)\right) \geq 0, \quad & i \in \{2, .., N_a\}, t_n \in \mathcal{T}, \quad & (4.5) \\
\text{(All wealth invested)} : \quad & \sum_{i=1}^{N_a} p_i\left(t_n, \boldsymbol{X}\left(t_n^-\right)\right) = 1, \quad & t_n \in \mathcal{T}. \quad & (4.6)
\end{aligned}
$$

(ii) **Borrowing premium** $b \geq 0$: If investor candidate asset $i = 1$ is shorted at time t_n (i.e. if $p_1\left(t_n, \boldsymbol{X}\left(t_n^-\right)\right) < 0$), then a constant borrowing premium $b \geq 0$ is added to the returns on asset $i = 1$ over the time interval $[t_n, t_{n+1}]$ to be paid by the investor. In other words, for asset $i = 1$, the return $R_1(t_n)$ incorporated in (4.2) is of the form

$$
\text{(Borrowing premium)} : \qquad R_1(t_n) = \begin{cases} \overline{R}_1(t_n), & \text{if } p_1\left(t_n, \boldsymbol{X}\left(t_n^-\right)\right) \geq 0 \\ \overline{R}_1(t_n) + b, & \text{if } p_1\left(t_n, \boldsymbol{X}\left(t_n^-\right)\right) < 0, \end{cases} \qquad (4.7)
$$

where $\overline{R}_1(t_n)$ is the (possibly inflation-adjusted) return on underlying asset $i = 1$ over $[t_n, t_{n+1}]$ without any added premiums. For long-only assets, we simply have $R_i(t_n) = \overline{R}_i(t_n)$, $i \in \{2, .., N_a\}, t_n \in \mathcal{T}$. Note that in the case of the benchmark strategy, no borrowing premium is applicable to any asset due to Assumption 4.2 below.

(iii) **Maximum leverage ratio** p_{max}: The total allocated proportion of wealth to the long-only assets $i \in \{2, .., N_a\}$ cannot exceed the maximum leverage ratio p_{max} ,

$$
\text{(Maximum leverage ratio)} : \qquad \sum_{i=2}^{N_a} p_i\left(t_n, \boldsymbol{X}\left(t_n^-\right)\right) \leq p_{max}, \qquad t_n \in \mathcal{T}. \qquad (4.8)
$$

(iv) **No trading in insolvency**: If the investor wealth is negative, i.e. if $W(t_n) < 0$ at any $t_n \in \mathcal{T}$, then all long asset positions (4.5) are liquidated and the total debt (the amount $W(t_n) < 0$) is allocated to the shortable asset (4.4). In such a scenario, no further trading occurs for the remainder of the investment time horizon ($t_m \in \mathcal{T}, t_m > t_n$), unless cash injections pay off the debt, and the portfolio wealth becomes positive. Total debt accumulates at a rate (4.7) which possibly includes a borrowing premium. More formally,

$$
\text{(No trading in insolvency)} : \qquad \text{If } W\left(t_n^-\right) < 0 \quad \Rightarrow \quad \boldsymbol{p}\left(t_n, \boldsymbol{X}\left(t_n^-\right)\right) = \boldsymbol{e}_1, \quad t_n \in \mathcal{T}, \qquad (4.9)
$$

where $\boldsymbol{e}_1 = (1, 0, ..., 0) \in \mathbb{R}^{N_a}$ is the standard basis vector $\mathbb{R}^{N_a}$ with 1 in the first position (corresponding to $i = 1$, the shortable asset as per (4.4)) and all other entries are zero. □

Note that (4.9) also implies $\boldsymbol{p}\left(t_m, \boldsymbol{X}\left(t_m^-\right)\right) = \boldsymbol{e}_1$ for all $t_m > t_n$, so that no further trading does indeed occur in the case of insolvency as required by Assumption 4.1(iv).

Recalling that $\mathcal{A}$ denotes the set of admissible controls and $\mathcal{Z}$ denoting the admissible control space, Assumption 4.1 implies that we have the following form for $\mathcal{Z}$ and $\mathcal{A}$, respectively:

$$
\mathcal{Z} \;=\; \left\{ \boldsymbol{z} \in \mathbb{R}^{N_a} \;\middle|\; \begin{array}{c} z_1 \in \mathbb{R}, \\ z_i \geq 0, \forall i \in \{2, .., N_a\}, \\ \sum_{i=1}^{N_a} z_i = 1, \\ \sum_{i=2}^{N_a} z_i \leq p_{max}. \end{array} \right\}, \qquad (4.10)
$$

and

$$\mathcal{A} \quad = \quad \left\{ \mathcal{P} = \left\{ \boldsymbol{p}\left(t_n, \boldsymbol{X}\left(t_n\right)\right), t_n \in \mathcal{T} \,\middle|\, \begin{array}{l} \boldsymbol{p}\left(t_n, \boldsymbol{X}\left(t_n\right)\right) \in \mathcal{Z}, \text{ if } W\left(t_n^-\right) \geq 0, \\ \boldsymbol{p}\left(t_n, \boldsymbol{X}\left(t_n\right)\right) = \boldsymbol{e}_1, \text{ if } W\left(t_n^-\right) < 0. \end{array} \right\} \right\}. \tag{4.11}$$

Note that with slight abuse of notation in (4.11), $\mathcal{Z}$ is the admissible control space in the case of solvency only.

Finally, in order ensure that the benchmarks align with typical investment benchmarks used in practice (see Remark 2.1) and to avoid pathological examples, Assumption 4.2 below specifies that no short-selling is allowed in the case of the benchmark strategy.

Assumption 4.2. (Benchmark: leverage restrictions) The benchmark strategy does not engage in the short-selling of any asset.

$$\text{(Long-only benchmark)} \qquad \hat{p}_j\left(t_n, \hat{\boldsymbol{X}}\left(t_n^-\right)\right) \geq 0, \qquad \forall j = 1, ..., \hat{N}_a. \tag{4.12}$$

As a result, the benchmark strategy has an implicit maximum leverage ratio of $p_{max} = 1$, with no borrowing premium being applicable, while benchmark insolvency is ruled out in the sense that $\hat{W}\left(t_n\right) \geq 0$ for all $t_n \in \mathcal{T}$ given dynamics (4.3). □

4.2 Neural network solution approach

The objective function in the case of the numerical solutions remains of the form (2.4),

$$\left(IR\left(\gamma\right)\right): \qquad \inf_{\mathcal{P} \in \mathcal{A}} E_{\mathcal{P}}^{t_0, w_0}\left[\left(W\left(T\right) - \left[\hat{W}\left(T\right) + \gamma\right]\right)^2\right], \qquad \gamma > 0, \tag{4.13}$$

with the main differences from the treatment in Section 3 being the following: (i) The set of admissible controls $\mathcal{A}$ is now given by (4.11). (ii) Rebalancing occurs at a strict discrete subset of times $t_n \in \mathcal{T} \subset [t_0 = 0, T]$. (iii) As discussed below, we no longer need the assumption of parametric models for the underlying assets, but can use market data directly.

To solve (4.13) numerically to obtain the optimal investment strategy $\mathcal{P}^* \in \mathcal{A}$, we follow the neural network-based solution approach of Ni et al. (2024), where a "leverage-feasible neural network" (LFNN) is constructed to approximate the investment strategy directly as a feedback control $\left(t_n, \boldsymbol{X}\left(t_n\right)\right) \rightarrow \mathcal{P}\left(t_n, \boldsymbol{X}\left(t_n\right)\right) := \boldsymbol{p}\left(t_n, \boldsymbol{X}\left(t_n\right)\right), \forall t_n \in \mathcal{T}$ in the case of admissible sets of the form (4.10)-(4.11). This approach forms part of a class of methods (see, for example, Buehler et al. (2019); Han and Weinan (2016); Mäkinen and Toivanen (2024); Reppen and Soner (2023); Reppen et al. (2023); Van Staden et al. (2023, 2024),) that does not require dynamic programming to solve problems such as (2.4), thereby avoiding the typical challenges such as evaluating high-dimensional conditional expectations and error amplifications over time-stepping.

Since more detailed information, including a convergence analysis, can be found in Ni et al. (2024), we give only a very short overview of the application of the LFNN approach in our setting. In this approach, the control *function* $\left(t_n, \boldsymbol{X}\left(t_n\right)\right) \rightarrow \boldsymbol{p}\left(t_n, \boldsymbol{X}\left(t_n\right)\right)$ is approximated by a single neural network (NN) with at least 3 features (inputs), namely $\left(t_n, \boldsymbol{X}\left(t_n\right)\right) = \left(t_n, W\left(t_n\right), \hat{W}\left(t_n\right)\right)$. Note that additional features such as trading signals can be incorporated in the NN inputs in settings where this is considered valuable. Let $\boldsymbol{F}(t, \boldsymbol{X}(t); \boldsymbol{\theta}) \equiv \boldsymbol{F}(\cdot, \boldsymbol{\theta})$ denote the NN, where $\boldsymbol{\theta} \in \mathbb{R}^{\eta_\theta}$ is the NN parameters, i.e. the NN weights and biases. Since the time t_n is used is an input into the NN, a single parameter vector $\boldsymbol{\theta}$ (equivalently, a single NN) is applicable to all rebalancing times, identifying this as a "global-in-time" approach in the taxonomy of Hu and Laurière (2023). One of the key contributions of Ni et al. (2024) is to construct the NN $\boldsymbol{F}(\cdot, \boldsymbol{\theta})$ with an output layer that guarantees, for all inputs $\left(t, \boldsymbol{X}\left(t\right)\right) = \left(t, W\left(t\right), \hat{W}\left(t\right)\right)$, that

$$\left(t, \boldsymbol{X}\left(t\right)\right) \quad \rightarrow \quad \boldsymbol{F}(t, \boldsymbol{X}(t); \boldsymbol{\theta}) \,\middle|\, \begin{array}{l} \boldsymbol{F}(t, \boldsymbol{X}(t); \boldsymbol{\theta}) \in \mathcal{Z}, \text{ if } W\left(t\right) \geq 0, \\ \boldsymbol{F}(t, \boldsymbol{X}(t); \boldsymbol{\theta}) = \boldsymbol{e}_1, \text{ if } W\left(t\right) < 0. \end{array} \tag{4.14}$$

As a result of (4.11), by using the approximation

$$\boldsymbol{p}(t, \boldsymbol{X}(t)) \quad \simeq \quad \boldsymbol{F}(t, \boldsymbol{X}(t); \boldsymbol{\theta}) \equiv \boldsymbol{F}(\cdot, \boldsymbol{\theta}), \tag{4.15}$$

we can therefore approximate the investor strategies as $\mathcal{P} = \left\{\boldsymbol{F}(t_n, \boldsymbol{X}(t_n); \boldsymbol{\theta}),\ t_n \in \mathcal{T}\right\}$ while being assured that

$\mathcal{P} \in \mathcal{A}$ where $\mathcal{A}$ is as per (4.11), without the need to impose constraints on the optimization problem itself. As a result, (4.13) can be solved as an *unconstrained* optimization problem over $\boldsymbol{\theta} \in \mathbb{R}^{\eta_\theta}$,

$$\inf_{\boldsymbol{\theta} \in \mathbb{R}^{\eta_\theta}} E^{t_0,w_0}_{\boldsymbol{F}(\cdot;\boldsymbol{\theta})} \left[\left(W\left(T;\boldsymbol{\theta}\right) - \left[\hat{W}\left(T\right) + \gamma \right] \right)^2 \right], \tag{4.16}$$

where the approximation (4.15) is used to obtain the asset allocation in the investor wealth dynamics (4.2) which depend on $\boldsymbol{\theta}$.

Parametric models for the underlying asset dynamics are no longer required. Instead, we use a finite set of samples from the set $Y = \left\{ Y^{(j)} : j = 1, ..., N_d \right\}$, where each element $Y^{(j)}$ denotes a time series of *joint* asset return observations $R_i, i \in \{1, .., N_a\}$, possibly adjusted for inflation and the application of a borrowing premium, observed at each $t_n \in \mathcal{T}$. Y represents the training data of the NN, so any $\boldsymbol{\theta} \in \mathbb{R}^{\eta_\theta}$ and returns path $Y^{(j)} \in Y$, the wealth dynamics (4.2) with approximation (4.15) generates a terminal wealth outcome $W^{(j)}\left(T;\boldsymbol{\theta}\right)$. The expectation in (4.16) is then approximated simply by

$$\min_{\boldsymbol{\theta} \in \mathbb{R}^{\eta_\theta}} \left\{ \frac{1}{N_d} \sum_{j=1}^{N_d} \left(W^{(j)}\left(T;\boldsymbol{\theta}\right) - \left[\hat{W}^{(j)}\left(T\right) + \gamma \right] \right)^2 \right\}, \tag{4.17}$$

where the optimal parameter vector $\boldsymbol{\theta}^*$ is obtained using stochastic gradient descent. The resulting IR-optimal strategy for (4.13) consistent with the constraints as outlined in Assumption 4.1 is therefore given by $\boldsymbol{p}^*(\cdot, \boldsymbol{X}(\cdot)) \simeq \boldsymbol{F}(\cdot, \boldsymbol{\theta}^*)$.

While the details underlying the construction of the data set Y are clearly of practical significance, we note that the approach of Ni et al. (2024) remains agnostic as to the how Y is constructed. It can be obtained using for example GAN-generated data sets (see e.g. Van Staden et al. (2024); Yoon et al. (2019)), or using Monte Carlo simulations if the underlying dynamics are specified for ground truth analysis purposes (see e.g. Van Staden et al. (2023)), or a version of bootstrap resampling of empirical market data, as we now discuss.

In practical applications, the use of empirical market data might be preferred for the construction of Y. However, since only a single historical path of asset returns is available, some form of data augmentation is typically used to obtain sufficiently rich training and testing data. For illustrative purposes, in Section 5 we use stationary block bootstrap resampling (Politis and Romano (1994)) to construct Y. This technique, designed for weakly stationary time series with serial dependence, is both popular in academic settings (Anarkulova et al. (2022)) and practical applications (Cavaglia et al. (2022); Cogneau and Zakalmouline (2013); Dichtl et al. (2016); Scott and Cavaglia (2017); Simonian and Martirosyan (2022)). Note that bootstrap resampling methods have been proposed for non-stationary time series (Politis (2003), Politis et al. (1999)), but this is not used in the illustrative investment results of Section 5.

5 Indicative investment results

The main objective of this section is to demonstrate the potential role of LETFs within long-term, diversified, dynamic, IR-optimal investment strategies subject to the investment constraints outlined in Section 4. As a result, we focus entirely on applying the numerical approach discussed in Subsection 4.2 to obtain the IR-optimal strategies based on the combination of a LETF or a VETF on a broad equity market index with bonds (T-bills and T-bonds).

5.1 Investment scenarios

The key investment parameters used for illustrative purposes throughout this section are outlined in Table 5.1. Note in particular that we use a relatively long investment time horizon (10 years) coupled with relatively infrequent (quarterly) rebalancing, and that the same implicit outperformance target γ is used in each of the scenarios to facilitate a fair comparison. This value of γ is chosen for general illustrative purposes only, and the conclusions remain qualitatively similar for different choices of γ.

Table 5.2 provides an overview of the benchmark and the investor's candidate assets. A 70/30 benchmark strategy is again used, since it aligns to the definition of popular investment benchmarks used in practice (see Remark 2.1). Note that the benchmark is defined in terms of the broad equity market index ("Market") with 70% of the wealth allocation, with the remaining 30% split between 30-day T-bills and 10-year T-bonds. As

Table 5.1: Key investment parameters for the illustrative results of Section 5.

Parameter	T	# rebalancing events (N_{rb})	Initial wealth (w_0)	Contributions (q_n)	Target (γ)
Value	10 years	40 (quarterly rebalancing)	\$ 100	\$ 5 per year (\$1.25/quarter)	125

in Section 3, we assume that the investor cannot invest directly in the broad equity market index ("Market"), but can gain exposure to this index via a VETF (expense ratio $c_v = 0.06\%$) or a LETF with multiplier $\beta = 2$ (expense ratio $c_\ell = 0.89\%$).

Table 5.2: Candidate assets and benchmark for the illustrative results of Section 5. A mark "✓" indicates that an asset is available for inclusion. Note that the investor cannot invest directly in the market portfolio ("Market"), but only indirectly via either the VETF or LETF, whereas the benchmark is defined directly in terms of "Market" in alignment with popular investment benchmarks.

Underlying assets		Benchmark	Investor candidate assets	
Label	Asset description		Using VETF	Using LETF
T30	30-day Treasury bill	15%	✓	✓
B10	10-year Treasury bond	15%	✓	✓
Market	Market portfolio (broad equity market index)	70%	-	-
VETF	Vanilla (unleveraged) ETF replicating the returns of the market portfolio, with expense ratio $c_v = 0.06\%$	-	✓	-
LETF	Leveraged ETF with daily returns replicating $\beta = 2$ times the daily returns of the market portfolio, with expense ratio $c_\ell = 0.89\%$	-	-	✓

Table 5.3 provides more detail on the leverage and borrowing premium scenarios considered, where we highlight the following:

- Investor portfolios formed with a LETF are never leveraged ($p_{max} = 1.0$), whereas portfolios formed with a VETF can use leverage up to a portfolio maximum of $p_{max} \in \{1.0, 1.2, 1.5, 2.0\}$ via the short-selling of 30-day T-bills (i.e. borrowing funds to invest in the VETF) with a borrowing premium $b \in \{0, 0.03\}$ potentially being applicable. This is done in order to compare the performance of an IR-optimal portfolio with a LETF and no portfolio-level leverage with that of an IR-optimal portfolio formed with a (potentially) leveraged VETF under various leverage assumptions.

- In terms of the selection of values for $p_{max} \in \{1.0, 1.2, 1.5, 2.0\}$ in the case of the VETF investor, note that Regulation T of the US Federal Reserve board requires at least 50% of the initial price of a stock position to be available on deposit, while brokerage firms are free to establish more stringent requirements. For the VETF investor, for illustrative purposes we will therefore mostly focus on the cases of $p_{max} = 1.0$ (no leverage) or $p_{max} = 1.5$, and for comparison purposes provide the additional examples using $p_{max} = 1.2$ and $p_{max} = 2.0$ in Appendix C.

- In terms of the selection of borrowing premiums $b \in \{0, 0.03\}$ for the VETF investor, we first note that all returns are inflation-adjusted (see Appendix B), and so these quantities should be interpreted net of inflation. The case of zero borrowing premium ($b = 0$) is provided for comparison purposes only, while the value of $b = 3\%$ is obtained from the examples in Ni et al. (2024), where it is based on a consideration of the average real return for T-bills and the average inflation-adjusted corporate bond yields for Moody's Aaa and Baa-rated bond issues.

The underlying data sets for the training and testing of the neural network giving the IR-optimal investment strategies using stationary block bootstrap resampling of empirical market data (see Section 4 and Appendix B) instead of calibrated process dynamics. In particular, we use all available inflation-adjusted market data over the time period January 1926 to December 2023, together with an expected block size of 3 months, to obtain 500,000 jointly bootstrapped asset return paths (see Li and Forsyth (2019); Van Staden et al. (2024) and Appendix B for more information). As in Ni et al. (2024), we use a shallow NN (2 hidden layers) with only the

Table 5.3: Maximum leverage and borrowing premium scenarios for the indicative investment results of Section 5.

Component of leverage scenario	Benchmark	Investor candidate assets	
		Using VETF	Using LETF
Maximum portfolio-level leverage ratio p_{max}	No leverage allowed ($p_{max} = 1.0$)	No leverage allowed ($p_{max} = 1.0$) as well as scenarios $p_{max} \in \{1.2, 1.5, 2.0\}$	No leverage allowed ($p_{max} = 1.0$)
Shortable asset to fund leveraged position (if applicable)	-	T30	-
Borrowing premiums: Scenarios for premium b over T30 return on leveraged positions (if applicable)	N/a	$b = 0$ or $b = 0.03$	N/a

minimal input features $(t, \boldsymbol{X}(t)) = \left(t, W(t), \hat{W}(t)\right)$, since that has been found sufficient to obtain a stable and accurate IR-optimal investment strategy in a setting where no additional market signals are used as inputs.

For illustrative purposes, we also present the investment results obtained from investing according to the IR-optimal investment strategy on selected historical data paths. This is discussed in more detail in Remark 5.1.

Remark 5.1. (Performance on single historical data paths) Since the future evolution of asset returns are not expected to replicate the past evolution of returns *precisely*, we consider illustrative investment results based on bootstrapped data sets as significantly more informative than using a single historical data path of asset returns to illustrate performance.

However, for purposes of concreteness and intuition, we do show the evolution of the LETF and VETF investor wealth obtained by implementing the corresponding IR-optimal portfolios and the benchmark on four historical data paths each spanning a period equal to the investment time horizon of 10 years:

(i) January 2000 until December 2009, which illustrates the impact on the portfolio wealth of both the DotCom bubble crash as well as the GFC period.

(ii) January 2005 until December 2014, which focuses on the GFC and the subsequent period of relatively slow market recovery.

(iii) January 2010 until December 2019, which illustrates the performance during the bull market of the 2010s, a period of very low interest rates and therefore cheap leverage.

(iv) January 2014 until December 2023, which combines an initial period of strong growth and low interest rates with the Covid-19 period and subsequent recovery, only to be followed by the bear market for stocks lasting from January to October 2022 and higher interest rates.

Note that while the historical path of returns enter the training data of the NN indirectly via bootstrap resampling, the probability that the actual historical data path itself appearing in the resulting bootstrapped data sets is vanishingly small (see Ni et al. (2022) for a proof), so that the historical data paths can themselves be considered as effectively "out-of-sample" for testing purposes. However, we emphasize that in this section the main focus remains on the investment results based on the much richer bootstrapped data set of returns data, which ensures a meaningful discussion of the implications for wealth *distributions*, for example, rather than individual wealth values from a single historical path. □

5.2 Comparison of investment results

Figure 5.1 illustrates the distributions of the IR-optimal terminal investor wealth $W_k^*(T), k \in \{v, \ell\}$ for different portfolios formed under the leverage scenarios as per Table 5.3, as well as the distribution of the benchmark terminal wealth $\hat{W}(T)$. With regards to Figure 5.1(a), we see that the IR-optimal strategy using the LETF (and bonds) achieves *partial stochastic dominance* (Atkinson (1987); Ni et al. (2024); Van Staden et al. (2021)) over the IR-optimal strategies using the VETF (and bonds), even if the VETF investment can be leveraged. Note that all IR-optimal strategies achieve partial stochastic dominance over the benchmark, which is to be expected considering the results of Van Staden et al. (2023).

While Figure 5.1 considers a maximum leverage ratio of $p_{max} = 1.5$ with a borrowing premium for short-selling of $b = 3\%$ (applicable to the leveraged VETF position), the results of Figure C.1 in Appendix C show

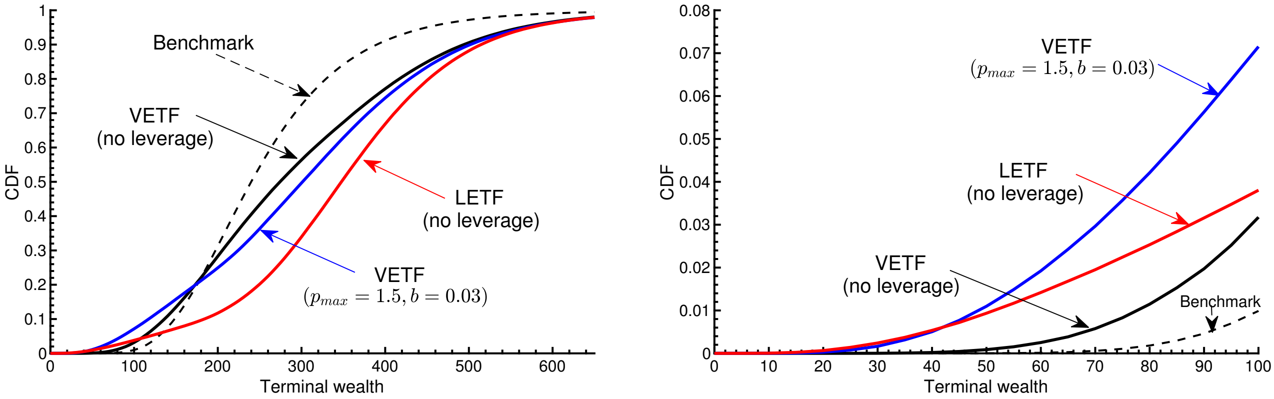

(a) CDFs of $W^*(T)$, $W_k^*(T)$, $k \in \{v, \ell\}$ **(b)** Extreme left tail: CDFs of $\hat{W}(T)$, $W_k^*(T)$, $k \in \{v, \ell\}$

Figure 5.1: CDFs of IR-optimal terminal wealth $W_k^*(T)$, $k \in \{v, \ell\}$ and the benchmark terminal wealth $\hat{W}(T)$.

that qualitatively similar results are obtained even if there are no borrowing premiums on short-selling, and leverage is allowed to increase to $p_{max} = 2.0$. At first glance, this might be somewhat surprising. In the case of continuous rebalancing, no investment constraints and zero costs, Proposition 3.4 shows that the IR-optimal investment results are identical, regardless of whether the investor uses a LETF with multiplier β or a VETF leveraged β times. However, the underlying assumptions of Proposition 3.4 are violated in the setting of this section, where we have discrete rebalancing and investment constraints. Therefore, while we expect to see at least some difference in the IR-optimal investment results obtained using the LETF and VETF-based strategies, we might not expect a difference of the magnitude seen in for example Figure 5.1 or Figure C.1.

The main driver of the difference in performance of the LETF-based strategy relative to that of the VETF-based strategy in this setting is a combination of (i) the call-like payoff of the LETF as underlying asset over a relatively short time horizon (e.g. 1 quarter) and (ii) the contrarian nature of the discretely-rebalanced IR-optimal investment strategy locking in the gains from rebalancing. First, holding the LETF position for a quarter amounts to holding a "continuously rebalanced" position in the equity index and bonds, resulting in the power law-type payoff discussed in Section 3.1. Note that the results of Section 3.1 were obtained using parametric models for the underlying assets. In contrast, in this section, we consider a data set obtained using the stationary block bootstrap resampling of historical data. As an aid to the intuition as to the implications of switching to bootstrapped historical data, Figure 5.2(a) compares the pathwise quarterly returns to two simple strategies using this data, namely investing all wealth in the LETF at the start of a quarter, as well as investing 200% of wealth in the VETF at the start of a quarter funded by borrowing 100% of wealth at the T-bill rate, and comparing outcomes at the end of the quarter with no intermediate trading. While the actual IR-optimal investment strategies are clearly not that straightforward, Figure 5.2(a) confirms that the call-like payoff of the LETF also holds in the bootstrapped historical data, resulting in a slight potential advantage relative to the 2x leveraged VETF strategy (Figure 5.2(b)). Given this payoff structure of the LETF, as discussed in Section 3, the IR-optimal LETF strategy responds to gains by reducing exposure to the LETF, thus locking in the results of prior quarters of good performance while reducing exposure to future possible losses by having lower exposure to the LETF. The compounding effect of applying the contrarian IR-optimal investment strategy quarter after quarter given returns to underlying assets as per Figure 5.2 ultimately results in the strong performance of the LETF-based strategy illustrated in Figure C.1 in Appendix C (for zero borrowing premiums and maximum leverage of $p_{max} = 2.0$) and the more realistic results in Figure 5.1 (where a borrowing premium of $b = 3\%$ applies and maximum leverage is $p_{max} = 1.5$).

However, as Figure 5.1(b) illustrates, there is no free lunch with regards to leverage, similar to what we observed in Figures 1.1, 3.1 and 3.2. In more detail, when considering the extreme left tails of the IR-optimal terminal wealth CDFs, whether leveraging an investment implicitly (via the LETF) or explicitly (via a leveraged VETF investment), the downside wealth outcomes are worse than using the VETF with no leverage (or simply the benchmark). Note that this is based on the distribution of empirical market data together with the implementation of IR-optimal investment strategies, and is not inconsistent with the observations regarding the downside protection offered by the LETFs in truly extreme cases (illustrated in Figures 3.2).

Figure 5.3 focuses on different measures of benchmark outperformance rather than investor wealth, with

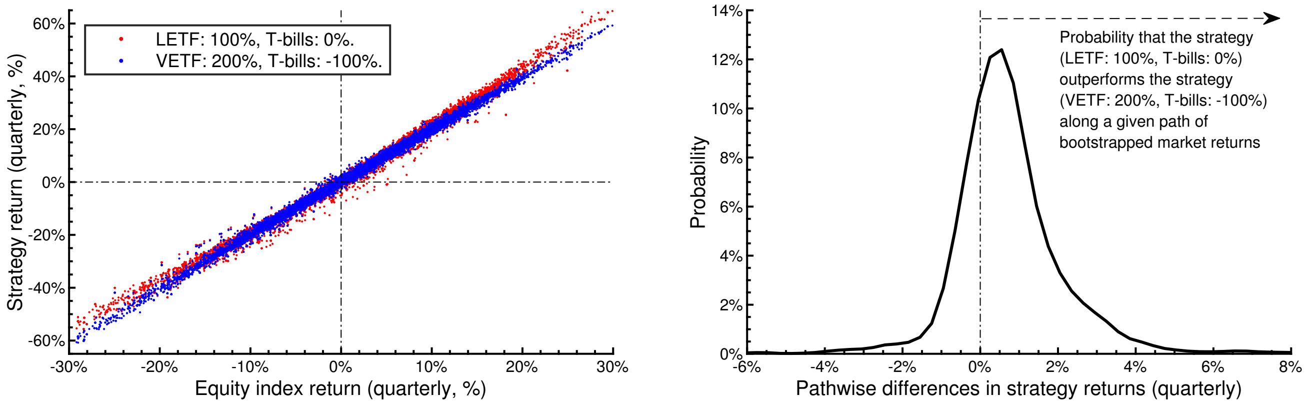

(a) Pathwise payoffs of simple strategies **(b)** Distribution - pathwise return differences

Figure 5.2: Pathwise comparison of the quarterly inflation-adjusted returns of two simple strategies using empirical data obtained by means of the stationary block bootstrap resampling of historical data as described in Section 5.1 and Appendix B. The strategies consist of (i) investing all wealth in the LETF at the start of a quarter and (ii) investing 200% of wealth in the VETF at the start of the quarter funded by borrowing 100% of wealth at the T-bill rate, and comparing outcomes at the end of the quarter with no intermediate trading. Figure 5.2(a) illustrates the quarterly returns of the simple strategies (y-axis) for a given level of equity index quarterly return (x-axis) over the quarter, which demonstrates qualitative similarities to the theoretical payoffs seen in Figure 3.1 in the case of jumps in the underlying equity market index. Figure 5.2(b) shows the distribution of pathwise quarterly return differences where, for each value of the x-axis in Figure 5.2(a), and therefore for a particular given path of (joint) asset returns, we calculate the vertical difference between the return of the simple strategy (i) where all wealth is invested in the LETF, minus the return of strategy (ii) where the 200% of wealth investment in the VETF is funded by borrowing at the T-bill rate. Figure 5.2(b) shows that the simple LETF-based strategy is empirically somewhat more likely to outperform the VETF-based strategy along any given path of underlying returns over a single quarter, although the return difference distribution in Figure 5.2(b) has a relatively small median value of only 59bps.

Figure 5.3(a) illustrating the CDF of the terminal pathwise wealth ratio $W_k^*(T)/\hat{W}(T), k \in \{v, \ell\}$ and Figure 5.3(b) illustrating the probability of benchmark outperformance over time. It is clear that IR-optimal portfolios formed using the LETF and no further leverage significantly improves the benchmark outperformance characteristics of the resulting strategy. Note that the results of Appendix C (Figure C.2 and Figure C.3) show that the conclusions of Figure 5.3 remain qualitatively applicable for different leverage and borrowing cost scenarios.

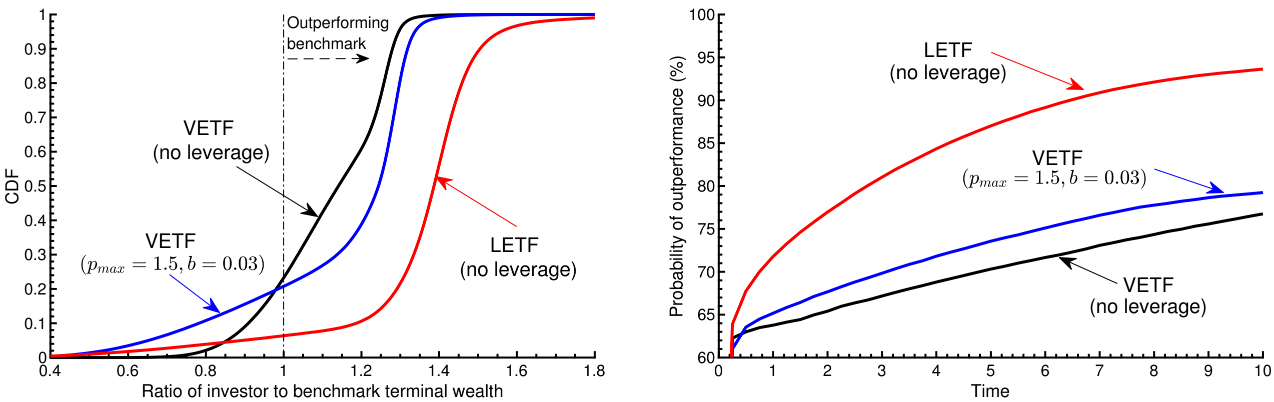

(a) CDF of ratio $W_k^*(T)/\hat{W}(T), k \in \{v, \ell\}$ **(b)** Probability $W_k^*(t) > \hat{W}(t), k \in \{v, \ell\}$ as a function of t

Figure 5.3: CDFs of IR-optimal terminal wealth ratios $W_k^*(T)/\hat{W}(T), k \in \{v, \ell\}$, and probability $W_k^*(t) > \hat{W}(t), k \in \{v, \ell\}$ of benchmark outperformance as a function of time t.

Figure 5.4 compares selected percentiles of the IR-optimal proportion of wealth in the LETF (no leverage) or leveraged VETF ($p_{max} = 1.5, b = 0.03$) over time, using the same scale in Figure 5.4(a) and Figure 5.4(b) for illustrative purposes. Figure 5.4(a) shows that the LETF investor initially allocates around 70% of wealth allocated to LETF, which quickly falls to around 40% or less around the middle of the investment time horizon for both the 20th and 50th percentiles. In the case of the portfolio with a leveraged position in the VETF

($p_{max} = 1.5$), Figure 5.4(b) shows that the median allocation to the VETF exceeds 100% of wealth for more than half the investment time horizon.

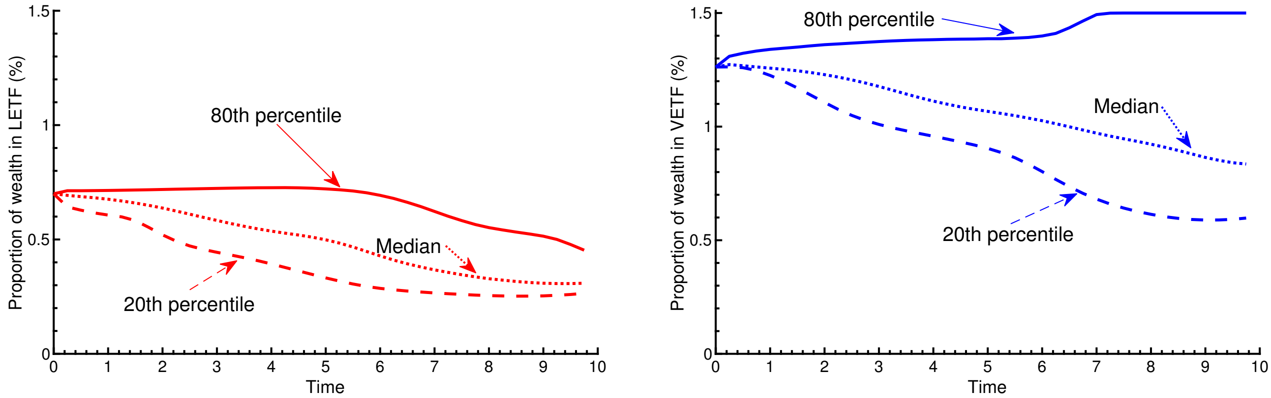

(a) Proportion of wealth in LETF (no leverage) **(b)** Proportion of wealth in VETF ($p_{max} = 1.5, b = 0.03$)

Figure 5.4: Numerical solutions, Scenario 1, basic assets only (LETF or VETF and bonds - see Table 5.2): Selected percentiles of the IR-optimal proportion of wealth in the LETF (no leverage) or leveraged VETF ($p_{max} = 1.5, b = 0.03$) over time. Note the same scale on the y-axis has been used to facilitate comparison.

Figure 5.4 shows that executing the IR-optimal strategy using a LETF offers the investor more flexibility: with lower levels of wealth tied up in the LETF compared to the allocation to a VETF, together with the higher volatility of LETF-based returns, the investor can "lock in" periods of good past returns by implementing a systematic de-risking of the portfolio to a lower allocation to LETF over time, increasing the allocation to bonds. While the leveraged VETF-based strategy essentially follows the same contrarian pattern, the 80th percentile in Figure 5.4(b) shows that it significantly harder for the leveraged VETF strategy to recover from periods of poor past returns in a setting of maximum leverage restrictions, borrowing costs and no trading in the event of bankruptcy.

In addition to the preceding results which are based on the bootstrap resampling of historical data, for illustrative purposes we consider the investment performance on selected single historical paths (see Remark 5.1) illustrated in Figure 5.5. Figures 5.5(a) and (b) show that both the LETF and VETF investors (regardless of VETF leverage) underperform the benchmark during the lowest points of the DotCom and GFC crashes, with the LETF investor experiences larger peak-to-trough declines but also faster post-crash recovery. Figure 5.5(c) illustrate that the LETF-based strategy remains only slightly ahead of the VETF-based strategies during periods of strong equity market performance and low interest rates, while Figure 5.5(d) shows that the LETF investor stays ahead despite the significant impact on portfolio wealth of the Covid-19 period and subsequent bear market of 2022.

6 Conclusion

In this paper, we investigated the potential of including a broad stock market index-based leveraged ETF (LETF) in long-term, dynamically-optimal investment strategies designed to maximize the outperformance over standard performance benchmarks in terms of the information ratio (IR).

Using both closed-form and numerical solutions, we showed that an investor can exploit the observation that LETFs offer call-like payoffs, and therefore could be a convenient way to add inexpensive leverage to the portfolio while providing extreme downside protection.

Under stylized assumptions including continuous rebalancing and no investment constraints, we derived the closed-form IR-optimal investment strategy for the LETF investor, which provided valuable intuition as to the contrarian nature of the strategy. In more practical settings of quarterly trading, leverage restrictions, no trading in the event of insolvency and the presence of margin costs on borrowing, we employed a neural network-based approach to determine the IR-optimal strategies. Our findings show that unleveraged IR-optimal strategies with a broad stock market LETF not only outperform the benchmark more often than possibly leveraged IR-optimal strategies derived using a VETF, but can achieve partial stochastic dominance over the benchmark and

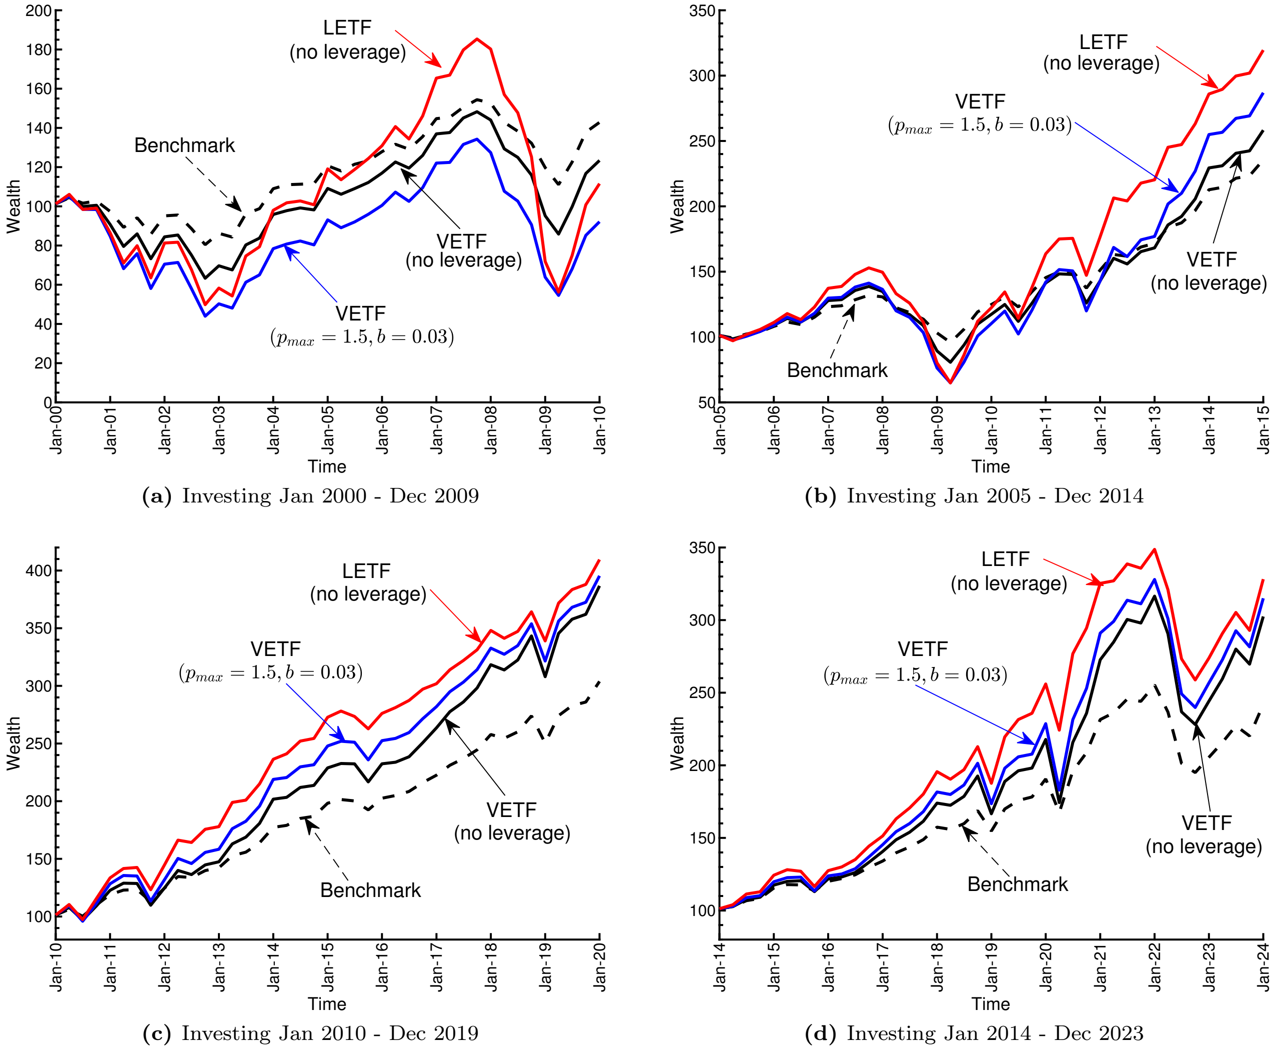

(a) Investing Jan 2000 - Dec 2009

(b) Investing Jan 2005 - Dec 2014

(c) Investing Jan 2010 - Dec 2019

(d) Investing Jan 2014 - Dec 2023

Figure 5.5: Evolution of portfolio wealth over time when investing according to the corresponding IR-optimal investment strategies on historical paths selected for the reasons as discussed in Remark 5.1.

(leveraged or unleveraged) VETF-based strategies in terms of terminal wealth.

Two important caveats are to be kept in mind regarding our results demonstrating the potential of LETFs: (i) The results and conclusions are associated with dynamic IR-optimal investment strategies, which are most emphatically *not* naive strategies like the buy-and-hold strategies over long time horizons often considered in the literature (see the Introduction for a discussion). In particular, critical to the investment outcomes are the rebalancing of the portfolio within the context of a contrarian investment strategy. (ii) The results emphasize that there is no free lunch with regards to leverage. Specifically, the extreme left tails of the IR-optimal terminal wealth CDFs confirm that whether leveraging an investment implicitly (via the LETF) or explicitly (via a leveraged VETF investment), the downside wealth outcomes are worse than using the VETF without any leverage, and therefore the upside outcomes of leverage is not without significant risks. Nevertheless, bootstrap resampling tests indicate that use of an optimal strategy using LETFs outperforms the benchmark $> 95\%$ of the time, which may make the extreme tail risk acceptable.

Despite the controversy surrounding the uses of LETFs for investment purposes in the literature, our results help to explain the empirical appeal of LETFs to institutional and retail investors alike, and encourage a reconsideration of the role of broad stock market LETFs within the context of more sophisticated investment strategies.

7 Declaration

Forsyth's work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) grant RGPIN-2017-03760. Li's work was supported by NSERC grant RGPIN-2020-04331.

The authors have no conflicts of interest to report.

References

Ahn, A., M. Haugh, and A. Jain (2015). Consistent pricing of options on leveraged ETFs. *SIAM Journal on Financial Mathematics* 6, 559–593.

Alekseev, A. G. and M. V. Sokolov (2016). Benchmark-based evaluation of portfolio performance: a characterization. *Annals of Finance* 12, 409–440.

Anarkulova, A., S. Cederburg, and M. S. O'Doherty (2022). Stocks for the long run? Evidence from a broad sample of developed markets. *Journal of Financial Economics* 143:1, 409–433.

Applebaum, D. (2004). *Lévy processes and stochastic calculus.* Cambridge University Press.

Atkinson, A. (1987). On the measurement of poverty. *Econometrica* 55(4), 749–764.

Avellaneda, M. and S. Zhang (2010). Path-dependence of leveraged ETF returns. *SIAM Journal on Financial Mathematics* 1(1), 586–603.

Bajeux-Besnainou, I., R. Portait, and G. Tergny (2013). Optimal portfolio allocations with tracking error volatility and stochastic hedging constraints. *Quantitative Finance* 13(10), 1599–1612.

Bansal, V. K. and J. F. Marshall (2015). A tracking error approach to leveraged ETFs: Are they really that bad? *Global Finance Journal* 26, 47–63.

Bednarek, Z. and P. Patel (2022). Just say no to leveraged ETFs. *Journal of Investment Management* 20(3), 53–69.

Bensoussan, A., K. C. Wong, S. C. P. Yam, and S. P. Yung (2014). Time-consistent portfolio selection under short-selling prohibition: From discrete to continuous setting. *SIAM Journal on Financial Mathematics* 5, 153–190.

Bertrand, P. and J.-L. Prigent (2022). On the diversification and rebalancing returns: performance comparison of constant mix versus buy and hold strategies SSRN 4153690.

Bolshakov, A. and L. B. Chincarini (2020). Manager skill and portfolio size with respect to a benchmark. *European Financial Management* 26(1), 176–197.

Buehler, H., L. Gonon, J. Teichmann, and B. Wood (2019). Deep hedging. *Quantitative Finance* 19(8), 1271–1291.

Canadian Pension Plan (2022). Our investment strategy. `www.cppinvestments.com/the-fund/how-we-invest/our-investment-strategy`.

Carver, A. B. (2009). Do leveraged and inverse ETFs converge to zero? *ETFs and Indexing* 1, 144–149.

Cavaglia, S., L. Scott, K. Blay, and S. Hixon (2022). Multi-asset class factor premia: A strategic asset allocation perspective. *The Journal of Portfolio Management* 48:9, 14–32.

Charupat, N. and P. Miu (2011). The pricing and performance of leveraged exchange-traded funds. *Journal of Banking and Finance* 35(4), 966–977.

Cheng, M. and A. Madhavan (2009). The dynamics of leveraged and inverse exchange-traded funds. *Journal of Investment Management* 7, 43–62.

Cogneau, P. and V. Zakalmouline (2013). Block bootstrap methods and the choice of stocks for the long run. *Quantitative Finance* 13, 1443–1457.

Dai, M., S. Kou, H. M. Soner, and C. Yang (2023). Leveraged exchange-traded funds with market closure and frictions. *Management Science* 69(4), 2517–2535.

Dang, D. and P. 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.

DeVault, L., H. Turtle, and K. Wang (2021). Blessing or curse? institutional investment in leveraged ETFs. *Journal of Banking and Finance* 129(106169).

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

Forsyth, P. (2020). Multiperiod mean conditional value at risk asset allocation: Is it advantageous to be time consistent? *SIAM Journal on Financial Mathematics* 11(2), 358–384.

Forsyth, P. and K. Vetzal (2017). Dynamic mean variance asset allocation: Tests for robustness. *International Journal of Financial Engineering* 4:2. 1750021 (electronic).

Forsyth, P., K. Vetzal, and G. Westmacott (2019). Management of portfolio depletion risk through optimal life cycle asset allocation. *North American Actuarial Journal* 23(3), 447–468.

Giese, G. (2010). On the risk-return profile of leveraged and inverse ETFs. *Journal of Asset Management* 11, 219–228.

Goetzmann, W., J. Ingersoll, M. Spiegel, and I. Welch (2002). Sharpening Sharpe ratios. *NBER Working Paper 9116* .

Goetzmann, W., J. Ingersoll, M. Spiegel, and I. Welch (2007). Manipulation and manipulation-proof performance measures. *Review of Financial Economics* 20, 1503–1546.

Government Pension Fund Global (2022). Investment strategy. `www.nbim.no/en/the-fund/how-we-invest/investment-strategy/`.

Guasoni, P. and E. Mayerhofer (2023). Leveraged funds: Robust replication and performance evaluation. *Quantitative Finance* 23(7-8), 1155–1176.

Han, J. and E. Weinan (2016). Deep learning approximation for stochastic control problems. *NIPS Deep Reinforcement Learning Workshop* .

Hassine, M. and T. Roncalli (2013). Measuring performance of exchange-traded funds. *The Journal of Index Investing* 4(3), 57–85.

Hill, J. and G. Foster (2009). Understanding returns of leveraged and inverse funds. *Journal of Indexes* September/October, 40–58.

Hu, R. and M. Laurière (2023). Recent developments in machine learning methods for stochastic control and games. *Working paper* .

Israelsen, C. K. and G. F. Cogswell (2007). The error of tracking error. *Journal of Asset Management* 7(6), 419–424.

Jarrow, R. A. (2010). Understanding the risk of leveraged ETFs. *Finance Research Letters* 7(3), 135–139.

Kashyap, A. K., N. Kovrijnykh, J. Li, and A. Pavlova (2021). The benchmark inclusion subsidy. *Journal of Financial Economics* Forthcoming.

Knapp, C. (2023). Leveraged ETFs: The data says we're missing out. *LinkedIn* https://linkedin.com/ pulse/leveraged-etfs-data-says-were-missing-out-christian-knapp.

Korn, R. and C. Lindberg (2014). Portfolio optimization for an investor with a benchmark. *Decisions in Economics and Finance* 37, 373–384.

Kou, S. G. (2002). A jump-diffusion model for option pricing. *Management Science* 48(8), 1086–1101.

Lehalle, C.-A. and G. Simon (2021). Portfolio selection with active strategies: how long only constraints shape convictions. *Journal of Asset Management* .
URL `https://doi.org/10.1057/s41260-021-00219-z`

Leung, T., M. Lorig, and A. Pascucci (2017). Leveraged ETF implied volatilities from ETF dynamics. *Mathematical Finance* 27(4), 1035–1068.

Leung, T. and M. Santoli (2012). Leveraged etfs: admissible leverage and risk horizon. *Journal of Investment Strategies* 2(1), 39–61.

Leung, T. and R. Sircar (2015). Implied volatility of leveraged ETF options. *Applied Mathematical Finance* 22(2), 162–188.

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

Li, Y. and P. Forsyth (2019). A data-driven neural network approach to optimal asset allocation for target based defined contribution pension plans. *Insurance: Mathematics and Economics* 86, 189–204.

Mackintosh, P. (2008). Double trouble. *The Journal of Beta Investment Strategies* 1, 25–31.

Mäkinen, R. A. E. and J. Toivanen (2024). Short communication: Monte carlo expected wealth and risk measure trade-off portfolio optimization. *SIAM Journal on Financial Mathematics* 15(2), SC41–SC53.

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

Merton, R. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3, 125–144.

Ni, C., Y. Li, P. Forsyth, and R. Carroll (2022). Optimal asset allocation for outperforming a stochastic benchmark target. *Quantitative Finance* 22:9, 1595–1626.

Ni, C., Y. Li, and P. A. Forsyth (2024). Neural network approach to portfolio optimization with leverage constraints: a case study on high inflation investment. *Quantitative Finance* 24(6), 753–777.

Oksendal, B. and A. Sulem (2019). *Applied Stochastic Control of Jump Diffusions*. Springer, 3rd edition.

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

Politis, D., J. Romano, and M. Wolf (1999). *Subsampling*. Springer-Verlag. Chapter 4: Subsampling for nonstationary time series.

Politis, D. N. (2003). The impact of bootstrap methods on time series analysis. *Statistical Science* 18:2, 219–230.

Reppen, A. M. and H. M. Soner (2023). Deep empirical risk minimization in finance: looking into the future. *Mathematical Finance* 33(1), 116–145.

Reppen, A. M., H. M. Soner, and V. Tissot-Daguette (2023). Deep stochastic optimization in finance. *Digital Finance* 5, 91–111.

Scott, L. and S. Cavaglia (2017). A wealth management perspective on factor premia and the value of downside protection. *The Journal of Portfolio Management* 43:3, 1–9.

Simonian, J. and A. Martirosyan (2022). Sharpe parity redux. *The Journal of Portfolio Management* 48:9, 183–193.

Strub, M., D. Li, and X. Cui (2019). An enhanced mean-variance framework for robo-advising applications SSRN 302111.

Sullivan, R. N. (2009). The pitfalls of leveraged and inverse ETFs. *CFA Magazine* 20(3).

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

Van Staden, P. M., D. Dang, and P. Forsyth (2021). On the distribution of terminal wealth under dynamic mean-variance optimal investment strategies. *SIAM Journal on Financial Mathematics* 12(2), 566–603.

Van Staden, P. M., P. A. Forsyth, and Y. Li (2023). Beating a benchmark: dynamic programming may not be the right numerical approach. *SIAM Journal on Financial Mathematics* 14(2).

Van Staden, P. M., P. A. Forsyth, and Y. Li (2024). A global-in-time neural network approach to dynamic portfolio optimization. *Applied Mathematical Finance* (online), 1–33.
URL https://doi.org/10.1080/1350486X.2024.2410200

Vigna, E. (2014). On efficiency of mean-variance based portfolio selection in defined contribution pension schemes. *Quantitative Finance* 14(2), 237–258.

Vigna, E. (2020). On time consistency for mean-variance portfolio selection. *International Journal of Theoretical and Applied Finance* 23(6).

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

Wagalath, L. (2014). Modelling the rebalancing slippage of leveraged exchange-traded funds. *Quantitative Finance* 14(9), 1503–1511.

Yoon, J., D. Jarrett, and M. Van der Schaar (2019). Time-series generative adversarial networks. *33rd Conference on Neural Information Processing Systems (NeurIPS 2019)* https://proceedings.neurips.cc/paper/2019/file/c9efe5f26cd17ba6216bbe2a7d26d490-Paper.pdf.

Zhao, Y. (2007). A dynamic model of active portfolio management with benchmark orientation. *Journal of Banking and Finance* 31, 3336–3356.

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

Appendix A: Proofs of main results

The proofs of the key results of Section 3 are presented in this appendix.

A.1: Proof of Theorem 3.1

For a given deterministic benchmark strategy $\hat{\varrho}_s(t)$, consider an arbitrary admissible investor strategy $\varrho_\ell(t) := \varrho_\ell(t, \boldsymbol{X}_\ell(t)) \in \mathcal{A}_0$, where we omit the dependence of ϱ_ℓ on $\boldsymbol{X}_\ell(t) = \left(W_\ell(t), \hat{W}(t), \hat{\varrho}_s(t)\right)$ for notational simplicity. Considering the objective functional of the IR problem (3.26) at a given point $(t, w, \hat{w}) \in [t_0, T] \times \mathbb{R}^2$ for a given and fixed value of $\gamma > 0$, define

$$J(t, w, \hat{w}; \varrho_\ell) \quad = \quad E_{\varrho_\ell}^{t,w,\hat{w}}\left[\left(W_\ell(T) - \left[\hat{W}(T) + \gamma\right]\right)^2\right]. \tag{A.1}$$

If we proceed informally and assume that J sufficiently smooth, then the application of Itô's lemma for jump processes (Oksendal and Sulem (2019)) gives

$$
\begin{aligned}
& E_{\varrho_\ell}^{t,w,\hat{w}}\left[\int_t^{t+h} dJ\left(u, W_\ell(u), \hat{W}(u); \varrho_\ell\right)\right] \\
= \; & E_{\varrho_\ell}^{t,w,\hat{w}}\left[\int_t^{t+h}\left(\frac{\partial J}{\partial t}+\frac{\partial J}{\partial w}\cdot\left\{W_\ell(u)\cdot\left[r+\varrho_\ell(u)\left\{\beta\left(\mu-\lambda\kappa_1^s-r\right)-c_\ell\right\}\right]+q\right\}\right)\cdot du\right] \\
& +E_{\varrho_\ell}^{t,w,\hat{w}}\left[\int_t^{t+h}\frac{\partial J}{\partial \hat{w}}\cdot\left\{\hat{W}(u)\cdot\left[r+\left(\mu-\lambda\kappa_1^s-r\right)\hat{\varrho}_s(u)\right]+q\right\}\cdot du\right] \\
& +E_{\varrho_\ell}^{t,w,\hat{w}}\left[\int_t^{t+h}\frac{1}{2}\left(\frac{\partial^2 J}{\partial \hat{w}^2}\cdot\left[\hat{\varrho}_s(u)\hat{W}(u)\sigma\right]^2+\frac{\partial^2 J}{\partial w^2}\cdot\left[\varrho_\ell(u)W_\ell(u)\beta\sigma\right]^2\right)\cdot du\right] \\
& +E_{\varrho_\ell}^{t,w,\hat{w}}\left[\int_t^{t+h}\frac{\partial^2 J}{\partial w\partial \hat{w}}\cdot\left[\varrho_\ell(u)W_\ell(u)\beta\sigma\right]\left[\hat{\varrho}_s(u)\hat{W}(u)\sigma\right]\cdot du\right] \\
& +E_{\varrho_\ell}^{t,w,\hat{w}}\left[\lambda\int_t^{t+h}\left[\int_0^\infty \phi\left(u, W_\ell(u^-), \hat{W}(u^-), \xi^s; \varrho_\ell\right)G(\xi^s)\,d\xi^s - J\left(W_\ell(u^-), \hat{W}(u^-), u; \varrho_\ell\right)\right]du\right] \qquad \text{(A.2)}
\end{aligned}
$$

where all partial derivatives are evaluated at $\left(u, W_\ell(u), \hat{W}(u); \varrho_\ell\right)$, and

$$
\begin{aligned}
& \phi\left(u, W_\ell(u^-), \hat{W}(u^-), \xi^s; \varrho_\ell\right) \\
= \; & J\left(W_\ell(u^-)+\varrho_\ell(u)W_\ell(u^-)\beta\left(\xi^\ell-1\right), \hat{W}(u^-)+\hat{\varrho}_s(u)\hat{W}(u^-)\left(\xi^s-1\right), u; \varrho_\ell\right). \qquad \text{(A.3)}
\end{aligned}
$$

Recall that the LETF jump multiplier ξ^ℓ is a function (3.6) of the underlying index S jump multiplier ξ^s, so ϕ in (A.4) can be interpreted as a function of ξ^s if all other values are held fixed.

Continuing to proceed informally, dividing (A.2) by $h > 0$, taking limits as $h \downarrow 0$ and assuming the limits and expectations could be interchanged, an application of the dynamic programming principle results in the PIDE (3.27)-(3.28).

While providing the necessary intuition, the preceding arguments are merely informal. However, since similar arguments (see Applebaum (2004); Oksendal and Sulem (2019)) can be applied to a suitably smooth test function instead of the objective functional in order to formally prove (3.27)-(3.28), the details are therefore omitted.

A.2: Proof of Proposition 3.2

The quadratic terminal condition (3.28) suggests an ansatz for the value function V of the form

$$
V(t, w, \hat{w}, \hat{\varrho}_s) \;=\; A(t)w^2+\hat{A}(t)\hat{w}^2+D(t)w\hat{w}+F(t)w+\hat{F}(t)\hat{w}+C(t), \qquad \text{(A.4)}
$$

where $A, \hat{A}, D, F, \hat{F}$ and C are deterministic but unknown functions of time. Since (A.4) implies partial derivatives of the form

$$
\frac{\partial V}{\partial w}=2A(t)w+F(t)+D(t)\hat{w}, \qquad \frac{\partial^2 V}{\partial w^2}=2A(t), \qquad \text{and} \qquad \frac{\partial^2 V}{\partial w\partial \hat{w}}=D(t), \qquad \text{(A.5)}
$$

substituting (A.4)-(A.5) into the HJB PIDE (3.27) results in the pointwise supremum $\varrho_\ell^* = \varrho_\ell^*(t, w, \hat{w}, \hat{\varrho}_s)$ obtained from the first-order condition that satisfies the relationship

$$
\begin{aligned}
\varrho_\ell^*\cdot w \;=\; & -\left(\frac{\beta\left[\mu+\lambda\left(\kappa_1^\ell-\kappa_1^s\right)-r\right]-c_\ell}{\beta^2\left(\sigma^2+\lambda\kappa_2^\ell\right)}\right)\left[w+\frac{F(t)}{2A(t)}+\frac{D(t)}{2A(t)}\cdot\hat{w}\right] \\
& -\left[\frac{\sigma^2+\lambda\kappa_\chi^{\ell,s}}{\beta\left(\sigma^2+\lambda\kappa_2^\ell\right)}\right]\frac{D(t)}{2A(t)}\cdot\hat{\varrho}_s\hat{w}. \qquad \text{(A.6)}
\end{aligned}
$$

Substituting (A.6) into (3.29) to obtain $\mathcal{H}\left(\varrho_{\ell}^{*};t,w,\hat{w},\hat{\varrho}_{s}\right)$, the PIDE (3.27)-(3.28) implies the following set of ordinary differential equations (ODEs) for the unknown functions A, D and F on $t \in [t_0, T]$,

$$\frac{d}{dt}A\left(t\right) = -\left(2r - \frac{\left(\beta\left[\mu+\lambda\left(\kappa_1^{\ell}-\kappa_1^{s}\right)-r\right]-c_{\ell}\right)^2}{\beta^2\left(\sigma^2+\lambda\kappa_2^{\ell}\right)}\right)A\left(t\right), \qquad A\left(T\right)=1, \tag{A.7}$$

$$\frac{d}{dt}D\left(t\right) = -\left(2r - \frac{\left(\beta\left[\mu+\lambda\left(\kappa_1^{\ell}-\kappa_1^{s}\right)-r\right]-c_{\ell}\right)^2}{\beta^2\left(\sigma^2+\lambda\kappa_2^{\ell}\right)} + K_{\beta}^{\ell,s}\cdot\hat{\varrho}_s\left(t\right)\right)D\left(t\right), \qquad D\left(T\right)=-2, \tag{A.8}$$

$$\frac{d}{dt}F\left(t\right) = -2qA\left(t\right) - \left(r - \frac{\left(\beta\left[\mu+\lambda\left(\kappa_1^{\ell}-\kappa_1^{s}\right)-r\right]-c_{\ell}\right)^2}{\beta^2\left(\sigma^2+\lambda\kappa_2^{\ell}\right)}\right)F\left(t\right) - qD\left(t\right), \qquad F\left(T\right)=-2\gamma, \tag{A.9}$$

where the constant $K_{\beta}^{\ell,s}$ is given by (3.33)

Note that the derivation of (A.8) as an ODE requires the benchmark strategy $\hat{\varrho}_s$ to be deterministic (in the case of closed-form solutions) as per Assumption 3.2. Solving ODEs (A.7)-(A.9), we obtain

$$A\left(t\right) = \exp\left\{\left(2r - \frac{\left(\beta\left[\mu+\lambda\left(\kappa_1^{\ell}-\kappa_1^{s}\right)-r\right]-c_{\ell}\right)^2}{\beta^2\left(\sigma^2+\lambda\kappa_2^{\ell}\right)}\right)\left(T-t\right)\right\}, \tag{A.10}$$

$$D\left(t\right) = -2\cdot\exp\left\{\left(2r - \frac{\left(\beta\left[\mu+\lambda\left(\kappa_1^{\ell}-\kappa_1^{s}\right)-r\right]-c_{\ell}\right)^2}{\beta^2\left(\sigma^2+\lambda\kappa_2^{\ell}\right)}\right)\left(T-t\right) + K_{\beta}^{\ell,s}\cdot\int_t^T\hat{\varrho}_s\left(u\right)du\right\}, \tag{A.11}$$

$$\begin{aligned} F\left(t\right) &= 2\exp\left\{\left(r - \frac{\left(\beta\left[\mu+\lambda\left(\kappa_1^{\ell}-\kappa_1^{s}\right)-r\right]-c_{\ell}\right)^2}{\beta^2\left(\sigma^2+\lambda\kappa_2^{\ell}\right)}\right)\left(T-t\right)\right\}\times \\ &\quad \left[-\gamma + \frac{q}{r}\left(e^{r(T-t)}-1\right) - q\cdot\int_t^T\exp\left\{r\left(T-v\right)+K_{\beta}^{\ell,s}\cdot\int_v^T\hat{\varrho}_s\left(u\right)du\right\}dv\right]. \end{aligned} \tag{A.12}$$

Substituting (A.10)-(A.12) into (A.6) and simplifying, we obtain the optimal fraction of wealth ϱ_{ℓ}^{*} to invest in the LETF (3.30) as per Proposition 3.2.

A.3: Expressions for κ_1^{ℓ}, κ_2^{ℓ} and $\kappa_{\chi}^{\ell,s}$

For the purposes of illustrating the closed-form solutions of Section 3, the broad equity market index S is assumed to have dynamics (3.2) with jumps as modelled in Kou (2002). As a result, with p_{up} denoting the probability of an upward jump given that a jump occurs, $y = \log\xi^s$ is assumed in Kou (2002) to follow an asymmetric double-exponential distribution with PDF $g\left(y\right)$,

$$g\left(y\right) = p_{up}\eta_1 e^{-\eta_1 y}\mathbb{I}_{\{y\geq 0\}} + \left(1-p_{up}\right)\eta_2 e^{\eta_2 y}\mathbb{I}_{\{y<0\}}, \tag{A.13}$$

where $p_{up} \in [0,1]$ and $\eta_1 > 1, \eta_2 > 0$. Equivalently, the PDF of ξ^s is given by

$$G\left(\xi^s\right) = p_{up}\eta_1\left(\xi^s\right)^{-\eta_1-1}\mathbb{I}_{[\xi^s\geq 1]}\left(\xi^s\right) + \left(1-p_{up}\right)\eta_2\left(\xi^s\right)^{\eta_2-1}\mathbb{I}_{[0\leq\xi^s<1]}\left(\xi^s\right). \tag{A.14}$$

Recall that we have defined κ_1^s and κ_1^s in (3.3), repeated here for convenience,

$$\kappa_1^s = \mathbb{E}\left[\xi^s-1\right], \qquad \kappa_2^s = \mathbb{E}\left[\left(\xi^s-1\right)^2\right]. \tag{A.15}$$

From the results in Kou (2002), we can obtain (A.15) for the distribution (A.14) using the results

$$\mathbb{E}\left[\xi^s\right] = \frac{p_{up}\eta_1}{\eta_1-1} + \frac{\left(1-p_{up}\right)\eta_2}{\eta_2+1}, \qquad \mathbb{E}\left[\left(\xi^s\right)^2\right] = \frac{p_{up}\eta_1}{\eta_1-2} + \frac{\left(1-p_{up}\right)\eta_2}{\eta_2+2}. \tag{A.16}$$

However, since the LETF experiences slightly different jumps as per (3.6), which we repeat here for convenience,

$$\xi^{\ell} = \begin{cases} \xi^s & \text{if } \xi^s > \left(\beta-1\right)/\beta, \\ \dfrac{\left(\beta-1\right)}{\beta} & \text{if } \xi^s \leq \left(\beta-1\right)/\beta, \end{cases} \tag{A.17}$$

we cannot use the results (A.15) directly. Instead, expressions for κ_1^ℓ, κ_2^ℓ and $\kappa_\chi^{\ell,s}$ are obtained using the results the following lemma.

Lemma A.1. *(κ_1^ℓ, κ_2^ℓ and $\kappa_\chi^{\ell,s}$ in the Kou (2002) model) Suppose the jump multiplier ξ^s in the S-dynamics (3.2) has PDF $G(\xi^s)$ given by (A.14), and a LETF with returns multiplier $\beta > 1$ and dynamics (3.8) has jump multiplier ξ^ℓ, which is defined in terms of ξ^s as per (A.17). Then the quantities*

$$\kappa_1^\ell = \mathbb{E}\left[\xi^\ell - 1\right], \qquad \kappa_2^\ell = \mathbb{E}\left[\left(\xi^\ell - 1\right)^2\right], \qquad \kappa_\chi^{\ell,s} = \mathbb{E}\left[\left(\xi^\ell - 1\right)\left(\xi^s - 1\right)\right], \tag{A.18}$$

required by the IR-optimal investment strategy in Proposition 3.2 can be obtained using the following expressions:

$$\mathbb{E}\left[\xi^\ell\right] = \frac{p_{up}\eta_1}{\eta_1 - 1} + (1 - p_{up})\,\eta_2 \cdot \left[\frac{\vartheta^{\eta_2+1}}{\eta_2} + \left(\frac{1 - \vartheta^{\eta_2+1}}{\eta_2 + 1}\right)\right], \tag{A.19}$$

$$\mathbb{E}\left[\left(\xi^\ell\right)^2\right] = \frac{p_{up}\eta_1}{\eta_1 - 2} + (1 - p_{up})\,\eta_2 \cdot \left[\frac{\vartheta^{\eta_2+2}}{\eta_2} + \left(\frac{1 - \vartheta^{\eta_2+2}}{\eta_2 + 2}\right)\right], \tag{A.20}$$

$$\mathbb{E}\left[\xi^\ell \xi^s\right] = \frac{p_{up}\eta_1}{\eta_1 - 2} + (1 - p_{up})\,\eta_2 \cdot \left[\frac{\vartheta^{\eta_2+2}}{\eta_2 + 1} + \left(\frac{1 - \vartheta^{\eta_2+2}}{\eta_2 + 2}\right)\right], \tag{A.21}$$

where $\vartheta = (\beta - 1)/\beta$.

Proof. Consider (A.21). Since $\beta > 1$ and $\vartheta = (\beta - 1)/\beta$, we have $0 < \vartheta < 1$. Therefore, using the definition of the LETF jump multiplier (A.17) and the PDF $G(\xi^s)$, we have

$$\begin{aligned}
\mathbb{E}\left[\xi^\ell \xi^s\right] &= \int_0^\infty \xi^\ell \xi^s \cdot G(\xi^s)\, d\xi^s \\
&= \vartheta \cdot \int_0^\vartheta \xi^s \cdot G(\xi^s)\, d\xi^s + \int_\vartheta^1 (\xi^s)^2 \cdot G(\xi^s)\, d\xi^s + \int_1^\infty (\xi^s)^2 \cdot G(\xi^s)\, d\xi^s.
\end{aligned} \tag{A.22}$$

Standard results (see (A.16) and Kou (2002)) for the Kou model gives

$$\int_1^\infty (\xi^s)^2 \cdot G(\xi^s)\, d\xi^s = \frac{p_{up}\eta_1}{\eta_1 - 2}. \tag{A.23}$$

Using (A.23) and writing the first two terms of (A.22) in terms of the log jump multiplier $y = \log \xi^s$ with PDF $g(y)$ as per (A.13), we have

$$\begin{aligned}
\mathbb{E}\left[\xi^\ell \xi^s\right] &= \vartheta \cdot \int_{-\infty}^{\log\vartheta} e^y g(y)\, dy + \int_{\log\vartheta}^0 e^{2y} g(y)\, dy + \frac{p_{up}\eta_1}{\eta_1 - 2} \\
&= (1 - p_{up})\,\eta_2 \vartheta \cdot \int_{-\infty}^{\log\vartheta} e^{(\eta_2+1)y} dy + (1 - p_{up})\,\eta_2 \int_{\log\vartheta}^0 e^{(\eta_2+2)y} dy + \frac{p_{up}\eta_1}{\eta_1 - 2}.
\end{aligned} \tag{A.24}$$

Simplifying (A.24) gives (A.21). Since (A.19) and (A.20) can be obtained using similar arguments, the details are omitted. □

A.4: Proof of Corollary 3.3

For purposes of intuition, we first give informal arguments as to how the results of Corollary 3.3 relate to the results of Proposition 3.2. Recall that the VETF has returns multiplier $\beta = 1$ (i.e. the VETF simply aims to replicate the returns of S before costs) and expense ratio $c_v > 0$. Note that if we let $\beta \downarrow 1$ in (3.6), we have

$$\lim_{\beta \downarrow 1} \xi^\ell = \xi^s \qquad \text{a.s.}, \tag{A.25}$$

from which it follows that

$$\lim_{\beta \downarrow 1} \kappa_1^\ell = \kappa_1^s, \quad \lim_{\beta \downarrow 1} \kappa_2^\ell = \kappa_2^s, \quad \lim_{\beta \downarrow 1} \kappa_\chi^{\ell,s} = \kappa_2^s. \tag{A.26}$$

Therefore, comparing the VETF and LETF investor wealth dynamics (3.23)-(3.24) in the case of identical expense ratios (i.e. $c_\ell = c_v$), identical but not necessarily optimal investment strategies ($\varrho_\ell = \varrho_v$) and the identical initial wealth, we have

$$\lim_{\beta\downarrow 1} W_\ell(t) \quad = \quad W_v(t) \quad \text{a.s. } \forall t \in [t_0, T], \quad \text{if } W_\ell(t_0) = W_v(t_0), c_\ell = c_v, \quad \text{and } \varrho_\ell = \varrho_v. \tag{A.27}$$

In other words, if we let $\beta \downarrow 1$ in the LETF investor wealth dynamics (3.24), we recover the VETF investor wealth dynamics (3.23). Continuing to proceed informally, the results of Corollary 3.3 can therefore be obtained by letting $\beta \downarrow 1$ in the results of Proposition 3.2, provided we use the VETF expense ratio c_v in both (3.24) and (3.23). Note that the definition (3.33) of $K_\beta^{\ell,s}$, results (A.26) and setting $c_\ell = c_v$ imply that

$$\begin{aligned} \lim_{\beta\downarrow 1} K_\beta^{\ell,s} &= \lim_{\beta\downarrow 1} \left[\mu - r - \frac{\left(\beta\left[\mu + \lambda\left(\kappa_1^\ell - \kappa_1^s\right) - r\right] - c_v\right)\left(\sigma^2 + \lambda\kappa_\chi^{\ell,s}\right)}{\beta\left(\sigma^2 + \lambda\kappa_2^\ell\right)} \right] \\ &= c_v, \end{aligned} \tag{A.28}$$

which confirms that the functions g_v and h_v ((3.35)-(3.36)) can be obtained from the functions the functions g_ℓ and h_ℓ ((3.31)-(3.32)) if identical expense ratios are used.

The preceding discussions were merely informal. More formally, the proof of Corollary 3.3 proceeds along the same lines as the proof of Proposition 3.2, except that VETF investor wealth dynamics (3.23) is used instead of (3.24), and details are therefore omitted.

A.5: Proof of Proposition 3.4

Suppose we have zero expense ratios, i.e. $c_v = c_\ell = 0$, and there are no jumps in the underlying S-dynamics (i.e. $\lambda = 0$ in (3.2)). Substituting these values in the deterministic functions g_ℓ and h_ℓ ((3.31) and (3.32)) in the case of a LETF and the deterministic functions g_v and h_v ((3.35) and (3.36)) in the case of the VETF, we have

$$g_\ell(t) = g_v(t) = 1, \tag{A.29}$$

and

$$h_\ell(t) = h_v(t) = 0. \tag{A.30}$$

In the case of the LETF investor, the optimal control (3.30) now satisfies

$$\beta \cdot W_\ell^*(t) \cdot \varrho_\ell^*(t, \boldsymbol{X}_\ell^*(t)) \quad = \quad \left(\frac{\mu - r}{\sigma^2}\right) \cdot \left[\gamma e^{-r(T-t)} - \left(W_\ell^*(t) - \hat{W}(t)\right)\right] + \hat{\varrho}_s(t)\hat{W}(t), \tag{A.31}$$

whereas in the case of the VETF investor, the optimal control (3.30) now becomes

$$W_v^*(t) \cdot \varrho_v^*(t, \boldsymbol{X}_v^*(t)) \quad = \quad \left(\frac{\mu - r}{\sigma^2}\right) \cdot \left[\gamma e^{-r(T-t)} - \left(W_v^*(t) - \hat{W}(t)\right)\right] + \hat{\varrho}_s(t)\hat{W}(t). \tag{A.32}$$

Using (A.31) and (A.32), define the auxiliary process $Q(t)$ as

$$Q(t) \quad = \quad e^{-rt} \cdot \left[W_\ell^*(t) - W_v^*(t)\right], \qquad t \in [t_0 = 0, T], \tag{A.33}$$

with $Q(t_0) = e^{-rt_0}\left[W_\ell^*(t_0) - W_v^*(t_0)\right] = w_0 - w_0 = 0$.

Substituting the optimal controls in this special case ((A.31) and (A.32)) into the wealth dynamics (3.23)-(3.24) and recalling that there are no jumps, we obtain the dynamics

$$\frac{dQ(t)}{Q(t)} = \left(\frac{\mu - r}{\sigma}\right)^2 \cdot dt - \left(\frac{\mu - r}{\sigma}\right) \cdot dZ(t). \tag{A.34}$$

Since $Q(t_0) = 0$, dynamics (A.34) therefore imply that $Q(t) = 0, \forall t \geq t_0$, so that in the special case of zero costs, we have

$$W_\ell^*(t) \quad = \quad W_v^*(t), \qquad \forall t \in [t_0, T], \tag{A.35}$$

which confirms (3.39).

Using (A.35) to write $W^*(t) \coloneqq W_\ell^*(t) = W_v^*(t)$ in this special case, the difference in controls (A.31) and

(A.32) satisfy

$$\begin{aligned}\left[\beta \cdot \varrho_{\ell}^{*}\left(t, \boldsymbol{X}_{\ell}^{*}(t)\right)-\varrho_{v}^{*}\left(t, \boldsymbol{X}_{v}^{*}(t)\right)\right] \cdot W^{*}(t) = & \beta \cdot W_{\ell}^{*}(t) \cdot \varrho_{\ell}^{*}\left(t, \boldsymbol{X}_{\ell}^{*}(t)\right)-W_{v}^{*}(t) \cdot \varrho_{v}^{*}\left(t, \boldsymbol{X}_{v}^{*}(t)\right) \\ = & -\left(\frac{\mu-r}{\sigma^{2}}\right)\left[W_{\ell}^{*}(t)-W_{v}^{*}(t)\right] \\ = & 0, \qquad \forall t \geq t_{0},\end{aligned} \tag{A.36}$$

thereby verifying (3.38).

Finally, given the form of the optimal control (A.32) for the VETF in this special case, together with Assumption 3.1 and wealth dynamics (3.23), imply that we can obtain the optimal Information Ratio in the case of the VETF as given by (3.40) (see Van Staden et al. (2023)). However, since (A.35), this is also the optimal IR using the LET in this special case, thereby confirming (3.40) and completing the proof of Proposition 3.4.

Appendix B: Source data and parameters

In this appendix, we provide details regarding the source data used to obtain the indicative investment results presented in Section 3 and Section 5.

Returns data for US Treasury bills and bonds, as well as the broad equity market index, were obtained from the CRSP[8]. In more detail, the historical time series are as follows:

(i) T30 (30-day Treasury bill): CRSP, monthly returns for 30-day Treasury bill.

(ii) B10 (10-year Treasury bond): CRSP, monthly returns for 10-year Treasury bond.

(iii) Market (broad equity market index): CRSP, monthly and daily returns, including dividends and distributions, for a capitalization-weighted index consisting of all domestic stocks trading on major US exchanges (the VWD index).

CRSP data was obtained for the historical time period 1926:01 to 2023:12. All time series were inflation-adjusted using inflation data from the US Bureau of Labor Statistics[9].

B.1: Constructing VETF and LETF returns time series

LETFs were only introduced in 2006 (Bansal and Marshall (2015)), whereas the first VETFs were listed in the US in the 1990s. In order to obtain longer time series of returns for indicative investment results of Section 5, we construct a proxy returns time series for a VETF and LETF referencing a broad equity market index as follows:

(i) Obtain daily returns for the underlying broad equity market index referenced by the VETF and LETF. For this purpose, we used daily returns for the CRSP capitalization-weighted index consisting of all domestic stocks trading on major US exchanges (the VWD index - see above), with historical data that is available since January 1926. We prefer to use a time series that is as long as possible, since this would include additional periods of exceptional market volatility such as 1929-1933.

(ii) Multiply each daily return by the returns multiplier β, where we used $\beta = 2$ for the LETF and $\beta = 1$ for the VETF, and construct a time series of monthly returns.

(iii) Adjust the time series of VWD returns using (3.5) to reflect the ETF expense ratios $c_k, k \in \{v, \ell\}$ and the observed T-bill rate r. As per Table 3.1, we assumed expense ratios of $c_\ell = 0.89\%$ p.a. for the LETF and $c_v = 0.06\%$ p.a. for the VETF to reflect typical values observed in the market.

(iv) Inflation-adjust the time series using inflation data from the US Bureau of Labor Statistics.

[8] Calculations were based on data from the Historical Indexes 2024©, 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.

[9] The annual average CPI-U index, which is based on inflation data for urban consumers, were used - see `http://www.bls.gov.cpi`

Note that a proxy time series of LETF returns is also similarly constructed in Bansal and Marshall (2015), although a number of details (such as inflation adjustment of returns and choice of underlying index) differ. As noted in Bansal and Marshall (2015), the construction of such a proxy returns time series assumes that the ETF managers achieve a negligible tracking error with respect to the underlying index. We observe that these assumptions are often made out of necessity in the literature concerning LETFs (e.g. Bansal and Marshall (2015), Leung and Sircar (2015)). Furthermore, given improvements in designing replication strategies for LETFs that remain robust even during periods of market volatility (see for example Guasoni and Mayerhofer (2023)), this appears to be a reasonable assumption for ETFs written on major stock market indices as considered in this paper.

We emphasize that the proxy time series for VETF and LETF returns are only used for bootstrapping the data sets for the numerical solutions implementing the data-driven neural network approach (Section 4 and Section 5), and *not* for the closed-form solutions of Section 3. This follows since closed-form solutions in Section 3 assume parametric dynamics for the underlying assets including the broad equity market index, from which the LETF and VETF dynamics can be constructed using (3.8) and (3.4), respectively.

B.2: Calibrated parameters for closed-form solutions

For the closed-form solutions of Section 3, using the CRSP data for 30-day T-bills and the broad equity market index (VWD index) for the period 1926:01 to 2023:12 as outlined above, the filtering technique as per Dang and Forsyth (2016); Forsyth and Vetzal (2017) for calibrating inflation-adjusted Kou (2002) jump-diffusion processes resulted in the calibrated process parameters as presented in Table B.1. Given the specified dynamics (3.1)-(3.2) of the risk-free asset B and equity market index S (with parameters as in Table B.1), we can obtain the dynamics of the LETF (3.8) and VETF (3.4).

Table B.1: Closed-form solutions: Calibrated, inflation-adjusted parameters for asset dynamics (3.1) and (3.2), assuming the Kou (2002) jump-diffusion model with $G(\xi^s)$ given by (A.14). The calibration methodology of Dang and Forsyth (2016); Forsyth and Vetzal (2017) is used with a jump threshold parameter value of 3.

Assumption for S-dynamics	Calibrated parameters						
	r	μ	σ	λ	p_{up}	η_1	η_2
Jump-diffusion (Kou model)	0.0031	0.0873	0.1477	0.3163	0.2258	4.3591	5.5337
GBM dynamics (no jumps)	0.0031	0.0819	0.1850	-	-	-	-

Note that the values of the remaining parameters for the parametric dynamics can be calculated by substituting the values of Table B.1 into the results of Appendix A.3. This gives $\kappa_1^s = -0.0513$, $\kappa_2^s = 0.0884$, $\kappa_1^\ell = -0.0500$, $\kappa_2^\ell = 0.0870$ and $\kappa_\chi^{\ell,s} = 0.0876$.

Appendix C: Additional numerical results

Additional leverage and borrowing cost scenarios are analyzed as a supplement to the results of Section 5. We consider scenarios where the maximum leverage allowed decreases from $p_{max} = 1.5$ to $p_{max} = 1.2$, or increases to $p_{max} = 2.0$, and where zero borrowing costs might be applicable (as opposed to borrowing costs of $b = 0.03$ throughout Section 5).

As discussed in Section 5, we observe that the IR-optimal portfolio of the LETF investor achieves partial stochastic dominance over the corresponding IR-optimal portfolio of the VETF investor with the same outperformance target γ, even if the VETF investment can be leveraged and borrowing costs on short-selling decrease to zero (see for example Figure C.1). For some intuition as to the underlying explanation, please refer to Section 5 and in particular Figure 5.2 and the associated discussion.

In summary, the results of this appendix confirm that the conclusions of Section 5 do not appear to be sensitive in a qualitative sense to the specific maximum leverage or borrowing costs parameters used in the numerical analysis, provided these values are within a reasonable range.

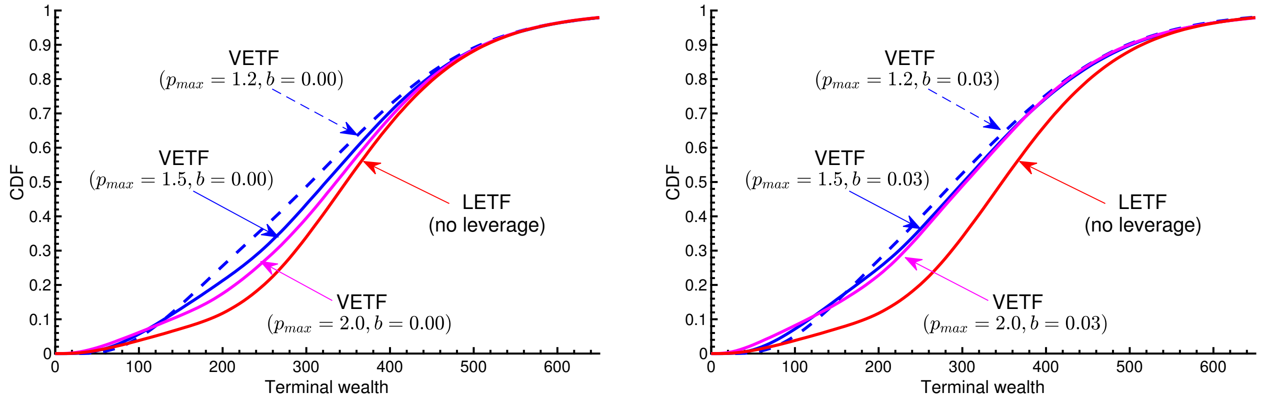

(a) CDF of $W(T)$ - Zero borrowing costs **(b)** CDF of $W(T)$ - With borrowing costs

Figure C.1: Effect of leverage and borrowing cost assumptions on CDFs of IR-optimal terminal wealth $W_k^*(T), k \in \{v, \ell\}$.

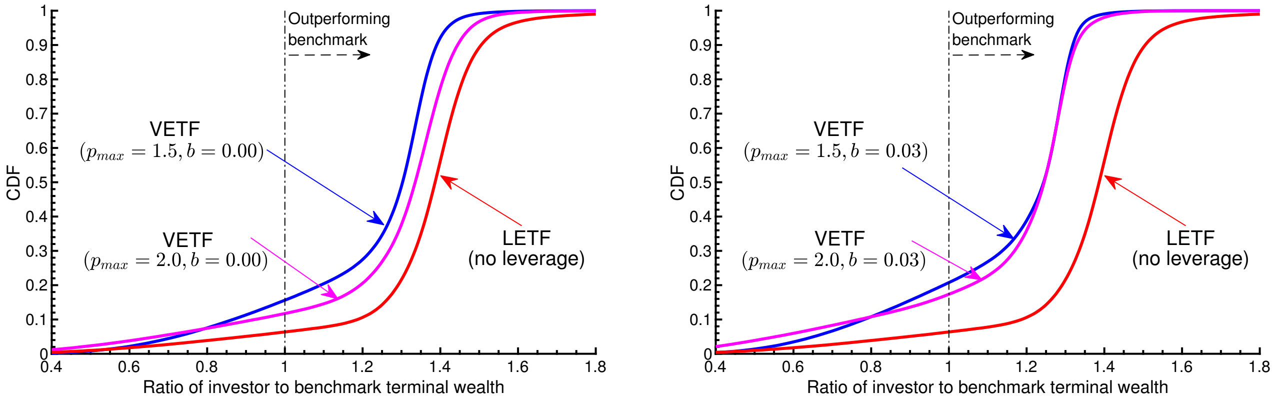

(a) CDF of ratio $W^*(T)/\hat{W}(T)$ - Zero borrowing costs **(b)** CDF of ratio $W(T)/\hat{W}(T)$ - With borrowing costs

Figure C.2: Effect of leverage and borrowing cost assumptions on the CDFs of IR-optimal terminal wealth ratios $W_k^*(T)/\hat{W}(T), k \in \{v, \ell\}$.

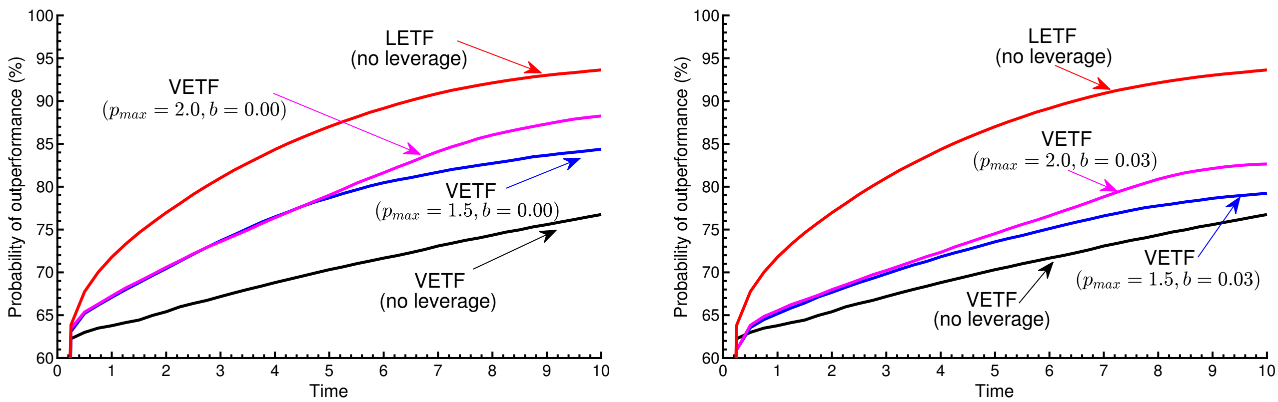

(a) Zero costs **(b)** With costs

Figure C.3: Effect of leverage and borrowing cost assumptions on the probability $W_k^*(t) > \hat{W}(t), k \in \{v, \ell\}$ of benchmark outperformance as a function of time t.