Vol. 00, No. 0, Xxxxx 0000, pp. 000–000
ISSN 0000-0000 | EISSN 0000-0000 | 00 | 0000 | 0001

INFORMS
DOI 10.1287/xxxx.0000.0000

Leveraged ETFs with Market Closure and Frictions

Min Dai
Department of Applied Mathematics, The Hong Kong Polytechnic University, Hung Hom, Kowloon, Hong Kong

Steven Kou
Questrom School of Business, Boston University, 595 Commonwealth Avenue, Boston, MA 02215, USA

Halil Mete Soner
Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08540, USA

Chen Yang
Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, Shatin, N.T., Hong Kong

Although leveraged ETFs are popular products for retail investors, how to hedge them poses a great challenge to financial institutions. We develop an optimal rebalancing (hedging) model for leveraged ETFs in a comprehensive setting, including overnight market closure and market frictions. The model allows for an analytical optimal rebalancing strategy. The result extends the principle of "aiming in front of target" introduced by Gârleanu and Pedersen (2013) from a constant weight between current and future positions to a time-varying weight because the rebalancing performance is monitored only at discrete time points, but the rebalancing takes place continuously. Empirical findings and implications for the weekend effect and the intraday trading volume are also presented.

Key words: daily rebalancing, leveraged ETFs, market closure, frictions.

1. Introduction

Leveraged ETFs (LETFs for short hereafter) are exchange-traded funds that aim to achieve a daily target return equal to a stated multiple of the daily return of an underlying asset, where the multiple β can be $+2$, $+3$ (for Bull LETFs), or -1, -2, -3 (for Bear LETFs), fixed for each fund. In 2006 ProShares introduced the first 12 LETFs on 4 U.S. equity indices (S&P 500, NASDAQ 100, Dow Jones Industrial Average, and S&P MidCap 400), each with three multiples $+2, -1, -2$. Since then, they have become popular, especially for retail investors. Indeed, LETFs provide investors with leveraged (both long and short) opportunities without directly accessing margin accounts or financial derivatives. As of March 2021, there are a total number of 108 LETFs on U.S. equity indices, with a total of 52.7 billion USD assets under management.

However, one important property of LETFs often misunderstood by retail investors is that the stated daily multiple cannot be translated directly to the exact multiple over a multi-day period

due to the discrete compounding effect; see, e.g., Avellaneda and Zhang (2010) and Jarrow (2010). For example, assume the underlying asset has a constant daily return of 0.1%, and the LETF hits the target multiple $\beta = 2$ every day. Then over 252 days, the LETF achieves a return of $(1 + 0.001 \times 2)^{252} - 1 = 65.45\%$, which is 2.285 (not 2) times of $(1 + 0.001)^{252} - 1 = 28.64\%$, the annual return of the underlying asset.

Even daily, the actual daily returns of the fund's net asset value tend to deviate from the target returns, a phenomenon called (daily) *slippage*. The LETF slippage has been documented in the empirical literature. Taking LETFs on the S&P 500 index as an example, Tang and Xu (2013) find that from 2006 to 2010, the LETFs with stated $\beta = 2, -1, -2$ achieved an average daily actual net asset value return of 0.0326%, −0.0097%, and −0.0247%, as compared to the average target return of 0.0471%, −0.0236%, −0.0471%, respectively.

The existing literature suggests several factors to explain the LETF slippage. Tang and Xu (2013) and Henderson and Buetow (2014) attribute the slippage to the interest cost; more precisely, to achieve a target leveraged exposure to the underlying asset, a Bull (resp. Bear) fund has to borrow (resp. lend) at the LIBOR rate, which brings down (resp. pushes up) the daily fund asset return and causes the slippage. Avellaneda and Dobi (2012), Wagalath (2014), and Guasoni and Mayerhofer (2017) instead ascribe to transaction costs incurred during the daily LETF rebalancing.

In this paper, we develop a model for the optimal daily rebalancing of LETFs under market frictions and overnight jump risk, which yields an explanation for the LETF daily slippage. More precisely, (i) an LETF fund usually rebalances near the market closing time, when it can better estimate the position needed for the target exposure for the next day; (ii) adjusting the fund position rapidly during a short period may incur a high cost due to the market price impact. Both can cause suboptimal exposure at the market closing, which leads to daily slippage. Additionally, in our model most intraday rebalancing activities occur right after market opening and before market closure, and the slippage tends to be the largest on Mondays, thus providing an explanation of the U-shape intraday trading volume and the "weekend effect" of the slippage of LETFs.

1.1. Intuition

Financial institutions face the challenge of hedging when issuing LETFs in the presence of market frictions and overnight market closure. To illustrate this, assume that the interest rate and transaction costs are zero and consider the following simple example. To reach the targeted daily return today, which is β times the daily return of the asset, an LETF fund has to achieve today's targeted position that corresponds to a leverage ratio of β at the previous market closure. Assume $\beta = 2$, and at the previous market closure, the underlying asset price and total asset are \$100 and \$100 million, respectively, and the fund has an ideal position of 2 million shares in the underlying

asset. Then by keeping this targeted position unchanged throughout today until market closure, the fund will reach the targeted return today. Assume the underlying asset value drops by 1% to \$99 at today's market closure, then the fund's total asset would drop by 2% to \$98 million and achieve the targeted return, which is ideal. However, the preparation for tomorrow's targeted position becomes a problem. To achieve the target, the fund would need to have tomorrow's position of $2 \times \$98$ million / $\$99 = 1.9798$ million shares at today's market closure by selling 0.0202 million shares right before today's market closure.

This selling is unrealistic and leads to daily slippage for at least two reasons: (1) In the presence of market frictions, selling a large number of shares right before the market closure is costly, as a sudden portfolio adjustment within an infinitesimal period would yield a high cost due to price impact and other transaction costs, leading to a suboptimal leverage ratio and potentially large slippage; (2) The slippage caused by a sub-optimal position at market closure can be amplified by the risk of overnight jumps in prices, as such jumps may move the fund's opening position further away and require larger rebalancing of positions.

In view of the two challenges, we show that the fund should start moving its aim continuously towards tomorrow's target much earlier than the market closure, in anticipation of tomorrow's target. In other words, the fund's aim is "in front of today's target". The proposed strategy can reduce the slippage significantly, as confirmed by our numerical results later.

1.2. Our Contribution and Literature Review

The contribution of this paper is fourfold. First, to the best of our knowledge, we are the first to study the daily rebalancing problem in a comprehensive setting, including market frictions and the overnight market closure. Empirical studies (e.g., Stoll and Whaley (1990), Lockwood and Linn (1990)) show that the volatility is much lower during market closure as compared to trading hours, although the expected return is on a similar level. Also, the trading strategy during trading hours needs to adjust for market closure when no trading is allowed (cf. Dai et al. (2015)). This is especially important for LETFs since if the target leverage ratio is not exactly met at market closing time, the fund bears the risk that an overnight price jump may lead to a large return deviation right at the next market opening.[1] We find that to reduce the overnight jump risk, it is optimal to perform a larger portion of rebalancing before market closure. Additionally, adjusting the aim for the overnight risk can lead to a slippage reduction of as much as 24%.

[1] A fund may continue rebalancing via trading derivatives such as futures after the normal trading hours of the stock market. However, during this period, the trading volume of futures is typically extremely low, making any large rebalancing impractical due to the potentially high price impact. Therefore, it is still optimal for the fund to finish the major, if not all, rebalancing before stock market closing time. Indeed, many researchers such as Cheng and Madhavan (2009), Bai et al. (2012), Tuzun (2014) pointed out that the LETFs carry out their rebalancing during the final hour before market closing.

Without the risk of market closure, several researchers focused on modeling the slippage of LETFs under price impact either in a pure discrete setting or a pure continuous setting. Wagalath (2014) proposed a model with an endogenous price impact and derived an analytic formula for the rebalancing slippage, assuming that rebalancing happens once per day and the return deviation is measured daily at the market closing time. Based on an approximate continuous model, Guasoni and Mayerhofer (2017) studied the fund's trade-off between the short-term and long-term deviation. With proportional transaction costs, they derived the optimal trading boundary, which has an explicit form in terms of an asymptotic expansion. In contrast, our model incorporates several other additional factors, including the risk of market closure due to overnight price jump.

To help unveil the optimal intraday rebalancing strategy, we use a delicate mixture of continuous and discrete-time modeling: The intraday trading strategies and market frictions are modeled in a continuous-time setting that allows gradual adjustment across the trading hours, while the performance, i.e., the return deviation, is monitored daily in a discrete-time manner. This mixture modeling is consistent with the LETFs in reality and complements the existing literature either in pure discrete-time or continuous-time settings. Furthermore, while the two papers mentioned above assume a linear price impact in the trading amount (i.e., proportional transaction costs), we consider a quadratic price impact in the trading amount.

Second, our rebalancing strategy contributes to the empirical study of LETF slippage. The slippage of LETFs was studied empirically in Tang and Xu (2013) and Henderson and Buetow (2014), both of which found that LETFs' realized return deviated significantly from the target return and explained the slippage via the interest cost incurred due to borrowing and lending. Several papers also linked the LETF slippage to the market friction costs. Using the 2008 financial crisis period, Shum and Kang (2013) observed a significant slippage for LETFs on international indices that tend to have lower liquidity. Avellaneda and Dobi (2012) found that the slippage tends to be prominent during volatile periods.

However, the majority of empirical literature studies the slippage by assuming the fund follows a simple suboptimal rebalancing strategy. For example, Tang and Xu (2013) and Henderson and Buetow (2014) assume that the fund achieves the target leverage ratio exactly at market closing and keeps the position until the subsequent market closing; however, this strategy is suboptimal if the interest rate is positive, even in the absence of market frictions (cf. Theorem 1). Our study complements these papers by studying the slippage based on the optimal rebalancing strategy in a comprehensive setting, including non-zero interest rate, market frictions, and market closure. Furthermore, to the best of our knowledge, we are the first to find the "weekend effect" of LETFs slippage empirically; that is, the slippage tends to be the largest on Monday. By incorporating

market closure, our model can produce implications consistent with the empirical finding. Our model can also suggest a testable implication of a U-shaped intraday rebalancing volume of LETFs.

Third, we demonstrate the principle of "aiming in front of the target and moving gradually towards the aim", proposed in Gârleanu and Pedersen (2013, 2016), under the setting of hedging periodic cash flows. In the mean-variance optimization setting of Gârleanu and Pedersen (2013, 2016), one immediately moves to the optimal Markowitz portfolio and stays on it by continuous rebalancing in the absence of market frictions. However, in the presence of frictions, they showed that it is optimal to gradually adjust the portfolio towards an aim, which is not the current Markowitz portfolio but rather a weighted average of the current and future Markowitz portfolios. In our hedging setting, the optimal position also moves towards the aim gradually in a mean-reverting manner, where the aim stays close to the current target that minimizes today's return deviation in the morning hours and moves in front of the current target, converging to tomorrow's target in the afternoon hours, so as to prepare for the future.

While the weight between current and future targets remains constant in Gârleanu and Pedersen (2013, 2016), the weight in our paper is time-varying. This distinct feature arises from the special structure of our model involving the end-of-day discrete monitoring and intraday continuous rebalancing. It provides an extra economic insight on how much the aim should be "in front of target" in the presence of competing goals. Moreover, we report two benefits of applying this principle in LETFs' daily rebalancing: (1) By aiming in front of the current target and looking into the future, our optimal rebalancing strategy reduces the average daily slippage significantly compared to the one-day strategy that only focuses on minimizing today's deviation; (2) Moving gradually towards the target smooths out the daily rebalancing and results in a large decrease in the end-of-day trading volume, as compared to the one-shot strategy that only adjusts immediately before market closure.

Fourth, from a technical viewpoint, we obtain an analytic solution for the rebalancing strategy whose coefficients are determined by a system of periodic ordinary differential equations (ODEs)[2]. This is made possible because the quadratic structure of the model reduces the high-dimensional problem into a system of nonlinear ODEs, which leads to analytic tractability. [3] Specifically, one

[2] In a system of standard ordinary differential equations, the terminal or initial conditions are specified exogenously. However, in a system of periodic ordinary differential equations, the terminal or initial conditions are determined endogenously using periodicity requirements.

[3] The LETF rebalancing problem in this paper can be linked to the optimal execution problem, which aims to achieve a pre-specified position level within a fixed amount of time (e.g., from market opening to market closing) under frictions (see Bertsimas and Lo (1998), **?**, Frei and Westray (2015)). However, while the LETF rebalancing problem also aims at achieving a position level at the market closure subject to market friction costs, there are some key differences: (i) For our LETF rebalancing problem, the position target at market closure is *unknown* right after the market opening and only fully revealed at market closure, as it depends on the realized fund value X and underlying

only needs to solve the system once to determine the coefficients, which can then be used any time in the future. Since the periodic ODEs are nonlinear, we verify carefully that the solution exists and does not blow up and show that an iterative algorithm solving the periodic ODEs (along with the endogenous terminal conditions) has a unique fixed point solution.

The goal of LETFs is to achieve a fund return that is β times the underlying index return on a daily basis. Thus, even if they accomplish this goal perfectly every day, the cumulative returns are unlikely β times the underlying index returns over multi-day periods due to the compounding effect. However, this goal of LETFs was misunderstood by many investors, and the compounding effect was ignored. Indeed, both the U.S. Securities and Exchange Commission and Financial Regulatory Industry Authority warned investors of the significance of the multi-period compounding effect for a buy-and-hold strategy over an extended period. There are analytical and approximated solutions for the multi-period compounding effect, first proposed in Cheng and Madhavan (2009) and extended to a more general setting in Avellaneda and Zhang (2010) and Haugh (2011). Unlike this strand of literature, our primary focus is on the daily slippage rather than deviation for multi-day returns, and our objective is to minimize such daily slippages over intraday rebalancing rather than the compounding effect. We provide a discussion in Online Supplement H for a direct link to the compounding effect in the above literature.

The remainder of this paper is organized as follows. Section 2 introduces the background and notations of the problem. Section 3 presents the main result of this paper, an optimal LETF rebalancing model with market frictions, via numerical examples. Section 4 shows LETF's position under the optimal rebalancing strategy. Section 5 explains the usage of the principle of aiming in front of target in LETFs rebalancing, as well as its benefits. Sections 6 and 7 show the empirical results on slippage, intraday rebalancing pattern, and compounding effect. Section 8 concludes. All technical results and related proofs are given in the appendices.

2. Basic Setting and Notations

We consider an infinite horizon with $0 = t_0 < t_1 < \cdots < t_{2i} < t_{2i+1} < \cdots$, where t_{2i} and t_{2i+1} are respectively the market opening and closing of day i, $i = 0, 1, \cdots$. Hence, market is open during

price S at market closure. As a result, one has to learn the future target as time goes on and watch for the impact of the past strategy from the market opening on the fund value X. In contrast, in optimal execution literature, the target position at the end of the period is deterministic and known right at the beginning of trade. (ii) The measurement of the deviation from the benchmark in our problem is different from that in the optimal execution literature with performance benchmarks (e.g., the VWAP in Frei and Westray (2015)), which leads to significant technical differences. In optimal execution, the deviation is measured as the aggregated instantaneous distance between stock position (number of shares) and an exogenous target position over a finite horizon. However, in our problem, the deviation is measured as the aggregated distance between simple returns at an infinite set of discrete-time points (i.e., daily market closure). This, together with the continuous-time rebalancing, creates a periodic structure in our problem, which is not covered by the optimal execution literature.

time intervals $[t_{2i}, t_{2i+1}]$ when trading is allowed; market is closed during time intervals (t_{2i+1}, t_{2i+2}) when no trading is allowed. Denote the length of daytime (i.e., trading hours) for each day as T, and the length of nighttime (i.e., market closure) as δT. Therefore the length of each day is $(1+\delta)T$, and $t_{2i} = i(1+\delta)T$, $t_{2i+1} = T + i(1+\delta)T$ for $i = 0, 1, \cdots$. Denote the total daytime as $\cup_{i=0}^{\infty}[t_{2i}, t_{2i+1}]$, and the total nighttime as $\cup_{i=0}^{\infty}(t_{2i+1}, t_{2i+2})$. The underlying asset price S evolves as a geometric Brownian motion with different constants during daytime and nighttime:

$$dS_u = \mu(u)S_u du + \sigma(u)S_u dW_u, \tag{1}$$

where $\mu(u) = \mu_d$ and $\sigma(u) = \sigma_d$ for daytime, $\mu(u) = \mu_n$ and $\sigma(u) = \sigma_n$ for nighttime, and the parameters μ_d, μ_n, σ_d, and σ_n are all positive constants.[4] The fund NAV is invested in θ shares of S, and the remaining in a risk-free account with a constant interest rate $r > 0$.[5]

Thus, the dynamics of the LETF's total net asset value (NAV), X, is given by

$$dX_u = \theta_u dS_u + (X_u - \theta_u S_u)r du, \quad u \in [t_{2i-1}, t_{2i+1}), \forall i \geq 0 \tag{2}$$

$$X_{t_{2i+1}} = (1-\gamma)X_{t_{2i+1}-}. \tag{3}$$

Here, (3) corresponds to the industry practice of deduction of management fee from the NAV at market closing $t_{2i+1}, i \geq 0$, where γ is the daily management fee rate (γ typically has an annualized value of about 1%, although it varies across funds). Furthermore, we require that the position θ is adapted, and constant during nighttime when no trading is allowed.

The LETF's daily return on i-th day is calculated using the NAV immediately before the market closing t_{2i+1} and the NAV at the market closing time t_{2i-1} on the previous day:[6] $R_i^X = \frac{X_{t_{2i+1}-} - X_{t_{2i-1}}}{X_{t_{2i-1}}}$. To emphasize the role of a daily management fee, $X_{t_{2i+1}-}$ and $X_{t_{2i+1}}$ denote the before-fee NAV and the after-fee NAV (that is, the NAV publicly announced after market closing).

[4] As the fund rebalancing is assumed to cause a temporary price impact, the fundamental price S here is not changed by the rebalancing.

[5] In practice, the LETFs also use (total return) swaps and futures to achieve their daily target exposure. For example, as of April 1, 2021, Ultra S&P500 from ProShares (a Bull LETF on S&P500 with a multiple of +2) has a total exposure (notional plus profit and loss) to swaps and futures of $4.01 billion, in addition to the $2.93 billion exposure to stocks. By using the total return swaps, the tracking problem is effectively transferred to the swap counterparty. However, the impact on the index is still present, albeit caused by the tracking effort of the swap counterparty. Also, the slippage cost to the fund manager is still present, since the counterparty can pass on the hedging cost to the fund (e.g., in forms of the LIBOR spread). By using futures, the impact is passed on similarly to the underlying asset via arbitrageurs; see Avellaneda and Dobi (2012), Wagalath (2014), and Cheng and Madhavan (2009) for more details. Such slippage cost can be passed on to the fund or to both the fund and the counterparty, depending on variables such as the fund's bargaining power compared to the counterparty. For example, when the fund has a very high bargain power overall, only a tiny fraction of the costs will be passed on to the fund.

[6] This agrees with the definition of LETFs daily return in practice. For instance, ProShare states in its prospectus that "A single day is measured from the time the Fund calculates its NAV to the time of the Fund's next NAV calculation," and "the NAV of each Fund ... is generally determined each business day as of the close of regular trading on the Exchange on which it is listed."

Therefore, R_i^X denotes the daily before-fee NAV return. Denote the daily underlying asset return as $R_i^S = \frac{S_{t_{2i+1}} - S_{t_{2i-1}}}{S_{t_{2i-1}}}$, so that the target return is βR_i^S. For the current day (i.e., $i = 0$), the notation $X_{t_{2i-1}}$ in R_0^X (resp. $S_{t_{2i-1}}$ in R_0^S) is the NAV (resp. underlying value) observed at the last market closing before the current day, and it will be denoted as $\bar{x}$ (resp. $\bar{s}$) throughout the paper.

Definition 1 (Slippage). The *slippage* on i-th day is the distance between LETF's daily before-fee NAV return and target return:

$$D_i = |R_i^X - \beta R_i^S|. \tag{4}$$

Note that to achieve the investment objective[7] exactly, D_i should be zero on every trading day.

Trading in the market is costly; trading at an instantaneous speed φ_u incurs a temporary price impact of $C(\varphi_u, S_u)du = \frac{1}{2}\tilde{\Lambda} S_u^2 \varphi_u^2 du$, where $\tilde{\Lambda}$ is a nonnegative constant, and the speed φ_u is such that $d\theta_u = \varphi_u du$. This type of quadratic price impact cost is also used in Obizhaeva and Wang (2013), Rogers and Singh (2010), Gârleanu and Pedersen (2016) and Moreau et al. (2017), and supported empirically by Breen et al. (2002) (see also Grossman and Miller (1988) and Greenwood (2005) for the justification in the multi-asset case). In the classic Kyle (1985) model, the equilibrium price set by the market maker increases linearly in the trader's amount of order; therefore, to the trader, the cost of price impact is quadratic in the trading amount. In our model, this price impact can be interpreted as the liquidity cost incurred during trading, e.g., from the presence of front-runners, especially immediately before the market closing when other market participants have a good estimation of LETFs' direction of rebalancing.[8] We also assume that C is quadratic in the underlying price S, which can be understood as that the price impact is higher when the underlying price is higher. This is especially important in the current problem, since the influence of asset price on its price impact is nonnegligible over an infinite horizon. Also, the multiplicative factor S^2 gives a natural scaling of the price impact and brings a technical convenience allowing for an analytical solution to the problem.

[7] As stated in the prospectus of ProShares fund, "The fund seeks daily investment results, before fees and expenses, that correspond to two times (2x) the daily performance of the index. The fund does not seek to achieve its stated investment objective over a period of time greater than a single day." Here, "two times" refers to the multiple β and is changed accordingly for other funds with different multiples.

[8] There is a strand of empirical literature on the impact of LETF rebalancing on the underlying index. Tuzun (2014) finds that daily LETF rebalancing leaves an imprint on various U.S. equity categories, triggering price reactions and increased volatility, especially on volatile days. Bai et al. (2012) show similar impacts for LETFs on real estate indices. More recently, Charupat and Miu (2016) give partial support that LETFs rebalancing contributes to the index movement before market closing, especially for indices with less liquidity, such as Russell 2000. One may refer to Acharya and Pedersen (2005) and Brunnermeier and Pedersen (2005) for the impact of liquidity risk on asset pricing or trading, and also Cont et al. (2014) and Cartea et al. (2015) about the micro-foundations behind it at the limit order book level.

The fund's objective is to minimize the difference between the daily simple return of X and the target return at market closing of each trading day over an infinite horizon, with an additional penalty on the trading speed φ. Specifically, the fund minimizes the total cost

$$\sum_{i=0}^{\infty} e^{-\rho t_{2i}} \left(\frac{1}{2} X_{t_{2i-1}}^2 D_i^2 + \nu \int_{t_{2i}}^{t_{2i+1}} \frac{1}{2} \tilde{\Lambda} S_u^2 \varphi_u^2 du \right). \tag{5}$$

Here, $\frac{1}{2} X_{t_{2i-1}}^2 D_i^2$ represents the daily slippage cost, which is larger for a larger slippage.[9] The factor $X_{t_{2i-1}}^2$ provides a natural scaling for the deviation cost; indeed, a larger fund will hold a larger position and thus have a larger price impact, and therefore such scaling is required to keep the relative importance between the two types of costs. $\rho > 0$ is a subjective discount rate, representing the weight in the trade-off between optimizing short-term performance or long-term performance: the larger ρ is, the greater is the emphasis on short-term costs as compared to the long-term costs. As a special case, $\rho = +\infty$ means that the manager only cares about today. The summation of the first and second terms in the parenthesis corresponds to the daily total cost. In particular, the first term represents the daily return deviation cost from the fund's slippage, and the second term represents the daily accumulated cost from the market frictions incurred by the intraday fund rebalancing. Note that the treatment of incorporating friction cost as a separate penalty term has been used extensively in the finance literature, e.g., Gârleanu and Pedersen (2013, 2016), or in macro-economy literature, e.g., Hansen and Sargent (2013).[10] The parameter $\nu \geq 0$ serves the role of prioritizing the goals of minimizing the deviation cost or minimizing the market friction costs on each day. For example, a small ν can reflect the case where, when using the futures and swaps, only a small fraction of price impact cost is transferred from the counterparty to the fund. Taking $\nu = 0$ means that the manager ignores the market friction costs and focuses solely on minimizing the slippage. To simplify notations, in the following we denote $\Lambda = \nu \tilde{\Lambda}$.

3. Theoretical Results

In this section, we solve the cost and deviation minimization problem introduced in Section 2, in the case without market frictions ($\Lambda = 0$) and with market frictions ($\Lambda > 0$). In either case, we derive the explicit value function whose coefficients are the unique solution to a system of ordinary differential equations, and the resulting explicit optimal rebalancing strategy.

[9] Note that this cost is not actually paid by the fund in practice; instead, one can regard this as a cost from reputational damage. Large return deviations are costly to LETFs in the sense that they make the fund less attractive and drive away investors.

[10] An alternative formation is to pay the cost out of X. However, this alternative formulation has one major drawback: In case that X achieved a return higher than the target return right before market closing, the fund manager can reduce the deviation by deliberately incurring cost from market frictions and pulling down X. This strategy seems unethical, since it throws away fund's value via triggering market friction cost intentionally. Also, this alternative formulation brings difficulty to the modeling, since the optimal strategy may not be well-defined.

3.1. The Case without Market Frictions

Without market frictions, i.e., $\Lambda = 0$, the total cost (5) reduces to the aggregated daily slippage. At the current time $0 \le t < t_1$, the cost minimization problem becomes

$$\begin{aligned} V(t,s,x,\bar{s},\bar{x}) = \inf_{\theta} E\Bigg[& \frac{1}{2}\bar{x}^2 \left(\beta\left(\frac{S_{t_1}}{\bar{s}} - 1\right) - \left(\frac{X_{t_1-}}{\bar{x}} - 1\right)\right)^2 \\ & + \sum_{i=1}^{\infty} e^{-\rho t_{2i}} \frac{1}{2} X_{t_{2i-1}}^2 \left(\beta\left(\frac{S_{t_{2i+1}}}{S_{t_{2i-1}}} - 1\right) - \left(\frac{X_{t_{2i+1}-}}{X_{t_{2i-1}}} - 1\right)\right)^2 \Bigg], \end{aligned} \tag{6}$$

where $S_t = s$ and $X_t = x$, and recall that $\bar{s}$ and $\bar{x}$ are the *reference values* of the underlying asset and NAV observed at the last market closing before today, respectively.[11] The first and second terms inside the expectation are the deviation on the first day and the aggregated deviation from the second day, respectively. Note that at $t = 0$, this minimizes the total cost as defined in (5).

Besides keeping track of X and S at any time, we also need to keep track of their reference values $\bar{x}$ and $\bar{s}$, in order to calculate the daily return at today's market closing. On each day, the current value of S and X, together with their reference values, provide an expectation of the daily return and are hence important in the rebalancing decision. Note that $\bar{s}$ and $\bar{x}$ are not necessarily equal to the time-t underlying value s and NAV x, respectively, even at $t = 0$ due to the overnight jump.

THEOREM 1. *(The minimal cost and optimal rebalancing strategy without market frictions). The minimal cost* (6) *is given as* $V(t,s,x,\bar{s},\bar{x}) = \bar{x}^2 V\left(t, \frac{s}{\bar{s}}, \frac{x}{\bar{x}}, 1, 1\right)$, *where* $V(t,s,x,1,1) = a(t)x^2 + b(t)xs + c(t)s^2 + d(t)x + e(t)s + f(t)$, $\quad t \in [0, t_1)$, *and the vector of coefficients* (a,b,c,d,e,f) *is the unique solution to a system of periodic ODEs* (F.3) – (F.8) *with endogenous terminal conditions* (F.10). *Furthermore, the optimal position level is*

$$\theta_t^{fl} = \begin{cases} -\frac{1}{2\sigma^2 a(t)} \left[b(t)\sigma^2 \frac{\bar{x}}{\bar{s}} + (\mu - r)\left(b(t)\frac{\bar{x}}{\bar{s}} + d(t)\frac{\bar{x}}{S_t} + 2a(t)\frac{X_t}{S_t} \right) \right] & t \in [0, t_1) \\ -\left[M\left(\frac{d(0)}{2a(0)} + e^{r\delta T} \right) + \frac{b(0)}{2a(0)} \left(1 + e^{r\delta T} M\right) \right] \frac{X_{t_1}}{S_{t_1}} & t \in [t_1, t_2). \end{cases} \tag{7}$$

where $M = \frac{e^{\mu_n \delta T} - e^{r\delta T}}{(e^{\mu_n \delta T} - e^{r\delta T})^2 + e^{2\mu_n \delta T}(e^{\sigma_n^2 \delta T} - 1)}$. *In particular, if* $r = 0$, *then*

$$\theta_t^{fl} = \frac{\beta \bar{x}}{\bar{s}}, \quad \text{if } t \in [0, t_1) \text{ and } \frac{\beta X_{t_1}}{S_{t_1}} \text{ if } t \in [t_1, t_2). \tag{8}$$

Theorem 1 shows that the minimum cost is a quadratic function of the normalized stock price $s/\bar{s}$ and portfolio value $x/\bar{x}$. The optimal position level is determined by the expression of θ^{fl}, and (2) – (3) with θ replaced by θ^{fl}. During the daytime (i.e., $t \in [0, t_1)$), the current optimal position level θ_t^{fl} is continuous with respect to t, depending on the coefficients of the minimal cost function,

[11] We keep the terms $\bar{x}^2$ and $X_{t_{2i-1}}^2$ due to tractability and comparison with the general cases with market friction.

the current underlying and fund values $S_t = s$ and $X_t = x$, as well as the reference values $\bar{s}$ and $\bar{x}$ for $[0, t_1)$. Right at the market closing time t_1, it is optimal to perform a lump-sum trade so as to adjust the position to $\theta_{t_1}^{fl}$. The position is then kept at this level during nighttime. At the next market opening, the current time t reset to 0, and the trading continues following (7), based on updated reference values $\bar{x} = X_{t_1}$ and $\bar{s} = S_{t_1}$.

The special case with $r = 0$ gives a much simpler solution, given by (8). That is, it is optimal to perform a *one-shot* rebalancing at every market closing time based on the underlying and fund value at market closing, and then keep the position unchanged until the next market closing, so on and so forth. If the interest rate is zero in the market, then this one-shot strategy leads to zero slippage. Indeed, this agrees with the common practice by fund managers to adjust the portfolio near market closing on each day to keep the target leverage ratio (cf. Avellaneda and Dobi (2012), Wagalath (2014)).

It is worth noting that even without market frictions, when $r > 0$, ignoring the interest rate and sticking to this one-shot rebalancing strategy (8) may lead to a non-zero daily slippage $D_i = |1-\beta|(e^{rT}-1)$.[12] Indeed, the positive interest rate means that the fund pays the funding cost for doing leverage and receives interest for short-selling. Actually, this is the motivation for Tang and Xu (2013) and Henderson and Buetow (2014) to explain slippage in terms of interest rate: the fund ignores the interest cost when rebalancing, and therefore the interest cost will affect the fund's return and result in the slippage. In contrast, Theorem 1 considers not just the interest rate but also the market closure. Nevertheless, since the daily interest rate is typically small, the amount of rebalancing required before market closing is small without market frictions, as will be illustrated numerically in Section 4.

3.2. The Case with Market Frictions

Now we study the optimal daily rebalancing strategy that takes market frictions into account. For example, with a quadratic trading cost, it is no longer feasible to do lump-sum trading, since it incurs an infinite cost. To incorporate market frictions, we impose a mild technical requirement that θ_u is absolutely continuous with trading speed φ_u, i.e.,

$$d\theta_u = \varphi_u du, \quad \theta_t = z. \tag{9}$$

Here, the adapted admissible control variable $\varphi \in \mathcal{A}$, where the set of admissible strategies $\mathcal{A}$ consists of control variable φ that results in a finite expected market frictions cost and equals 0 during nighttime.[13]

[12] Specifically, $R_i^X = \theta^*_{t_{2i-1}} \cdot \frac{S_{t_{2i+1}} - S_{t_{2i-1}}}{X_{t_{2i-1}}} + \frac{X_{t_{2i-1}} - \theta^*_{t_{2i-1}} S_{t_{2i-1}}}{X_{t_{2i-1}}}(e^{rT}-1) = \frac{\beta X_{t_{2i-1}}}{S_{t_{2i-1}}} \frac{S_{t_{2i+1}} - S_{t_{2i-1}}}{X_{t_{2i-1}}} + (1-\beta)(e^{rT}-1) = \beta \cdot R_i^S + (1-\beta)(e^{rT}-1)$.

[13] More precisely, we define $\mathcal{A} = \left\{\varphi : \varphi_u = 0 \text{ for } u \in \cup_{i=0}^{\infty}(t_{2i+1}, t_{2i+2}) \text{ and } E\left[\sum_{i=0}^{\infty} e^{-\rho t_{2i}} \int_{t_{2i}}^{t_{2i+1}} \varphi_u^2 S_u^2 du\right] < \infty\right\}$.

For $0 \le t < t_1$, $\bar{s}, s \in \mathbb{R}^+$, $\bar{x}, x, z \in \mathbb{R}$, $\Lambda > 0$, the value function is defined in a similar way as (6)

$$\begin{aligned} &V(t,s,x,z,\bar{s},\bar{x}) \\ &= \inf_{\varphi\in\mathcal{A}} E\left[\frac{1}{2}\bar{x}^2\left(\beta\cdot\left(\frac{S_{t_1}}{\bar{s}}-1\right)-\left(\frac{X_{t_1-}}{\bar{x}}-1\right)\right)^2+\int_t^{t_1}\frac{1}{2}\Lambda S_u^2\varphi_u^2 du\right. \\ &\left.+\sum_{i=1}^{\infty}e^{-\rho t_{2i}}\left(\frac{1}{2}X_{t_{2i-1}}^2\left(\beta\cdot\left(\frac{S_{t_{2i+1}}}{S_{t_{2i-1}}}-1\right)-\left(\frac{X_{t_{2i+1}-}}{X_{t_{2i-1}}}-1\right)\right)^2+\int_{t_{2i}}^{t_{2i+1}}\frac{1}{2}\Lambda S_u^2\varphi_u^2 du\right)\right], \end{aligned} \tag{10}$$

where the expectation is computed under the initial value $S_t = s$, $X_t = x$ and $\theta_t = z$.

Compared to (6), the value function (10) has the additional terms for the market friction costs. Since it is now only feasible to adjust the position θ at a finite speed φ, the initial position z becomes relevant for the calculation of the market friction costs and is hence required as a state variable. To see why the initial position matters, starting from a position that is farther away from the optimal position, the manager needs to perform a larger rebalancing to push the position towards the optimal one, which in turn triggers a higher cost from the market frictions.

The following theorem gives an explicit form of the value function $V(t,s,x,z,\bar{s},\bar{x})$ in the case with market friction, as well as the optimal rebalancing strategy.

THEOREM 2. *(The optimal value function and the optimal rebalancing strategy with market frictions). The value function is given as* $V(t,s,x,z,\bar{s},\bar{x}) = \bar{x}^2 V\left(t, \frac{s}{\bar{s}}, \frac{x}{\bar{x}}, \frac{\bar{s}}{\bar{x}}z, 1, 1\right)$, *where*

$$\begin{aligned} V(t,s,x,z,1,1) = a(t)x^2 + (b_0(t)+b_1(t)z)xs + (c_0(t)+c_1(t)z+c_2(t)z^2)s^2 + d(t)x \\ + (e_0(t)+e_1(t)z)s + f(t), \end{aligned} \tag{11}$$

where the coefficients are determined as the unique solution to a system of periodic ODEs (A.9) – (A.18) *subject to endogenous terminal condition* (D.1). *Furthermore, the optimal rebalancing strategy* φ^* *equals*

$$\varphi_t^* = -\frac{1}{\Lambda}\left(b_1(t)\frac{X_t}{S_t} + c_1(t)\frac{\bar{x}}{\bar{s}} + 2c_2(t)\theta_t^* + e_1(t)\frac{\bar{x}}{S_t}\right), \quad t \in [0, t_1), \tag{12}$$

where the optimal position θ^* *satisfies* $d\theta_t^* = \varphi_t^* dt$.

The normalized value function (11) for $\bar{s} = \bar{x} = 1$ is a polynomial of degree 4 with respect to (s,x,z), which is different from the quadratic value function derived in Gârleanu and Pedersen (2013). Specifically, the underlying value s appears in (11), and z appears as a multiplicative factor in front of terms involving s (i.e., s^2, s, xs) with the highest order matching the order of s. In contrast, z also does not appear in the value function in Gârleanu and Pedersen (2013) since it does not involve the underlying value s.

The optimal rebalancing strategy φ^* depends on the current underlying and fund values $S_t = s$ and $X_t = x$ as well as the reference values $\bar{s}$ and $\bar{x}$. In addition, φ^* depends on the current position

level, so that the optimal position and optimal trading speed form a feedback system. During nighttime, $\varphi^* = 0$ since no trading is allowed, and at the next market opening, the reference values are updated to $\bar{s} = S_{t_1}$ and $\bar{x} = X_{t_1}$ by definition, t reset to 0, and the optimal trading speed is again determined by (12). Since one only needs to solve the ODE system once to determine the coefficients, this optimal strategy can be implemented very efficiently.

From a technical point of view, the ODE system (A.9) – (A.18) for the coefficients of (11) (and also the system (F.3) – (F.8) for the frictionless case) are nonlinear and has a periodic terminal condition. Even in the one-period subproblem, the global existence of solution is not guaranteed by the classic theory, since it may blow up in finite time. In Lemma 1 of Online Supplement, we verify carefully that the solution indeed exists in one period, given a suitable terminal value. Also, since the periodic ODEs are solved in iteration, one needs to show that this iteration indeed has a unique fixed point to guarantee convergence. Proposition 3 in Online Supplement shows that, as long as the discount rate is positive, we indeed have a fixed point. Finally, the uniqueness of the solution is guaranteed by the verification theorem, Proposition 4. In the remaining of the paper, we shall discuss the financial implications of our rebalancing strategy mainly via numerical illustration.

4. Optimal Daily Position

This section shows the position under the optimal rebalancing strategy, where the rebalancing starts every day before market closing and finishes after the next day's market opening.

Definition 2 (Strategies for Comparison). We define the following three rebalancing strategies, all of which are given via (12) but with different coefficients b_1, c_1, c_2, and e_1.

Periodic-DN The *Periodic Day-and-Night* strategy is the optimal strategy that considers an infinite horizon and market closure, given as (12) whose coefficients are calculated under $\delta = 17.5/6.5$ (noting the total trading hours is 6.5) and $\rho = 0.6$;

One-Day The *One-Day* strategy is the myopic suboptimal strategy that only aims at minimizing today's cost and ignores the presence of market closure, given as (12) whose coefficients are calculated under $\rho = +\infty$;

Periodic-D The *Periodic Day-Only* strategy is the suboptimal strategy that considers an infinite horizon but ignores market closure, given as (12) whose coefficients are calculated under $\delta = 0$ and $\rho = 0.6$.[14]

[14] For $\rho = +\infty$, the value function (10) simplifies to

$$V(t,s,x,z,\bar{s},\bar{x}) = \inf_{\theta\in\mathcal{A}} E\left[\frac{1}{2}\bar{x}^2\left(\beta\left(\frac{S_{t_1}}{\bar{s}}-1\right)-\left(\frac{X_{t_1-}}{\bar{x}}-1\right)\right)^2+\int_t^{t_1}\frac{1}{2}\Lambda S_u^2\varphi_u^2 du\right].$$

The corresponding coefficients can be calculated using (A.9) – (A.18) and plugging $\rho = +\infty$ in (A.19).

In the following, we use Monte Carlo simulation to estimate the average slippage. For this purpose, we simulate the underlying asset price S via (1) during daytime and its jumps overnight, using the default parameter values to be specified. Along the simulated path of S, one can calculate the net asset value X of the LETF for each day via (2), (3), (9), and (12), and then calculate the daily slippage using (4). Note that the reference values $X_{t_{2i-1}}$ and $S_{t_{2i-1}}$ in (12) are given as model input on day $i=0$ (i.e., $t\in[0,t_1)$).

Since the daytime volatility σ_d is k times the nighttime volatility σ_n, we have $\sigma_d = k\sigma\cdot\sqrt{\frac{T+\delta T}{k^2T+\delta T}}$ and $\sigma_n=\sigma\cdot\sqrt{\frac{T+\delta T}{k^2T+\delta T}}$, where σ is the average volatility. Here we choose $k=3$ and $\sigma=0.2$. On the other hand, we use the same expected return for daytime and nighttime as suggested by empirical studies such as Stoll and Whaley (1990) and Lockwood and Linn (1990) and take $\mu_d=\mu_n=0.1$. The default values for other parameters are[15] $\delta=17.5/6.5$, $\nu=1$, $\beta=2$, $\rho=0.6$, and $r=\gamma=0.01$. All these parameter values are annualized.[16] Furthermore, we take $\Lambda=\nu\tilde{\Lambda}=10^{-6}$ to match the level of magnitude of price impact documented in Robert et al. (2012).[17] Finally, we consider $s=x=\bar{s}=\bar{x}=1, z=2$ (so the position is optimal at t_0, and no overnight jump occurs before t_0).

We first look at the position level under the optimal strategy, which directly determines the performance. To this end, we simulate a sample path of the underlying asset value over a 5-day period using the default parameter values. Due to the overnight jumps, the underlying value at the market closing time is not necessarily equal to the value at the next market opening. Then, via (9) and (12), we calculate the optimal position θ^{PDN} based on the Periodic-DN strategy, θ^{O} based on the One-Day strategy, and θ^{PD} based on the Periodic-D strategy, as shown in Figure 1. We also plot the frictionless target θ^{fl}, that is, the optimal position without market frictions as given by (7) in Theorem 1. Finally, we overlay the 5-day trajectory of the underlying asset as the dotted line corresponding to the right-hand side axis.

Figure 1 shows that, for all three strategies, the fund needs to increase (resp. decrease) the position θ by market closing if the underlying asset return is positive (resp. negative) during that day. The magnitude of such an increase or decrease in θ is also proportional to the magnitude of

[15] ρ is a subjective discount rate that represents the relative importance between short-term slippage (e.g., today and next few days) and long-term slippage. Note that ρ is not the rate for discounting cash flow. $\rho=0.6$ roughly means that the daily slippage one year later only accounts for half the importance of the same slippage today. ρ should not be very small (i.e., not on the level of interest rate), because if the current slippage is large, then it can already cause damage to the fund's reputation. In terms of "aiming in front of target", a larger ρ means the aim is "less in front of target", so as to achieve a better performance today while sacrificing future performance. In the extreme case where $\rho=+\infty$, the fund only cares about today's slippage and does not care about the future slippage at all.

[16] Specifically, by normalizing $T=1$ in the following calculations, the corresponding effective values are $\mu_d=\mu_n=\frac{0.1}{252(1+\delta)}$, $\sigma=\frac{0.2}{\sqrt{252\times(1+\delta)}}$, and $\rho=\frac{0.6}{252\times(1+\delta)}$, $r=\frac{0.01}{252\times(1+\delta)}$ per trading session (6.5 hours), and $\gamma=\frac{0.01}{252}$ per day.

[17] Specifically, Robert et al. (2012) reported a cost level TC that is 0.1009% of the dollar transaction amount for all the stock transactions in their sample. In our case, denote $\Delta\theta$ as the daily transaction amount in number of shares, $\frac{\tilde{\Lambda}}{2}S^2(\Delta\theta)^2=\mathrm{TC}^2$, meaning $\tilde{\Lambda}=2\times\left(\frac{\mathrm{TC}}{S\Delta\theta}\right)^2=2\times(0.1009\%)^2=2\times10^{-6}$.

daily asset price return. This is consistent with the changes in the frictionless target. To understand this, when the interest rate is low and there is no friction, the fund would ideally hold $\theta^{fl}_{t_{2i+1}}$ share of the risky asset at market closing t_{2i+1}, which provides a leverage ratio of β approximately. If the underlying asset has a positive (resp. negative) daily return, the magnitude of return of the fund value X is greater than that of the underlying asset. If θ remains unchanged, the fund would be under-leveraged (resp. over-leveraged) at market closing. Thus, to keep the leverage ratio close to β at market closing, the fund has to increase (resp. decrease) θ.

Figure 1 Comparison of intraday Positions from Three Strategies

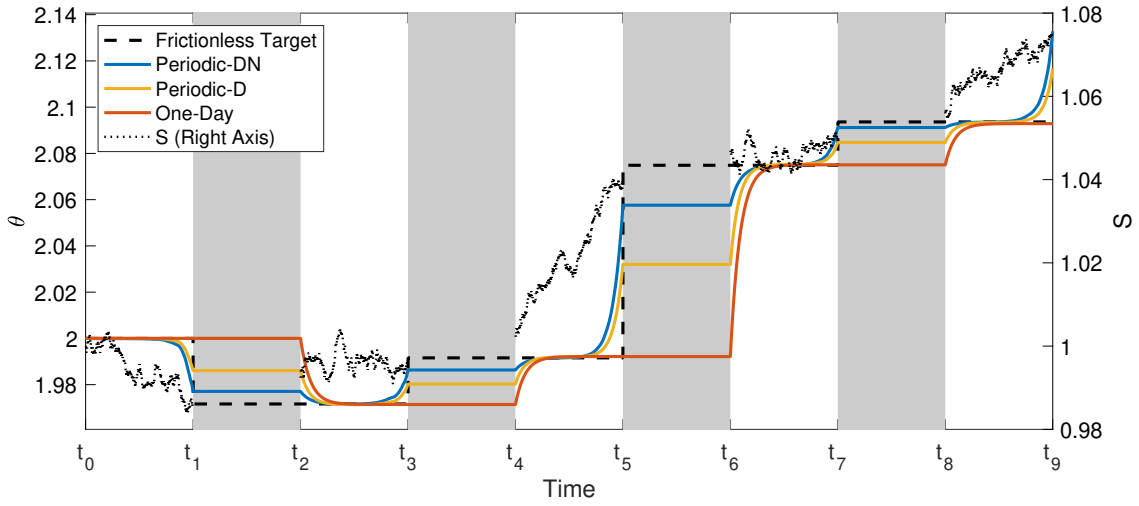

This figure compares the positions for the Periodic-DN strategy, Periodic-D strategy, and One-Day strategy based on a simulated asset price trajectory for 5 days (dotted line with right-hand side axis). The dashed line denotes the frictionless target (technically speaking, the 5-day trajectories of frictionless targets under the three strategies differ, as they depend on the realized trajectories of X. However, since such differences are negligible over the 5-day period in this figure, only the frictionless target corresponding to the Periodic-DN strategy is presented for a better illustration). The strategies are defined in Definition 2. Time periods with white and gray backgrounds denote daytime and nighttime, respectively. By comparing the difference between the dotted line and the solid line, it appears that the Periodic-DN strategy produces a level of position that is on average closest to the frictionless target, especially during nighttime.

Figure 1 shows two advantages of the Periodic-DN strategy. First, it results in a smaller overall distance of the position level to the frictionless target. Indeed, for the Periodic-DN and Periodic-D strategies, during daytime, the position first moves towards and then stays close to the target after the market opening. Then, right before market closing, it moves away from today's target and tries to shoot for the predicted target for the next day. Due to the presence of market frictions, the position level does not hit tomorrow's target. As a result, further rebalancing is required after the next day's market opening. In contrast, the position level of the One-Day strategy does not change significantly in the late trading hours and has a much larger distance from tomorrow's target before the market closing time and during nighttime. As a result of this myopic strategy, the position is significantly off the target level right after tomorrow's market opening, and one requires a greater effort to finish the rebalancing in the next day.

Second, by comparing the position levels of the Periodic-DN and Periodic-D strategies in Figure 1, we can observe that accounting for overnight risk leads to a better end-of-day position. Indeed, by accounting for the overnight risk, the fund adjusts its position at a faster pace late in daytime, so that the fund has a better position at the market closing and during nighttime. Consequently, the fund is exposed to a smaller overnight jump risk, and a smaller rebalancing is required for the position to reach the target after tomorrow's market opening. The discussion in Section 6.1 confirms that this indeed leads to a smaller slippage as compared to the Periodic-D strategy that ignores the overnight jump risk.

Figure 1 also illustrates that the frictionless target position remains largely constant over daytime and jumps to the target level for the next day at the market closing. Therefore, although continuous rebalancing is optimal during daytime for $r > 0$ as suggested by Theorem 1, the amount of rebalancing is negligible if there is no friction.

5. Aiming in Front of Target

What is the underlying reason for the difference in the position levels as depicted in Figure 1, especially before the market closing time? To answer this question, we show in this section that the optimal strategy follows the principle of aiming in front of target and moving gradually towards aim. To this end, we rewrite the optimal rebalancing strategy (12) as

$$d\theta_t^* = \kappa_t^* \left(\bar{\theta}_t^* - \theta_t^*\right) dt, \tag{13}$$

for $t \in [0, t_1)$, where

$$\bar{\theta}_t^* = -\frac{b_1(t)}{2c_2(t)}\frac{X_t}{S_t} - \frac{c_1(t)}{2c_2(t)}\frac{\bar{x}}{\bar{s}} - \frac{e_1(t)}{2c_2(t)}\frac{\bar{x}}{S_t}, \quad \kappa_t^* = \frac{2c_2(t)}{\Lambda}. \tag{14}$$

Equation (13) shows that the optimal position exhibits a mean-reverting pattern, with *speed* $0 \leq \kappa^* < \infty$, and *aim* $\bar{\theta}^*$. The finite speed κ^* is consistent with the principle of "moving gradually toward aim" studied in Gârleanu and Pedersen (2013), due to the existence of market friction. The speed is lower for larger friction; if there is no friction ($\Lambda = 0$), the speed κ^* is infinite, meaning that the position follows the aim exactly. The distance between the current position level and the aim $\bar{\theta}^*$, together with κ^*, decides the optimal trading speed.

Also, from the analytic form (11) of V, the aim can be represented as

$$\bar{\theta}_t^* = \arg\min_z V(t, S_t, X_t, z, \bar{s}, \bar{x}).$$

Thus, the aim is the position that minimizes the expected optimal total future costs given the current information. This interpretation is consistent with Gârleanu and Pedersen (2016). More precisely, the aim portfolio (given by Equation (11) in Gârleanu and Pedersen (2016)) is exactly

the minimizer of the value function (given by Equation (6) in their paper). Their aim can be represented as a weighted average of the future expected frictionless targets. However, a similar representation for $\bar{\theta}_t^*$ here is difficult because our frictionless target depends on the past strategy via the current fund value, creating a "feedback" loop because the position tries to follow an aim that is itself affected by the position. This is in stark contrast to their model, where the frictionless Markowitz target does not depend on the past strategy. The coefficients $b_1,\ c_1,\ c_2$, and e_1 appearing in $\bar{\theta}_t^*$ are determined by the system of Riccati ODEs (A.9) – (A.18), which is in general difficult to solve explicitly, although numerical solutions are readily available. Also, for $\Lambda \in (0,\infty)$, $b_1,\ c_1,\ c_2$, and e_1 can be decoupled from other coefficients only in the special case $\mu - r = \sigma = 0$; however, this special case provides little extra economic insight as both S and X will have a deterministic growth at the risk-free rate.

Next we illustrate the relationship between the aim $\bar{\theta}_t^*$ for the Periodic-DN strategy with the frictionless target via a special example without market closure (that is, $\delta = 0$). From the terminal condition (A.19) of the ODE system in Online Supplement A.2, we know that

$$b_1(t_1-) = e^{-\rho t_1}(1-\gamma)(b_1(0)+c_1(0)+e_1(0)),\quad c_1(t_1-) = e_1(t_1-) = 0,\ \text{and}\ c_2(t_1-) = e^{\rho t_1}c_2(0).$$

As a result,

$$\bar{\theta}_0^* = \pi\frac{\beta\bar{x}}{\bar{s}}\quad\text{and}\quad \bar{\theta}_{t_1-}^* = \pi\frac{\beta X_{t_1}}{S_{t_1}},\quad \text{where } \pi = -\frac{b_1(0)+c_1(0)+e_1(0)}{2\beta c_2(0)}.$$

Numerically, we find that π is close to 1. This means that, for each day, the aim $\bar{\theta}_t^*$ is near the current day's frictionless target at market opening, and near the next day's frictionless target $\frac{\beta X_{t_1}}{S_{t_1}}$ at market closing. Furthermore, the transition of aim through daytime is continuous thanks to the continuity of the coefficients in (14). In the following, we show via numerical illustrations that the above observation also holds in the general case with market closure.

To illustrate how the position level moves towards the aim, Figure 2 plots the aim for the Periodic-DN strategy over the position and stock price trajectory in Figure 1. Figure 2 shows that, consistent with (13), the position θ^{PDN} is guided by the moving aim $\bar{\theta}^{PDN}$ and moves towards the aim. However, due to the market frictions, the position can only do so with a finite speed. As a result, it does not reach the aim at market closing (e.g., at t_5), but only reaches it halfway on the next day (between t_6 and t_7).

Since the position level is guided by the aim due to (13) and illustrated in Figure 2, to explain the differences of the position levels in Figure 1 from the three strategies, it is worth comparing their respective aims. Again, we calculate the coefficients b_1, c_1, c_2 and e_1 based on the Periodic-DN, Periodic-D and One-Day strategies. For illustration, we use the 5-day stock price sample path $\tilde{S}$

Figure 2 Position and Aim

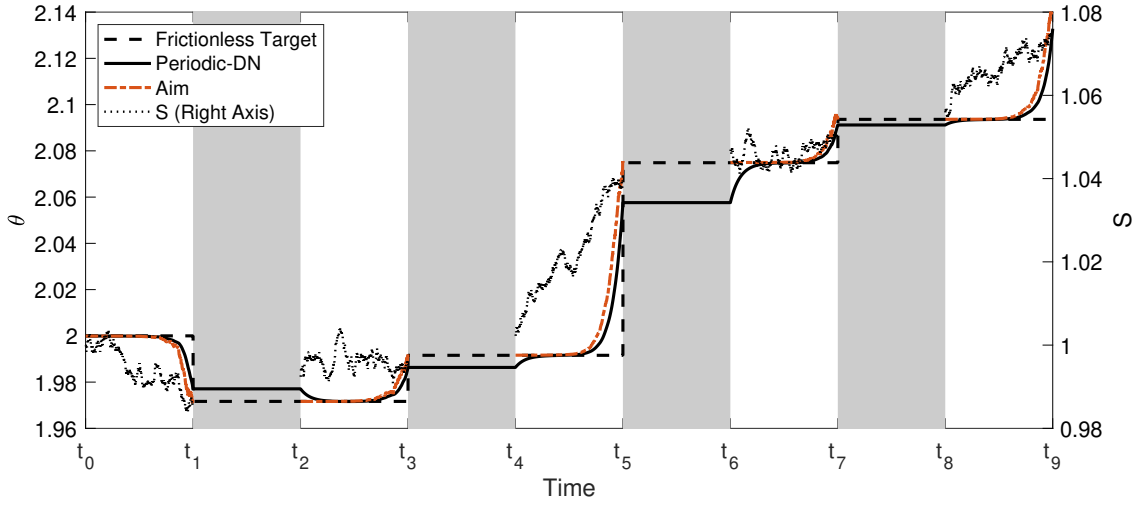

This figure compares the optimal position for the Periodic-DN strategy as in Figure 1 against its aim $\bar{\theta}^{PDN}$ (the dash-dotted curve).

as shown in Figure 1 (right-hand side axis), and calculate the path of $\tilde{X}$ under the Periodic-DN strategy using θ^{PDN} and (2) – (3). We then calculate the path for the aims $\bar{\theta}^{PDN}$, $\bar{\theta}^{O}$ and $\bar{\theta}^{PD}$ via (14), using the coefficients from the corresponding strategy, $\tilde{X}$ and $\tilde{S}$, as shown in Figure 3.[18] We also include the frictionless target for comparison.

We now look at $\bar{\theta}^{O}$ for the One-Day strategy. Figure 3 shows its one key characteristic: it stays on the frictionless target θ^{fl}, which is the optimal position level without market frictions, every day before market closing. As a result, $\bar{\theta}^{O}$ is updated *discontinuously* at the transition between today at t_{2i-1} and tomorrow t_{2i}, when the target position changes. Such behavior of the aim is because the One-Day strategy only focuses on today and does not prepare for tomorrow. Were there no market frictions, such a myopic aim would not be a problem, since the position could always catch up with the aim even if the aim jumps (by rebalancing at an infinite speed). However, due to the presence of market frictions, the actual position level has to move towards this aim gradually after the market opening to catch up with the aim, as demonstrated in Figure 1. In other words, the myopic position level of the One-Day strategy illustrated in Figure 1 is rooted in its myopic aim. This results in a large slippage as we will see in Section 6.

Next, we look at the aim $\bar{\theta}^{PDN}$ for Periodic-DN strategy. In the early trading hours of each day, it overlaps with the frictionless target, similar to that of the One-Day strategy. This suggests that,

[18] Note that the paths of $\bar{\theta}^{O}$ and $\bar{\theta}^{PD}$ illustrated in Figure 3 are based on the same sample path of $\tilde{S}$, as well as the path of $\tilde{X}$ under Periodic-DN strategy, which are different from $\bar{\theta}^{O}$ and $\bar{\theta}^{PD}$ calculated based on the path of X by following the One-Day and Periodic-D strategies, respectively. The purpose of doing so is to get a fair comparison: under the Periodic-DN strategy, we want to examine what the aim would be at any time, if we were standing in One-Day and Periodic-D strategies' shoes. Note that $\bar{\theta}^{O}$ and $\bar{\theta}^{PD}$ are completely determined given the current underlying and fund value, as well as the reference values.

Figure 3 Comparison of Aim $\bar{\theta}^*$

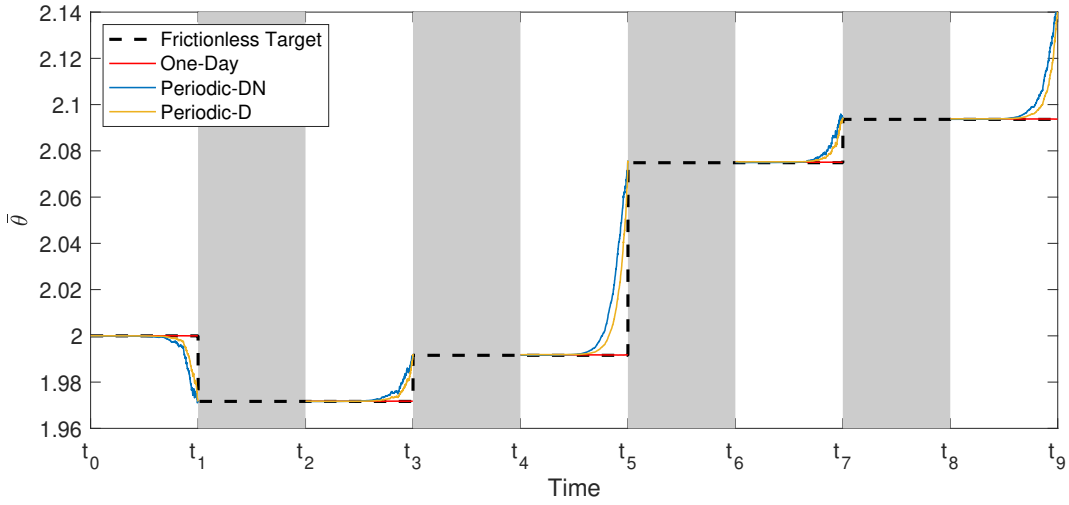

This figure illustrates the aims for the Periodic-DN strategy $\bar{\theta}^{PDN}$, the One-Day strategy $\bar{\theta}^{O}$, and the Periodic-D strategy $\bar{\theta}^{PD}$ over the simulated path given in Figure 1. The strategies are defined in Definition 2. This figure also includes the frictionless target calculated θ^{fl} defined in Theorem 1 for comparison. Note that the aims are only defined for the daytime. All of them are calculated under a simulated 5-day path for S and the corresponding X under the Periodic-DN strategy. By comparing with the curve of the frictionless target, it appears that the Periodic-DN strategy exhibits the strongest feature of aiming in front of target, while the One-Day strategy exhibits the least feature.

in the early hours, the Periodic-DN strategy mainly focuses on reaching the optimal position that minimizes today's slippage. However, $\bar{\theta}^{PDN}$ behaves dramatically differently from $\bar{\theta}^{O}$ in the late hours, when $\bar{\theta}^{PDN}$ moves away from today's frictionless target and towards tomorrow's target; in other words, the aim $\bar{\theta}^{PDN}$ is in front of the today's frictionless target. Furthermore, in stark contrast to the One-Day strategy whose aim is updated discontinuously, $\bar{\theta}^{PDN}$ is updated continuously and hits tomorrow's target exactly at market closing time. Therefore, as the market closing draws near, the Periodic-DN strategy starts to prepare for tomorrow by looking forward into the future and moving its aim continuously towards the position level that optimizes tomorrow's performance. It should be noted, however, that although the aim $\bar{\theta}^{PDN}$ hits tomorrow's target, the position level θ^{PDN} does not, as illustrated in Figure 1, since the position can only move at a finite speed due to the presence of market frictions. Nevertheless, θ^{PDN} is much closer to tomorrow's target than θ^{O}.

Finally, the target for Periodic-D strategy $\bar{\theta}^{PD}$ behaves similar to $\bar{\theta}^{PDN}$: it also starts the continuous update in the late hours and hit tomorrow's target exactly at market closing time. However, ignoring the overnight jump risk, $\bar{\theta}^{PD}$ moves towards tomorrow's target at a slower pace as compared to $\bar{\theta}^{PDN}$, which is shown as the difference between the blue and yellow curves in the second half of daytime of each day. Although the difference between $\bar{\theta}^{PD}$ and $\bar{\theta}^{PDN}$ does not appear to be large and eventually vanishes at market closing, it still results in a less optimal position at market closing and a larger exposure of the fund to the overnight jump risk as shown in Section 4 above, and a larger slippage (see Section 6.1 below).

The intraday transition of aim not only holds during the 5-day sample path in Figure 3, but also in general. To see this, we calculate ω_t such that $\bar{\theta}_t^{PDN} = (1-\omega_t)\bar{\theta}^O_{t_{2i-1}} + \omega_t \bar{\theta}^O_{t_{2i+1}}$, for any $t \in [t_{2i}, t_{2i+1}), i \geq 0$. In other words, ω is the weight of the aim between the current target and the future target; a value of ω close to 0 (resp. 1) means the aim is close to the current (resp. future) target. Figure 4 shows that the weight ω for both Periodic-DN and Periodic-D strategies transits from 0 to 1 continuously during daytime, consistent with the observation in Figure 3. In particular, the aim stays essentially on the current target during the first half of daytime, and moves towards the future target in the second half. This is a result of the tradeoff between the performance today and in the future. Indeed, since the current day's performance is only measured at market closing, moving the aim away from the current target too early would be detrimental to the performance, while moving too late would potentially cause a large deviation from tomorrow's target position and trigger a high friction cost. As a result, the transition in aim of the Periodic-DN strategy starts earlier than that of the Periodic-D strategy, by considering the overnight jump risk.

Figure 4 The weight between the current target and future target, ω, during daytime.

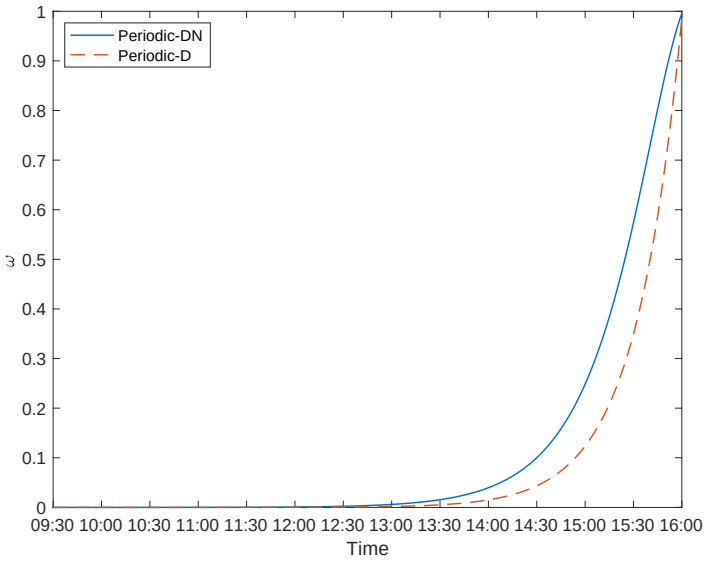

This figure shows the intraday change of the weight between the current target and future target, ω, averaged across days. We simulate 20000 sample paths based on the default parameters, each containing a 10-year period (2520 days including daytime and nighttime). On each sample path, the average ω at each time point during daytime is then calculated as the 1% trimmed mean over the 10-year period.

The intraday variation of weight ω is a distinct feature of our framework that was not observed in Gârleanu and Pedersen (2013, 2016), in which the weight of the aim remains constant. The fundamental reason is that in the Markowitz-type problem they considered, the performance monitoring and rebalancing occur at the same frequency: both are either at discrete time points or at any instant in the continuous-time version. Therefore, the current target in their model is also

continuously updated, resulting in the constant relative importance of the current target and future target, and hence, a constant weight. In contrast, our framework features discrete monitoring at daily market closing with a continuous-time rebalancing. As a result, the current target jumps at each market closing, a set of periodic and discrete time points. This results in the time-varying relative importance of current and future targets, and subsequently, a time-varying weight.

6. Statistics of Daily Slippage

In this section, we first extend the empirical findings on daily slippage in Tang and Xu (2013) by using the updated data till 2020. Then we show the overall slippage of different trading strategies and discuss the impact of market closure on slippage.

6.1. Daily Slippage

Due to the existence of market frictions, it is costly for LETFs to achieve the target leverage ratio exactly. Therefore, one can expect nonnegligible slippage. To check this, we calculate the empirical statistics of slippage (defined as in (4)) for the same set of twelve LETFs as in Tang and Xu (2013),[19] and update the time range to 2006 – 2020. For comparison, we also included 2011 – 2020, which is entirely uncovered in Tang and Xu (2013). Launched in 2006, these twelve LETFs were among the first-ever batch of LETFs introduced to the market. They have four underlying indices S&P500 (SPX), Dow Jones Industrial Average (INDU), NASDAQ 100 (NDX), and S&P MidCap 400 (MID), and for each index, there are one Bull LETF with multiple $\beta = +2$ and two Bear LETFs with multiples $\beta = -1, -2$. The results are reported in Panel A of Table 1, which confirms that the slippage is significant. For instance, during 2006 – 2020, the (-1x) fund on SPX has a mean slippage of 1.6917 bps. In comparison, the absolute values of daily target return (i.e., two times daily S&P 500) return have an average of 69.72 bps. This suggests that the slippage accounts for 1.94% of the daily target return, which cannot be ignored.

To quantify the slippage under Periodic-DN, Periodic-D, and One-Day strategies, we estimate the average daily slippage D_i via simulation. Consistent with Panel A, the (-2x) funds tend to have a much higher simulated slippage than the (2x) and (-1x) funds. Furthermore, the magnitude of slippage generated from the simulation result is in a similar range as that from historical data.

The Periodic-DN strategy results in the lowest slippage among three models for all four values of Λ and β, followed by the Periodic-D strategy, and the One-Day strategy has the highest slippage. For instance, for $\Lambda = 10^{-6}$ and (-2x) fund, the Periodic-DN strategy has an average slippage about

[19] Tang and Xu (2013) reported daily slippage using daily return calculated from LETFs market price, while here we do so using daily return from NAV. Indeed, the fund only has direct control over its NAV rather than its market price, and minimizing the slippage in terms of NAV is the goal of LETFs. However, as pointed out by Tang and Xu (2013), LETFs' total return deviation can be mainly attributed to NAV return deviation.

Table 1 Statistics for Slippage

Panel A: Historical Data

		2006-2020			2011-2020		
Index	Statistics	(2x)	(-1x)	(-2x)	(2x)	(-1x)	(-2x)
SPX	Mean	1.5624	1.6917	2.6820	1.4414	1.3503	2.3599
	Std dev	1.6581	1.6953	2.9280	1.6073	1.1008	2.5737
INDU	Mean	1.6429	1.7418	2.8077	1.5141	1.4084	2.4833
	Std dev	1.9110	1.8095	3.0932	1.7301	1.1693	2.4129
NDX	Mean	1.5846	1.7046	2.6196	1.4336	1.3381	2.1260
	Std dev	1.8105	1.6862	2.8148	1.7238	1.1794	2.3432
MID	Mean	1.7790	2.0298	2.7072	1.6594	1.8346	2.3208
	Std dev	2.3143	1.8698	3.3925	2.3639	1.5600	3.2562

Panel B: Simulation Results

		Periodic-DN			Periodic-D			One-Day		
Λ	Statistics	(2x)	(-1x)	(-2x)	(2x)	(-1x)	(-2x)	(2x)	(-1x)	(-2x)
10^{-6}	Mean	0.5590	0.8332	1.5702	0.6943	0.9512	2.0698	1.1037	1.3094	3.3194
	Std dev	0.2505	0.5818	1.6114	0.4634	0.6524	2.2348	0.9917	1.0394	3.1840
10^{-5}	Mean	0.7628	0.9568	2.1332	0.8194	1.0321	2.3934	1.2960	1.4195	3.7130
	Std dev	0.7171	0.6776	1.6809	0.9415	1.1309	2.4016	1.2503	1.7798	5.9288
10^{-4}	Mean	1.1308	1.2068	3.0859	1.1541	1.2460	3.2017	4.4654	2.1652	6.1326
	Std dev	1.1404	1.3398	3.7767	1.3626	1.3839	5.1612	2.8479	2.7922	7.2067
10^{-3}	Mean	1.9925	2.0507	5.8098	2.0027	2.0655	5.8542	5.6674	6.0283	18.8835
	Std dev	3.5890	3.6416	11.4236	1.7818	2.5431	8.7676	31.5373	28.1014	62.4603

Historical LETFs NAV returns are adjusted for dividends and the daily management fees. The slippage measure is defined as $D_i = |R_i^X - \beta R_i^S|$ and reported in bps. The statistics for simulation results are calculated over the daily slippage for 2520×20000 day (20000 sample paths, 2520 days per sample path), and the estimator is constructed as the 1% trimmed mean of the daily absolute deviations. The strategies are defined in Definition 2.

24% smaller than that of the Periodic-D strategy, and less than half of that of the Day-Only strategy. This indicates that aiming in front of target and accounting for the overnight risk indeed leads to a smaller slippage. In particular, the One-Day strategy's slippage is even higher compared to the Periodic-DN and Periodic-D strategy when Λ is larger. As discussed in Section 5, this can be attributed to the fact that the One-Day strategy has a myopic aim, so that it does not perform rebalancing in the latter half of daytime and leads to an inferior position at the next market opening. However, the amount of rebalancing after the next market opening is restricted by the presence of market frictions, especially for large Λ, causing a large slippage. Note that the difference between the slippage of the Periodic-D and Periodic-DN strategy diminishes as Λ becomes larger. Indeed, as will be shown in Section 7.2, when Λ is large (e.g., $\Lambda = 10^{-3}$), the U shape for the intraday trading volume of Periodic-DN strategy becomes flatter and less asymmetric as observed in Table 5. As a result, the fund puts less emphasis on the preparation for the overnight risk, and therefore the Periodic-DN and Periodic-D strategy have more similar rebalancing and slippage.

Table 1 demonstrates the advantage of Periodic-DN strategy over the Periodic-D and One-Day strategies using simulated underlying sample paths generated from the geometric Brownian motion model (1). However, does this advantage still exist in a real-world test? To this end, we repeat the slippage test in Panel B of Table 1 on the real intraday values of the four indices SPX, INDU, NDX, and MID. Specifically, for each index, we obtain the minutely values from September 17, 2020 to March 30, 2021. As a result, there are 132 trading days (excluding two days on which the

market closed earlier than usual), and 391 data points per day. On every trading day, we apply the three strategies on the intraday data and rebalance on a minutely basis. Note that since the closing index value on one day can differ from the opening value on the next day, the overnight price jump still exists, therefore we still expect that the Periodic-DN strategy outperforms the Periodic-D strategy. As shown in Table 2, the testing results show that the pattern in Panel B of Table 1 still holds. That is, for all four indices, Periodic-DN strategy generally leads to the lowest slippage, while the One-Day strategy leads to the highest slippage. This suggests that the benefit of aiming in front of target is still relevant in a real-world setting.[20]

6.2. Impact of Market Closure

Although the Periodic-DN strategy provides the optimal rebalancing strategy to better prepare for the market closure and overnight jump risk, the market closure can still be a contributing factor to the slippage under this strategy. To find out such cost, we compare the slippage presented in Panel B of Table 1 with a 17.5 hours market closure by following the Periodic-DN strategy, against the slippage in an otherwise identical market that opens for 24 hours a day (with the volatility $\sigma = 0.2$) by following the Day-Only strategy (which is the same as the optimal Periodic-DN strategy in this market as there is no market closure).

The result of this comparison is shown in Table 3. As expected, without market closure, the average slippage becomes lower across all Λ and β, suggesting that the presence of market closure is indeed costly to the fund and increases the slippage. Such reduction is most significant when Λ is small. For instance, with $\Lambda = 10^{-6}$, for the (2x) fund, the mean slippage reduces by 22% from 0.5590 bps to 0.4361 bps. However, for larger Λ, the reduction becomes less significant. This is because the larger frictions limit the fund's ability to rebalance during daytime, which makes the inability to rebalance during nighttime less constraining in comparison.

7. Empirical Findings and Implications

In this section, we discuss empirical findings and implications using our rebalancing model, including the weekend effect, intraday trading volume pattern, and compounding effect.

7.1. Weekend Effect

The long market closure before Monday, as it includes the weekend, can have a particular impact on the rebalancing[21]. Intuitively, our model suggests that the slippage will be largest on Monday,

[20] It should be noted that although Table 2 uses real index data instead of model-based simulated data (in Panel B of Table 1), it is still based on the market friction model in Section 2. As a result, the results may still be biased towards our model.

[21] In a different context Adrian et al. (2020) found that the market makers tend to speed up their inventory liquidation before the end of Friday.

Table 2 Statistics for Slippage: Real Index Data

SPX		Periodic-DN			Periodic-D			One-Day		
Λ	Statistics	(2x)	(-1x)	(-2x)	(2x)	(-1x)	(-2x)	(2x)	(-1x)	(-2x)
10^{-6}	Mean	0.3654	0.4124	1.0875	0.5034	0.5651	1.5506	0.9019	0.9580	2.7706
	Std dev	0.5260	0.4937	1.5211	0.6947	0.6629	2.0798	1.3043	1.3097	4.0467
10^{-5}	Mean	0.5498	0.5877	1.6329	0.6170	0.6745	1.8680	1.0383	1.0890	3.1715
	Std dev	0.8124	0.7868	2.4232	0.9078	0.8809	2.7414	1.5392	1.5601	4.7982
10^{-4}	Mean	0.8123	0.8322	2.3872	0.8432	0.8676	2.4920	1.4321	1.5058	4.4594
	Std dev	1.1891	1.1958	3.6399	1.2260	1.2344	3.7661	2.1041	2.1318	6.5470
10^{-3}	Mean	1.3015	1.3226	3.8783	1.3113	1.3334	3.9100	3.0570	3.5243	11.0877
	Std dev	1.7822	1.8208	5.5189	1.7982	1.8375	5.5707	2.8732	3.3166	10.4958
INDU		Periodic-DN			Periodic-D			One-Day		
Λ	Statistics	(2x)	(-1x)	(-2x)	(2x)	(-1x)	(-2x)	(2x)	(-1x)	(-2x)
10^{-6}	Mean	0.3687	0.4180	1.0888	0.5078	0.5387	1.4881	0.8972	0.9294	2.7095
	Std dev	0.6739	0.6247	1.9273	0.7097	0.6674	2.0845	1.2380	1.2169	3.7311
10^{-5}	Mean	0.5309	0.5673	1.5764	0.6091	0.6526	1.8419	0.9932	1.0533	3.0659
	Std dev	0.9372	0.8832	2.7302	0.9479	0.8970	2.7841	1.4831	1.4535	4.4740
10^{-4}	Mean	0.7494	0.7683	2.2060	0.7859	0.8084	2.3211	1.4303	1.4995	4.4730
	Std dev	1.2143	1.1889	3.6508	1.2367	1.2133	3.7345	2.2295	2.2588	6.9943
10^{-3}	Mean	1.2636	1.2642	3.7371	1.2744	1.2789	3.7804	3.1115	3.5568	11.1608
	Std dev	1.9035	1.9468	5.9640	1.9249	1.9669	6.0256	3.2244	3.4703	10.8910
NDX		Periodic-DN			Periodic-D			One-Day		
Λ	Statistics	(2x)	(-1x)	(-2x)	(2x)	(-1x)	(-2x)	(2x)	(-1x)	(-2x)
10^{-6}	Mean	0.6222	0.6833	1.9211	1.0449	1.0982	3.1821	1.9342	1.9776	5.8650
	Std dev	0.7968	0.8129	2.4495	1.5753	1.5951	4.8827	3.1340	3.1926	9.7768
10^{-5}	Mean	1.0538	1.0808	3.1362	1.2625	1.3023	3.7948	2.1339	2.1646	6.4374
	Std dev	1.4725	1.4858	4.4605	1.8680	1.8768	5.6847	3.3612	3.4246	10.4463
10^{-4}	Mean	1.5887	1.5865	4.6747	1.6471	1.6537	4.8735	2.7380	2.7687	8.2967
	Std dev	2.3181	2.3420	7.0119	2.4270	2.4506	7.3597	3.9347	3.9343	11.9107
10^{-3}	Mean	2.4024	2.3680	7.0371	2.4223	2.3927	7.1111	5.8675	6.5957	20.8493
	Std dev	3.5684	3.5462	10.5408	3.5969	3.5706	10.6150	7.4308	7.3282	22.7155
MID		Periodic-DN			Periodic-D			One-Day		
Λ	Statistics	(2x)	(-1x)	(-2x)	(2x)	(-1x)	(-2x)	(2x)	(-1x)	(-2x)
10^{-6}	Mean	0.4520	0.5105	1.4177	0.7708	0.7996	2.3238	1.5054	1.5233	4.5534
	Std dev	0.5555	0.5517	1.6456	0.8161	0.8064	2.4738	1.6891	1.6791	5.1004
10^{-5}	Mean	0.7412	0.7819	2.2658	0.9174	0.9464	2.7838	1.6482	1.6942	5.0633
	Std dev	0.9359	0.9290	2.8262	1.0295	1.0322	3.1520	1.8669	1.8723	5.7261
10^{-4}	Mean	1.2253	1.2223	3.6347	1.2965	1.2970	3.8609	2.5156	2.6480	8.0552
	Std dev	1.3860	1.4192	4.2927	1.4119	1.4448	4.3810	2.6038	2.7070	8.3896
10^{-3}	Mean	2.2765	2.2557	6.7799	2.2934	2.2761	6.8430	7.1193	9.6291	33.0852
	Std dev	2.4441	2.4877	7.5572	2.4615	2.5059	7.6145	5.9431	8.2481	28.6164

This table shows the slippage produced by Periodic-DN, Periodic-D and One-Day strategies based on real intraday index data rather than simulated index data. The slippage measure is defined as $D_i = |R_i^X - \beta R_i^S|$ and reported in bps. The mean and standard deviation are calculated from the 132 daily values from September 17, 2020 to March 30, 2021. The three strategies are applied on the minutely data of the underlying index over these 132 days.

as the long market closure during weekend increases the overnight jump risk. We call this weekend effect.

We first perform a simulation test to estimate the model slippage with a long market closure, for which we assume that a daily trading hour of 2.17 hours and market closure hours of 21.83 hours. This ratio matches the scenario from Friday market opening to Monday market opening, where there is a 6.5-hour trading period followed by a (17.5+48)-hour market closure.[22] By comparing

[22] In our model framework, the fund performance is monitored on a daily basis, and the slippage is measured via the daily fund NAV and the underlying returns from one market closing time to the next. Since the slippage during a weekend is measured via the fund NAV and the underlying return from the Friday market closing time to the Monday market closing time, it is natural to rescale the weekend into a "day" to be consistent with our model. Ideally the weekend should be treated differently from the remaining days without scaling; however, this ideal treatment would bring extra complexity to the model framework.

Table 3 Cost of Market Closure

		Long Market Closure			Normal Market Closure			No Market Closure		
Λ	Statistics	(2x)	(-1x)	(-2x)	(2x)	(-1x)	(-2x)	(2x)	(-1x)	(-2x)
10^{-6}	Mean	0.6178	0.8623	1.7074	0.5590	0.8332	1.5702	0.4361	0.7963	1.2853
	Std dev	0.5414	0.6648	1.7117	0.2505	0.5818	1.6114	0.1929	0.3271	0.4671
10^{-5}	Mean	0.8863	1.0366	2.4351	0.7628	0.9568	2.1332	0.5492	0.8352	1.5714
	Std dev	1.0452	1.0209	2.6054	0.7171	0.6776	1.6809	0.4302	0.4636	1.4761
10^{-4}	Mean	1.4290	1.5069	4.0712	1.1308	1.2068	3.0859	0.7972	0.9674	2.1760
	Std dev	1.6663	1.9474	4.5514	1.1404	1.3398	3.7767	0.5489	0.8965	2.9398
10^{-3}	Mean	2.6988	2.7629	8.0162	1.9925	2.0507	5.8098	1.3686	1.4333	3.8467
	Std dev	2.5767	4.0329	28.8919	3.5890	3.6416	11.4236	1.9423	1.6857	4.0751

The slippage measure is defined as $D_i = |R_i^X - \beta R_i^S|$ and reported in bps. The statistics for simulation results are calculated over the daily slippage for 2520×20000 day (20000 sample paths, 2520 days per sample path), and the estimator is constructed as the 1% trimmed mean of the daily absolute deviations. No Market Closure means 24-hour trading, Normal Market Closure means 6.5-hour trading and 17.5-hour market closure, and Long Market Closure means 2.17-hour trading and 21.83-hour market closure. The daily average $\sigma = 0.2$ in all three cases.

the slippages in the Long Market Closure columns and the Normal Market Closure columns in Table 3, our model predicts that the slippage should also be more prominent on Monday, as the fund will not be able to rebalance during the large proportion of market closure.

Next, we test the weekend effect empirically. To this end, we redo the slippage calculation in Panel A of Table 1 for 2006 – 2020 by grouping them on the day of the week. The results are shown in Table 4. For all 12 LETFs, the mean slippage on Monday is always the highest among all days of the week. Furthermore, the slippage on Monday is 23% to 33% higher than the average of remaining days for all LETFs. Therefore, the empirical results support the implication of the weekend effect from our model.

Table 4 Statistics for Historical Slippage Grouped on the Days of the Week

Index	Multiple	Statistics	Monday	Tuesday	Wednesday	Thursday	Friday
SPX	(2x)	Mean	1.8883	1.4806	1.3794	1.5509	1.5385
		Std dev	2.0741	1.6002	1.2943	1.5441	1.6803
	(-1x)	Mean	2.0368	1.5724	1.6107	1.6464	1.6175
		Std dev	2.3450	1.8130	1.4712	1.3370	1.2954
	(-2x)	Mean	3.2053	2.4861	2.5589	2.5508	2.6483
		Std dev	4.0658	2.9198	2.3393	2.2615	2.7407
INDU	(2x)	Mean	2.0597	1.5171	1.5227	1.6463	1.4989
		Std dev	2.2424	2.0524	1.6306	1.8397	1.6954
	(-1x)	Mean	2.1050	1.6379	1.6238	1.7639	1.6051
		Std dev	2.5846	1.7282	1.5834	1.6663	1.2247
	(-2x)	Mean	3.3248	2.6025	2.6729	2.8084	2.6684
		Std dev	4.0078	3.3140	2.5565	2.8942	2.4565
NDX	(2x)	Mean	1.8772	1.4737	1.4868	1.5926	1.5151
		Std dev	1.9163	1.9057	1.6145	1.7266	1.8549
	(-1x)	Mean	2.0534	1.5774	1.5997	1.6557	1.6631
		Std dev	2.3910	1.6385	1.4576	1.4291	1.3093
	(-2x)	Mean	3.2119	2.4287	2.4812	2.6331	2.3857
		Std dev	3.8823	2.7625	2.4076	2.4484	2.2843
MID	(2x)	Mean	2.1454	1.6006	1.7329	1.6869	1.7567
		Std dev	2.5895	1.7788	2.6767	2.0334	2.3592
	(-1x)	Mean	2.4490	1.9455	1.9124	1.9372	1.9351
		Std dev	2.4577	1.8259	1.6758	1.5145	1.7314
	(-2x)	Mean	3.3692	2.5688	2.3804	2.6562	2.6123
		Std dev	4.3976	3.4666	2.7779	2.6760	3.3758

The slippage measure is defined as $D_i = |R_i^X - \beta R_i^S|$ and reported in bps. The data period is 2006-2020 as in Panel A of Table 1, but the slippage is divided into five groups based on the day of the week. This table indicates the "weekend" effect that the slippage tends to be the largest on Monday.

7.2. Intraday Trading Volume

To investigate the intraday trading pattern, we compare strategies with three overnight risk profiles: *high overnight risk* (the benchmark case with $\delta = 17.5/6.5, k = 3$), *low overnight risk* (with $\delta = 17.5/6.5, k = 5$) and *no overnight risk* ($\delta = 0$). By taking a larger value of k, the underlying nighttime volatility is smaller compared to daytime, which leads to smaller overnight jumps on average.

Figure 5 Intraday Absolute Trading Speed Distribution

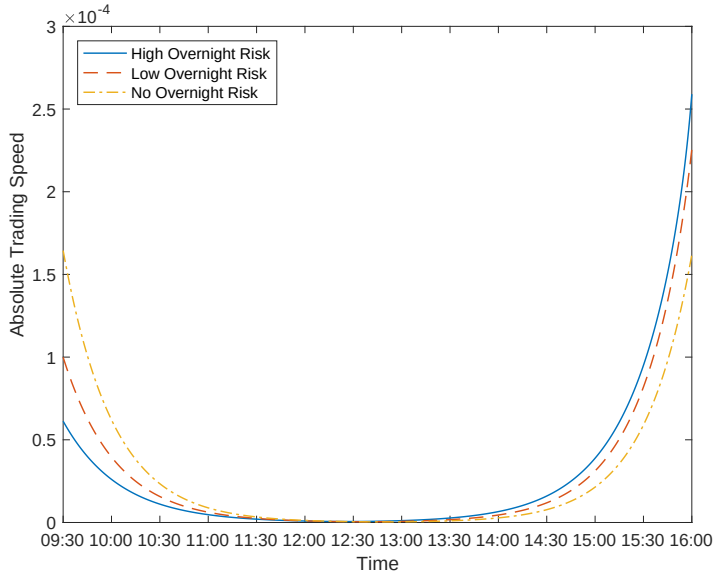

This figure shows the intraday distribution of the absolute trading speed $|\varphi_u|$ for three overnight risk profiles. The distribution is estimated via Monte Carlo simulation. We simulate 20000 sample paths based on the default parameters, each containing a 10-year period (2520 days including daytime and nighttime). We calculate the absolute trading speed over daytime of each day, and then calculate the average across all days on all sample paths.

Figure 5 shows under all three profiles, the intraday absolute trading speed exhibits the well-known U-shaped curve (see, for instance, Admati and Pfleiderer (1988), Dai et al. (2015)). When there is no overnight risk, the curve is symmetric, that is, the absolute trading speed after the market opening is about the same as that before market closing. In the benchmark case with high overnight risk, the trading speed is much higher before market closing, and this bias is less significant for the low overnight risk case. In other words, a larger portion of rebalancing occurs right before market closing when the overnight risk is greater. This is consistent with the observation based on one sample path in Section 4, reflecting the effort of accounting for overnight jump risk.

This implies that the fund's intraday trading activity is at a high level before market closing, a medium level after market opening, and at a low level in the middle of daytime. This conjecture may be tested in future empirical studies using high-frequency intraday data on LETF rebalancing activity, although the data set does not seem to be available now.

Table 5 Intraday Trading Volume

From To	9:30 10:00	10:00 10:30	10:30 11:00	11:00 11:30	11:30 12:00	12:00 12:30	12:30 13:00	13:00 13:30	13:30 14:00	14:00 14:30	14:30 15:00	15:00 15:30	15:30 16:00	15:55 16:00	Whole Day	Market Frictions
Periodic-DN Strategy																
$\Lambda = 10^{-6}$	0.0032	0.0014	0.0006	0.0002	0.0001	0.0001	0.0001	0.0001	0.0003	0.0008	0.0020	0.0049	0.0122	0.0031	0.0260	0.0005
$\Lambda = 10^{-5}$	0.0024	0.0018	0.0014	0.0011	0.0009	0.0008	0.0009	0.0011	0.0014	0.0018	0.0025	0.0034	0.0046	0.0009	0.0242	0.0014
$\Lambda = 10^{-4}$	0.0014	0.0013	0.0013	0.0012	0.0012	0.0012	0.0012	0.0012	0.0013	0.0014	0.0015	0.0016	0.0017	0.0003	0.0175	0.0060
$\Lambda = 10^{-3}$	0.0008	0.0008	0.0008	0.0008	0.0008	0.0008	0.0008	0.0008	0.0008	0.0008	0.0008	0.0009	0.0009	0.0001	0.0108	0.0205
Periodic-D Strategy																
$\Lambda = 10^{-6}$	0.0080	0.0030	0.0011	0.0004	0.0002	0.0001	0.0000	0.0001	0.0001	0.0004	0.0011	0.0029	0.0078	0.0019	0.0252	0.0004
$\Lambda = 10^{-5}$	0.0036	0.0027	0.0020	0.0015	0.0011	0.0009	0.0008	0.0008	0.0010	0.0013	0.0019	0.0026	0.0035	0.0007	0.0236	0.0014
$\Lambda = 10^{-4}$	0.0016	0.0015	0.0014	0.0013	0.0012	0.0012	0.0012	0.0012	0.0012	0.0012	0.0013	0.0014	0.0016	0.0003	0.0172	0.0059
$\Lambda = 10^{-3}$	0.0009	0.0009	0.0008	0.0008	0.0008	0.0008	0.0008	0.0008	0.0008	0.0008	0.0008	0.0008	0.0008	0.0001	0.0107	0.0190
One-Day Strategy																
$\Lambda = 10^{-6}$	0.0142	0.0060	0.0026	0.0011	0.0005	0.0002	0.0001	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0248	0.0007
$\Lambda = 10^{-5}$	0.0058	0.0045	0.0034	0.0026	0.0020	0.0015	0.0011	0.0008	0.0006	0.0004	0.0003	0.0002	0.0001	0.0000	0.0233	0.0019
$\Lambda = 10^{-4}$	0.0020	0.0018	0.0016	0.0014	0.0012	0.0011	0.0009	0.0008	0.0006	0.0005	0.0003	0.0002	0.0001	0.0000	0.0124	0.0036
$\Lambda = 10^{-3}$	0.0006	0.0006	0.0005	0.0005	0.0004	0.0004	0.0003	0.0003	0.0002	0.0002	0.0001	0.0001	0.0000	0.0000	0.0042	0.0044
One-Shot Strategy																
	0	0	0	0	0	0	0	0	0	0	0	0	0.0246	0.0246	0.0246	$+\infty$

This table compares the intraday trading volume for the Periodic-DN strategy, Periodic-D strategy and One-Day strategy (with $\Lambda = 10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}$) and the one-shot rebalancing strategy (the strategy ignoring the market frictions and assuming $r = 0$, as given in Theorem 1). The unit of the trading volume is the number of shares. The market friction is daily average, with unit 10^{-4} squared dollars. The statistics are calculated using the trading volume during each of the 13 half-hour periods for 2520×20000 days (20000 sample paths, 2520 days per sample path). Initial values: $s = x = 1$, $z = 2$.

Next, we compare the intraday trading volume of the Periodic-DN, Periodic-D, and One-Day strategies. We split the daytime on each trading day from market opening time 9:30 to market closing time 16:00 into 13 subintervals, each with 30 minutes. Then we calculate the total expected trading volume for each subinterval, the trading volume during the last 5 minutes, as well as total trading volume via simulation, as reported in Table 5. As a comparison, we also include the one-shot strategy that only performs a single lump-sum trade at the market closing each day.

For each Λ, the intraday trading volume exhibits the same U-shaped characteristics as in Figure 5. For each strategy, a larger Λ results in a smaller daily total trading volume, because larger market friction discourages rebalancing. On the other hand, a larger Λ also results in a flatter U-shape; in particular, for $\Lambda = 10^{-3}$ the intraday volume is almost constant, ranging from 0.0009 to 0.0010. Indeed, given a daily total trading volume, spreading the rebalancing more evenly across daytime leads to a smaller (quadratic) friction cost.

Compared to the One-Shot trading strategy (the strategy ignoring market frictions and assuming $r = 0$ given in Theorem 1), the Periodic-DN strategy has a significantly smaller trading volume near market closing. For instance, the Periodic-DN strategy for $\Lambda = 10^{-6}$ has a trading volume 0.0031 during the last five minutes (as opposed to 0.0246 for the one-shot strategy), implying a much smaller market friction cost right before market closing. On the other hand, Periodic-DN strategy has a higher volume right before market closing compared to the Periodic-D and One-Day strategy, due to accounting for overnight risk and aiming in front of target as discussed earlier.

On a daily level, the total trading volume and the market friction costs from the Periodic-DN and Periodic-D strategies are on the same level; however, both are greater than those from the One-Day strategy, especially when Λ is large. This suggests that the One-Day strategy does not rebalance sufficiently. Indeed, following the myopic aim, the One-Day strategy barely does any rebalancing in the second half of daytime. Therefore, at the beginning of the next day, the position is way off target, and the fund cannot do enough rebalancing in a limited amount of time due to the market frictions. As a result, under the myopic One-Day strategy, the fund ends up with insufficient rebalancing and large slippage, as discussed in Section 6.1. This is especially the case for large Λ, since the amount of rebalancing after the market opening needs to be cut down even further to keep a reasonable level of cost.

8. Conclusion

We study how to rebalance leveraged ETFs in a comprehensive setting, including overnight market closure and market frictions, and obtain analytical solutions in terms of a system of periodic ordinary differential equations. Interestingly, the optimal rebalancing (hedging) strategy confirms the principle of "aiming in front of target" in Gârleanu and Pedersen (2013), introduced in the setting of asset allocation, although the focus here is on hedging. Our optimal strategy yields a lower slippage and smoother trading pattern, compared to existing strategies. The framework in this paper may be extended, in principle, to study more general hedging problems with market frictions and market closure. The challenge is to get analytical solutions for specific hedging problems, although we expect that numerical procedures suggested in the current framework may still work.

We present empirical findings and implications for the weekend effect and return deviations during multi-day periods. The discussion about the intraday trading volume leads to the following conjecture from the optimal strategy of our model: LETFs' daily rebalancing activity will be at a high level before market closing, at a medium level after the market opening, and at a low level in the middle of daytime. This conjecture may be tested in future empirical studies using high-frequency intraday data on LETF rebalancing activity, although the data set does not seem available now.

Acknowledgments

Chen Yang acknowledges the support from CUHK Direct Grant No. 4055132 and a CUHK University Startup Grant. Dai is on leave from the National University of Singapore, where part of this work was done, and he acknowledges the support from the National Natural Science Foundation of China (NSFC) [Grant No. 12071333], the Hong Kong Polytechnic University [Grant No. P0039114], and the Singapore Ministry of Education [Grant Nos R-146-000-243/306/311-114 and R-703-000-032-112].

References

Acharya, V. V. and Pedersen, L. H. (2005). Asset pricing with liquidity risk. *Journal of Financial Economics*, 77(2):375–410.

Admati, A. R. and Pfleiderer, P. (1988). A theory of intraday patterns: Volume and price variability. *The Review of Financial Studies*, 1(1):3–40.

Adrian, T., Capponi, A., Fleming, M., Vogt, E., and Zhang, H. (2020). Intraday market making with overnight inventory costs. *Journal of Financial Markets*, 50:100564.

Amann, H. (1990). *Ordinary Differential Equations*. An Introduction to Nonlinear Analysis. DE GRUYTER, Berlin, New York.

Avellaneda, M. and Dobi, D. (2012). Structural slippage of leveraged ETFs. Available at SSRN 2127738.

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

Bai, Q., Bond, S. A., and Hatch, B. C. (2012). The impact of leveraged and inverse ETFs on underlying stock returns. In *46th Annual AREUEA Conference Paper*.

Bertsimas, D. and Lo, A. W. (1998). Optimal control of execution costs. *Journal of Financial Markets*, 1(1):1–50.

Breen, W. J., Hodrick, L. S., and Korajczyk, R. A. (2002). Predicting equity liquidity. *Management Science*, 48(4):470–483.

Brunnermeier, M. K. and Pedersen, L. H. (2005). Predatory trading. *The Journal of Finance*, 60(4):1825–1863.

Cartea, Á., Jaimungal, S., and Penalva, J. (2015). *Algorithmic and High-Frequency Trading*. Cambridge University Press.

Charupat, N. and Miu, P. (2016). *Leveraged Exchange-Traded Funds*. Springer.

Cheng, M. and Madhavan, A. (2009). The dynamics of leveraged and inverse exchange-traded funds. *Journal of Investment Management*, 16(4):43.

Cont, R., Kukanov, A., and Stoikov, S. (2014). The price impact of order book events. *Journal of financial econometrics*, 12(1):47–88.

Dai, M., Li, P., Liu, H., and Wang, Y. (2015). Portfolio choice with market closure and implications for liquidity premia. *Management Science*, 62(2):368–386.

Frei, C. and Westray, N. (2015). Optimal execution of a vwap order: a stochastic control approach. *Mathematical Finance*, 25(3):612–639.

Gârleanu, N. and Pedersen, L. H. (2013). Dynamic trading with predictable returns and transaction costs. *The Journal of Finance*, 68(6):2309–2340.

Gârleanu, N. and Pedersen, L. H. (2016). Dynamic portfolio choice with frictions. *Journal of Economic Theory*, 165:487–516.

Greenwood, R. (2005). Short-and long-term demand curves for stocks: theory and evidence on the dynamics of arbitrage. *Journal of Financial Economics*, 75(3):607–649.

Grossman, S. J. and Miller, M. H. (1988). Liquidity and market structure. *the Journal of Finance*, 43(3):617–633.

Guasoni, P. and Mayerhofer, E. (2017). Leveraged funds: Robust replication and performance evaluation. *Michael J. Brennan Irish Finance Working Paper Series Research Paper 19-1*.

Hansen, L. P. and Sargent, T. J. (2013). *Recursive Models of Dynamic Linear Economies*. Princeton University Press.

Haugh, M. B. (2011). A note on constant proportion trading strategies. *Operations Research Letters*, 39(3):172–179.

Henderson, B. J. and Buetow, G. W. (2014). The performance of leveraged and inverse leveraged exchange traded funds. *The Journal of Investment Management*, 12(1):69–92.

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

Keppo, J., Reppen, M., and Soner, H. M. (2021). Discrete dividend payments in continuous time. *Mathematics of Operations Research*, 46(3):895–911.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335.

Lockwood, L. J. and Linn, S. C. (1990). An examination of stock market return volatility during overnight and intraday periods, 1964–1989. *The Journal of Finance*, 45(2):591–601.

Moreau, L., Muhle-Karbe, J., and Soner, H. M. (2017). Trading with small price impact. *Mathematical Finance*, 27(2):350–400.

Obizhaeva, A. A. and Wang, J. (2013). Optimal trading strategy and supply/demand dynamics. *Journal of Financial Markets*, 16(1):1–32.

Pham, H. (2010). *Continuous-time Stochastic Control and Optimization with Financial Applications (Stochastic Modelling and Applied Probability)*. Springer.

Reid (1972). *Riccati Differential Equations*. Academic Press.

Robert, E., Robert, F., and Jeffrey, R. (2012). Measuring and modeling execution cost and risk. *The Journal of Portfolio Management*, 38(2):14–28.

Rogers, L. C. and Singh, S. (2010). The cost of illiquidity and its effects on hedging. *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics*, 20(4):597–615.

Shum, P. M. and Kang, J. (2013). Leveraged and inverse ETF performance during the financial crisis. *Managerial Finance*.

Stoll, H. R. and Whaley, R. E. (1990). Stock market structure and volatility. *Review of Financial Studies*, 3(1):37–71.

Tang, H. and Xu, X. E. (2013). Solving the return deviation conundrum of leveraged exchange-traded funds. *Journal of Financial and Quantitative Analysis*, 48(1):309–342.

Tuzun, T. (2014). Are leveraged and inverse etfs the new portfolio insurers? In *Paris December 2014 Finance Meeting EUROFIDAI-AFFI Paper*.

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

Online Supplement

Leveraged ETFs with Market Closure and Frictions

A. Derivation of the Explicit Value Function without Overnight Market Closure

For simplicity, we shall begin with the case without overnight market closure. We shall study the case with overnight market closure in Section D.

We take $t_k = kT,\ k = 0, 1, \cdots$ as the end of k-th day (and also the beginning of $(k+1)$-th day). The value function is

$$V(t, s, x, z; \bar{s}, \bar{x}) = \inf_{\varphi \in \mathcal{A}} J(t, s, x, z; \varphi, \bar{s}, \bar{x}), \tag{A.1}$$

where the cost functional is

$$\begin{aligned} &J(t, s, x, z; \varphi, \bar{s}, \bar{x}) \\ &= E\left[\frac{1}{2}\bar{x}^2\left(\beta\cdot\left(\frac{S_{t_1}}{\bar{s}}-1\right)-\left(\frac{X_{t_1-}}{\bar{x}}-1\right)\right)^2+\int_t^{t_1}\frac{1}{2}\Lambda S_u^2\varphi_u^2 du\right. \\ &\left.+\sum_{i=1}^{\infty}e^{-\rho i}\left(\frac{1}{2}X_{t_i}^2\left(\beta\cdot\left(\frac{S_{t_{i+1}}}{S_{t_i}}-1\right)-\left(\frac{X_{t_{i+1}-}}{X_{t_i}}-1\right)\right)^2+\int_{t_i}^{t_{i+1}}\frac{1}{2}\Lambda S_u^2\varphi_u^2 du\right)\right], \end{aligned} \tag{A.2}$$

and the set of admissible strategies is

$$\mathcal{A} = \left\{\varphi : E\left[\sum_{i=1}^{\infty} e^{-\rho t_i}\int_{t_i}^{t_{i+1}}\varphi_u^2 S_u^2 du\right] < \infty\right\}.$$

A similar discrete-continuous formulation is given in Keppo et al. (2021) in a different context.

From the dynamic programming principle and the time-homogeneous Markov property, for $0 \le t < T$, we have the following recursive formula

$$\begin{aligned} V(t, s, x, z; \bar{s}, \bar{x}) = \inf_{\varphi \in \mathcal{A}_1} E\left[\frac{1}{2}\bar{x}^2\left(\beta\cdot\left(\frac{S_T}{\bar{s}}-1\right)-\left(\frac{X_{T-}}{\bar{x}}-1\right)\right)^2\right. \\ \left.+\int_t^T\frac{1}{2}\Lambda S_u^2\varphi_u^2 du + e^{-\rho T}V(0, S_T, X_T, \theta_T; S_T, X_T)\right], \end{aligned}$$

where

$$\mathcal{A}_1 = \left\{\varphi : E\left[\int_0^T \varphi_u^2 S_u^2 du\right] < \infty\right\}.$$

And in particular, the following connection condition holds thanks to the homotheticity

$$\begin{aligned} &V(T-, s, x, z; \bar{s}, \bar{x}) \\ &= \frac{1}{2}\bar{x}^2\left(\beta\cdot\left(\frac{s}{\bar{s}}-1\right)-\left(\frac{x}{\bar{x}}-1\right)\right)^2 + e^{-\rho T}V(0, s, (1-\gamma)x, z; s, (1-\gamma)x) \\ &= \frac{1}{2}\bar{x}^2\left(\beta\cdot\left(\frac{s}{\bar{s}}-1\right)-\left(\frac{x}{\bar{x}}-1\right)\right)^2 + e^{-\rho T}(1-\gamma)^2 x^2 V\left(0, 1, 1, \frac{s}{(1-\gamma)x}z; 1, 1\right), \end{aligned} \tag{A.3}$$

This iterative formula is useful for giving the terminal condition for the HJB equation, as we will see below.

Denote

$$\begin{aligned}\mathcal{L}^{\varphi}U = &\frac{1}{2}z^2\sigma^2s^2\frac{\partial^2U}{\partial x^2}+\frac{1}{2}\sigma^2s^2\frac{\partial^2U}{\partial s^2}+z\sigma^2s^2\frac{\partial^2U}{\partial x\partial s}\\ &+((\mu-r)zs+rx)\frac{\partial U}{\partial x}+\mu s\frac{\partial U}{\partial s}+\varphi\frac{\partial U}{\partial z}.\end{aligned}$$

The HJB equation corresponding to $V(t,s,x,z;\bar{s},\bar{x})$ is given as

$$\frac{\partial U}{\partial t}+\inf_{\varphi\in\mathbb{R}}\left[\mathcal{L}^{\varphi}U+\frac{1}{2}\Lambda s^2\varphi^2\right]=0, \tag{A.4}$$

in $\mathcal{D}=\{(t,s,x,z):t\in[0,T),s\in\mathbb{R}^+,x,z\in\mathbb{R}\}$ with

$$\begin{aligned}U(T,s,x,z;\bar{s},\bar{x}) = &\frac{1}{2}\bar{x}^2\left(\beta\cdot\left(\frac{s}{\bar{s}}-1\right)-\left(\frac{x}{\bar{x}}-1\right)\right)^2\\ &+e^{-\rho T}(1-\gamma)^2x^2U(0,1,1,\frac{s}{(1-\gamma)x}z;1,1).\end{aligned} \tag{A.5}$$

The optimal strategy φ^* can be solved uniquely via the first order condition as

$$\varphi^*=-\frac{1}{\Lambda s^2}\frac{\partial U}{\partial z}, \tag{A.6}$$

and (A.4) can be equivalently written as

$$\begin{aligned}&\frac{\partial U}{\partial t}+\frac{1}{2}z^2\sigma^2s^2\frac{\partial^2U}{\partial x^2}+\frac{1}{2}\sigma^2s^2\frac{\partial^2U}{\partial s^2}+z\sigma^2s^2\frac{\partial^2U}{\partial x\partial s}\\ &\quad+((\mu-r)zs+rx)\frac{\partial U}{\partial x}+\mu s\frac{\partial U}{\partial s}-\frac{1}{2\Lambda s^2}\left(\frac{\partial U}{\partial z}\right)^2=0.\end{aligned} \tag{A.7}$$

A.1. Homotheticity of the Value Function

Since the value function involves time and five state variables, it is extremely difficult to calculate due to the curse of dimensionality. Fortunately, our formulation of the cost functional allows the following homotheticity, which significantly simplifies the problem and makes it analytically solvable.

PROPOSITION 1. *We have the following homotheticity:* $\forall\eta>0,\alpha\in\mathbb{R}$,

1. $V(t,\eta s,\eta x,z,\eta\bar{s},\eta\bar{x})=\eta^2V(t,s,x,z,\bar{s},\bar{x})$.
2. $V(t,s,\alpha x,\alpha z,\bar{s},\alpha\bar{x})=\alpha^2V(t,s,x,z,\bar{s},\bar{x})$.

Here, the first equation in Proposition 1 says that proportionally increasing stock price and fund value does not change the optimal number of shares or the optimal strategy. Intuitively, denominating the stock price and fund value in another currency does not change the problem. The second equation says that a proportionally increasing number of shares implies a proportional

increase in fund value, as well as the optimal strategy. Intuitively, a fund with twice the NAV needs to hold and trade twice the amount of stock. Combining these two properties, we have

$$V(t,s,x,z;\bar{s},\bar{x}) = \bar{x}^2 V\left(t, \frac{s}{\bar{s}}, \frac{x}{\bar{x}}, \frac{\bar{s}}{\bar{x}}z; 1, 1\right), \quad \forall \bar{s} > 0, \bar{x} \in \mathbb{R}.$$

Therefore, in the following we can assume without loss of generality that $\bar{x} = \bar{s} = 1$, and denote for short that $V(t,s,x,z) = V(t,s,x,z;1,1)$.

[Proof of Proposition 1] For part (1), equations (1) – (3) and (9) show that for any $\varphi \in \mathcal{A}$, $X_u^{t,\eta s,\eta x,z,\varphi} = \eta X_u^{t,s,x,z,\varphi}$, $S_u^{t,\eta s,\eta x,z,\varphi} = \eta S_u^{t,s,x,z,\varphi}$, and $\theta_u^{t,\eta s,\eta_x,z,\varphi} = \theta_u^{t,s,x,z,\varphi}$. Therefore, (A.2) shows that $\forall \varphi \in \mathcal{A}$,

$$J(t,\eta s,\eta x,z,\eta\bar{s},\eta\bar{x};\varphi) = \eta^2 J(t,s,x,z,\bar{s},\bar{x};\varphi),$$

from which the conclusion follows. For part (2), equations (1) – (3) and (9) show that for any $\varphi \in \mathcal{A}$, $X_u^{t,s,\eta x,\eta z,\eta\varphi} = \eta X_u^{t,s,x,z,\varphi}$, $S_u^{t,s,\eta x,\eta z,\eta\varphi} = S_u^{t,s,x,z,\varphi}$, and $\theta_u^{t,s,\eta z,\eta\varphi} = \eta\theta_u^{t,x,z,\varphi}$. Therefore, $\eta\varphi \in \mathcal{A}$, and (A.2) shows that $\forall \varphi \in \mathcal{A}$,

$$J(t,\eta s,\eta x,z,\eta\bar{s},\eta\bar{x};\varphi) = \eta^2 J(t,s,x,z,\bar{s},\bar{x};\varphi),$$

which again leads to the conclusion. Q.E.D.

A.2. Solving the HJB Equation

Next we focus on the equation for $V(t,s,x,z) := V(t,s,x,z;1,1)$ and $U(t,s,x,z) := U(t,s,x,z;1,1)$. Based on the quadratic form of HJB equation and terminal condition, we guess that

$$\begin{aligned} U(t,s,x,z) = a(t)x^2 + (b_0(t) + b_1(t)z)xs + (c_0(t) + c_1(t)z + c_2(t)z^2)s^2 \\ + d(t)x + (e_0(t) + e_1(t)z)s + f(t). \end{aligned} \tag{A.8}$$

By plugging (A.8) into the HJB equation (A.7), the coefficients satisfy the following nonlinear ODE system

$$\frac{\partial a}{\partial t} = \frac{1}{2\Lambda}(b_1)^2 - 2ra \tag{A.9}$$

$$\frac{\partial b_0}{\partial t} = \frac{1}{\Lambda}b_1c_1 - (\mu + r)b_0 \tag{A.10}$$

$$\frac{\partial b_1}{\partial t} = \frac{1}{\Lambda}2b_1c_2 - 2(\mu - r)a - (\mu + r)b_1 \tag{A.11}$$

$$\frac{\partial c_0}{\partial t} = -(\sigma^2 + 2\mu)c_0 + \frac{1}{2\Lambda}(c_1)^2 \tag{A.12}$$

$$\frac{\partial c_1}{\partial t} = -(\sigma^2 + 2\mu)c_1 - (\sigma^2 + \mu - r)b_0 + \frac{1}{\Lambda}2c_1c_2 \tag{A.13}$$

$$\frac{\partial c_2}{\partial t} = -\sigma^2 a - (\sigma^2 + 2\mu)c_2 - (\sigma^2 + \mu - r)b_1 + \frac{1}{\Lambda}2(c_2)^2 \tag{A.14}$$

$$\frac{\partial d}{\partial t} = \frac{1}{\Lambda}b_1e_1 - rd \tag{A.15}$$

$$\frac{\partial e_0}{\partial t} = -\mu e_0 + \frac{1}{\Lambda} c_1 e_1 \tag{A.16}$$
$$\frac{\partial e_1}{\partial t} = -(\mu - r)d - \mu e_1 + \frac{1}{\Lambda} 2c_2 e_1 \tag{A.17}$$
$$\frac{\partial f}{\partial t} = \frac{1}{2\Lambda}(e_1)^2. \tag{A.18}$$

Note that

$$\begin{aligned} x^2 U(0,1,1,\frac{s}{x}z;1,1) = (a(0)+b_0(0)+c_0(0)+d(0)+e_0(0)+f(0))x^2 \\ + (b_1(0)+c_1(0)+e_1(0))sxz + c_2(0)s^2z^2. \end{aligned}$$

Therefore, the terminal condition for the above ODE system is given as

$$\begin{aligned} a(T) &= \frac{1}{2} + e^{-\rho T}(1-\gamma)^2(a(0)+b_0(0)+c_0(0)+d(0)+e_0(0)+f(0)) \\ b_0(T) &= -\beta, \quad b_1(T) = e^{-\rho T}(1-\gamma)(b_1(0)+c_1(0)+e_1(0)) \\ c_0(T) &= \frac{1}{2}\beta^2, \quad c_1(T) = 0, \quad c_2(T) = e^{-\rho T}c_2(0), \quad d(T) = \beta - 1 \\ e_0(T) &= \beta(1-\beta), \quad e_1(T) = 0, \quad f(T) = \frac{1}{2}(\beta-1)^2. \end{aligned} \tag{A.19}$$

Note that the terminal condition depends on the time-0 value of the solution, which will be determined via a fixed point iteration.

From (A.6) and (A.8), the optimal strategy φ^* is given as

$$\varphi^*(t,s,x,z;\bar{s},\bar{x}) = -\frac{1}{\Lambda}\left(b_1(t)\frac{x}{s} + c_1(t)\frac{\bar{x}}{\bar{s}} + 2c_2(t)z + e_1(t)\frac{\bar{x}}{s}\right). \tag{A.20}$$

Our main result concerning the optimality of (A.20) is summarized in Theorem 3. This theorem is the special case of Theorem 2 without market closure.

THEOREM 3. *There exists a unique solution to the periodic ODE system* (A.9) – (A.18)*, subject to the endogenous terminal condition* (A.19)*, such that U defined in* (A.8) *is nonnegative for all $(t,s,x,z) \in \mathcal{D}$. Furthermore, $\varphi^* \in \mathcal{A}$ is an optimal policy and the optimal value function is given by $U(0,s,x,z)$. More precisely, we have*

$$U(0,s,x,z) = V(0,s,x,z) = J(0,s,x,z;\varphi^*), \quad \forall (s,x,z) \in \mathbb{R}^+ \times \mathbb{R}^2.$$

The proof of this theorem is given in the next section.

B. Proof of Theorem 3

In this section, we prove Theorem 3. In particular, we establish the existence and uniqueness of classical solutions to the ODE system, as well as a verification theorem for the optimal strategy.

B.1. Main Steps

In the following, we denote a vector of functions as $\mathbf{v_u}$ as the solution (provided the solution exists and is unique, for which we will prove later) to the ODE system (A.9) – (A.18) with a given exogenous terminal condition, where

$$\mathbf{v_u} = (a, b_0, b_1, c_0, c_1, c_2, d, e_0, e_1, f)^T, \quad \mathbf{v_u}(T) = \mathbf{u}.$$

Let

$$\begin{aligned} F(\mathbf{v}(t))(s,x,z) = a(t)x^2 + (b_0(t) + b_1(t)z)xs + (c_0(t) + c_1(t)z + c_2(t)z^2)s^2 \\ + d(t)x + (e_0(t) + e_1(t)z)s + f(t). \end{aligned} \tag{B.1}$$

Also, denote an initial point for iteration as

$$\mathbf{u}_0 = (1/2, -\beta, 0, \beta^2/2, 0, 0, \beta - 1, (1-\beta)\beta, 0, (\beta-1)^2/2)^T.$$

Starting with terminal condition $\mathbf{v_u}(T) = \mathbf{u}$, we then use

$$\bar{\mathbf{u}}^\rho = \mathbf{u}_0 + e^{-\rho T}\mathbf{G}(\mathbf{v_u}(0)) \tag{B.2}$$

as the terminal condition for the next iteration, where

$$\begin{aligned} &\mathbf{G}_1(\mathbf{v_u}(0)) = (1-\gamma)^2(a(0) + b_0(0) + c_0(0) + d(0) + e_0(0) + f(0)) \\ &\mathbf{G}_3(\mathbf{v_u}(0)) = (1-\gamma)(b_1(0) + c_1(0) + e_1(0)) \\ &\mathbf{G}_6(\mathbf{v_u}(0)) = c_2(0) \\ &\mathbf{G}_j(\mathbf{v_u}(0)) = 0, \quad j = 2, 4, 5, 7, 8, 9, 10. \end{aligned} \tag{B.3}$$

Denote $B^+ = \{\mathbf{u} \in \mathbb{R}^{10} : F(\mathbf{u})(s,x,z) \geq 0, \forall (s,x,z) \in \mathbb{R}^+ \times \mathbb{R}^2\}$.

B.1.1. One-period Problem To study the infinite horizon problem (10) we first consider the following one-period problem. Denote

$$\begin{aligned} J_1(t,s,x,z;w,\varphi) = E\Bigg[&\frac{1}{2}\left(\beta \cdot (S_T - 1) - (X_T - 1)\right)^2 \\ &+ \int_t^1 \frac{1}{2}\Lambda S_u^2 \varphi_u^2 du + e^{-\rho T} w(S_1, X_1, \theta_1)\Bigg], \end{aligned} \tag{B.4}$$

for $w \in C(\mathbb{R}^+ \times \mathbb{R}^2)$, and

$$V_1(t,s,x,z;w) = \inf_{\varphi \in \mathcal{A}_1} J_1(t,s,x,z;w,\varphi). \tag{B.5}$$

Since the right-hand side of the ODE system (A.9) – (A.18) is analytic, for a given terminal condition $\mathbf{u}$ there has a solution locally, i.e., the solution exists in $[t_{min}, T]$, for some $0 \leq t_{min} < T$ (cf. Proposition 7.4 of Amann (1990)). Proposition 2 shows that the solution is unique in $[t_{min}, T]$, and Lemma 1 below shows that actually $t_{min} = 0$.

PROPOSITION 2 **(One-period Verification Theorem)**. *Let* $\mathbf{w} \in B^+$. *Denote* $U_1(t,s,x,z) = F(\mathbf{v_w}(t))(s,x,z)$, $\forall t \in [t_{min}, T]$, *using* (B.1).

(1) For any $\varphi \in \mathcal{A}_1$, $U_1 \leq J_1(\cdot; F(\mathbf{w}-\mathbf{u}_0), \varphi)$ *on* $[t_{min}, T] \times \mathbb{R}^+ \times \mathbb{R}^2$.

(2) Suppose there exists $\varphi^* \in \mathcal{A}_1$ *such that*

$$\varphi^*(u) \in arg\min_{\varphi \in \mathcal{A}_1} \left[\mathcal{L}^{\varphi} U_1(u, S_u, X_u^*, \theta_u^*) + \frac{1}{2}\Lambda S^2 \varphi^2 \right], a.e. \text{ on } [t_{min}, T] \times \Omega$$

Then $U_1 = V_1(\cdot; F(\mathbf{w}-\mathbf{u}_0))$ *on* $[t_{min}, T] \times \mathbb{R}^+ \times \mathbb{R}^2$.

(3) φ^* *defined in* (A.20) *satisfies the condition in (2).*

Thus, starting from $\mathbf{w}$ *at time* $t \in [t_{min}, T]$, U_1 *is the optimal value function and* φ^* *is an optimal control policy. Since* F *is a polynomial, we also have the coefficients of the polynomial as the vector* $\mathbf{v_w}(t)$ *is uniquely determined on* $[t_{min}, T]$.

LEMMA 1. *If* $\mathbf{u} \in B^+$, *then:*

1. *Given the terminal condition* $\mathbf{u}$, *the ODE system* (A.9) – (A.18) *is solvable up to time 0; in other words,* $t_{min} = 0$, *and consequently* $\mathbf{v_u}(t)$ *exists and is unique,* $\forall t \in [0, T]$;
2. $\mathbf{v_u}(t) \in B^+$, $\forall t \in [0, T]$;
3. $\forall t \in [0, T]$, *the value function for the one-period problem equals* $F(\mathbf{v_u}(t))$, *that is,*

$$F(\mathbf{v_u}(t))(s,x,z) = V_1(t,s,x,z; F(\mathbf{u}-\mathbf{u}_0)).$$

Furthermore, for any $\rho \geq 0$, $\bar{\mathbf{u}}^\rho \in B^+$, *recalling the definition in* (B.2).

The proofs of Proposition 2 and Lemma 1 will be given in Section B.2.

B.1.2. Infinite-horizon Problem via Fixed Point Iteration Now we turn to the infinite-horizon problem. Note that Lemma 1 shows that the ODE system is uniquely determined with terminal condition $\mathbf{u} \in B^+$. Therefore, $\mathbf{v_u}(t) \in B^+$ is defined for $t \in [0, T]$, and so is $\bar{\mathbf{u}}^\rho \in B^+$, recalling the definition in (B.2). As a result, the following mapping

$$\mathcal{T}^\rho := \{B^+ \to B^+ : \quad \mathbf{u} \mapsto \bar{\mathbf{u}}^\rho\}$$

is well-defined. To show the existence of solution to the original ODE system, it is equivalent to show that $\mathcal{T}$ has a fixed point, which is given in the following proposition.

PROPOSITION 3. *For all* $\rho > 0$, *denote* $\mathbf{u}^{(N)} = (\mathcal{T}^\rho)^N \mathbf{u}_0$. *Then (1)* $\mathcal{T}^\rho$ *is continuous on* B^+; *(2) for any* $(s,x,z) \in \mathbb{R}^+ \times \mathbb{R}^2$, $F(\mathbf{u}^{(N)}(s,x,z))_{N \geq 1}$ *is nondecreasing; (3)* $\lim_{N\to\infty} \mathbf{u}^{(N)} = \mathbf{u}^{(\infty)}$ *exists; and (4)* $\mathbf{u}^{(\infty)}$ *is a fixed point of* $\mathcal{T}^\rho$ *in* B^+.

Note that Proposition 3 also gives a method to solve for the fixed point for the periodic ODE system, namely starting from an initial guess. Finally, we prove a verification theorem for the infinite-horizon problem. Recall from the definition of F, (B.1), that for any fixed point $\mathbf{u}^* \in B^+$ of $\mathcal{T}^\rho$, $F(\mathbf{v}_{\mathbf{u}^*}(t))$ gives a polynomial of the form of (A.8) with $\mathbf{v}_{\mathbf{u}^*}(t)$ as coefficients.

PROPOSITION 4 **(Infinite-horizon Verification Theorem, No Overnight Jump)**. *Let $\rho > 0$ and any fixed point $\mathbf{u}^* \in B^+$ of $\mathcal{T}^\rho$, and let $U^*(t,\cdot) = F(\mathbf{v}_{\mathbf{u}^*}(t))(\cdot)$.*

(1) For any $\varphi \in \mathcal{A}$, $U^*(0,s,x,z) \le J(0,s,x,z)$ on $\mathbb{R}^+ \times \mathbb{R}^2$.

(2) Suppose there exists $\varphi^* \in \mathcal{A}$ such that

$$\varphi^*(u) \in arg\min_{\varphi \in \mathcal{A}} \left[\mathcal{L}^\varphi U^*(u, S_u, X_u^*, \theta_u^*) + \frac{1}{2}\Lambda S_u^2 \varphi_u^2 \right], a.e. \text{ on } [0,T] \times \Omega. \tag{B.6}$$

Then $U^*(0,s,x,z) = V(0,s,x,z)$ on $\mathbb{R}^+ \times \mathbb{R}^2$.

(3) φ^* defined in (A.20) *satisfies the conditions in (2).*

Proposition 4 also implies the uniqueness of the fixed point $\mathbf{u}^* \in B^+$. Indeed, for any fixed point $\hat{\mathbf{u}} \in B^+$, $F(\mathbf{v}_{\hat{\mathbf{u}}}(0))$ corresponds to the same value function V, so

$$\hat{\mathbf{u}} = \mathbf{u}_0 + e^{-\rho T}\mathbf{G}(\mathbf{v}_{\hat{\mathbf{u}}}(0)) = \mathbf{u}_0 + e^{-\rho T}\mathbf{G}(\mathbf{v}_{\mathbf{u}^*}(0)) = \mathbf{u}^*.$$

The proofs of Propositions 3 and 4 are given in Section B.2.

B.2. Additional Proofs

In the following, we denote $Y = \theta S$, then

$$dX_u = \mu Y_u du + \sigma Y_u dW_u + r(X_u - Y_u)du \tag{B.7}$$

$$dY_u = (\mu Y_u + \varphi_u S_u)du + \sigma Y_u dW_u. \tag{B.8}$$

[Proof of Proposition 2] Fix $t \in [t_{min}, T]$. For part (1), for any $\varphi \in \mathcal{A}_1$, by Ito's formula,

$$\begin{aligned} &U_1(u \wedge \tau_n, S_{u\wedge\tau_n}, X_{u\wedge\tau_n}, \theta_{u\wedge\tau_n}) \\ &= U_1(t,s,x,z) + \int_t^{u\wedge\tau_n} \left(\frac{\partial V}{\partial t} + \mathcal{L}^{\varphi_r} \right)(r, S_r, X_r, \theta_r)dr \\ &\quad + Martingale, \quad \forall u \in [t,T), \end{aligned}$$

where $\tau_n := \inf\{r \ge t : \int_t^r \sigma^2 \left(\theta_r \frac{\partial V}{\partial x} + \frac{\partial V}{\partial s}\right)^2 S_r^2 dr \ge n\}$. Therefore,

$$\begin{aligned} &E\left[U_1(u \wedge \tau_n, S_{u\wedge\tau_n}, X_{u\wedge\tau_n}, \theta_{u\wedge\tau_n})\right] \\ &= U_1(t,s,x,z) + E\left[\int_t^{u\wedge\tau_n} \left(\frac{\partial V}{\partial t} + \mathcal{L}^{\varphi_r} \right)(r, S_r, X_r, \theta_r)dr \right] \\ &\ge U_1(t,s,x,z) - E\left[\frac{1}{2}\Lambda \int_t^{u\wedge\tau_n} S_r^2 \varphi_r^2 dr \right]. \end{aligned}$$

Note that U_1 satisfies

$$|U_1(t,s,x,z)| \leq C(1+|x|^2+|s|^2+|zs|^2),$$

for some constant $C>0$. From the classic moment estimate (cf. Theorem 1.3.16 in Pham (2010)) we have

$$E\left[|X_{u\wedge\tau_n}|^2, |\theta_{u\wedge\tau_n}S_{u\wedge\tau_n}|^2\right] \leq C'\left(|x|^2+|zs|^2\right) + C'e^{C'T}\left[\int_0^T |x|^2+|zs|^2+\varphi_u^2 S_u^2 du\right].$$

for some $C'>0$. Therefore, for some $C_0, C_1 > 0$,

$$\begin{aligned} &E|U_1(u\wedge\tau_n, S_{u\wedge\tau_n}, X_{u\wedge\tau_n}, \theta_{u\wedge\tau_n})| \\ &\leq CE(1+|X_{u\wedge\tau_n}|^2+|S_{u\wedge\tau_n}|^2+|\theta_{u\wedge\tau_n}S_{u\wedge\tau_n}|^2) \\ &\leq C_0 + C_1 E\left[\int_0^T \varphi_u^2 S_u^2 du\right]. \end{aligned}$$

Also,

$$E\left[\int_t^{u\wedge\tau_n} \varphi_r^2 S_r^2 dr\right] \leq E\left[\int_t^T \varphi_r^2 S_r^2 dr\right].$$

Therefore, thanks to the dominated convergence theorem, by first sending $n\to\infty$ and then $u\to T$, we have

$$E[F(\mathbf{w})(S_T, X_T, \theta_T)] \geq U_1(t,s,x,z) - E\left[\frac{1}{2}\Lambda\int_t^T \varphi_r^2 S_r^2 dr\right], \forall \varphi\in\mathcal{A},$$

and $U_1 \leq J_1$ follows.

For part (2), repeating the proof of 1) shows that

$$U_1(t,s,x,z) = J_1(t,s,x,z;\varphi^*) \geq V_1(t,s,x,z),$$

which, together with 1), shows $U_1 = V_1$.

For part (3), note that φ^* defined in (A.20) together with V defined in (A.8) solves the HJB equation. Furthermore,

$$\varphi_u^* S_u = -\frac{1}{\Lambda}\left(b_1(u)X_u + c_1(u)\frac{\bar{x}}{\bar{s}}S_u + 2c_2(u)\theta_u S_u + e_1(u)\bar{x}\right),$$

and

$$\begin{aligned} dY_u &= \mu Y_u du + \sigma Y_u dW_u + \varphi_u^* S_u du \\ &= \left(\mu - \frac{2}{\Lambda}c_2(u)\right)Y_u du + \sigma Y_u dW_u - \frac{1}{\Lambda}\left(b_1(u)X_u + c_1(u)\frac{\bar{x}}{\bar{s}}S_u + e_1(u)\bar{x}\right)du. \end{aligned}$$

Due to the continuity of b_1, c_1, c_2, and e_1, the drift and diffusion terms of the SDEs for (X,Y) are Lipschitz in (X,Y) and satisfy a linear growth condition. Therefore, again from the classic SDE moment estimate, we have

$$E[\max_{t\le s\le T}(|X_s|^2+|Y_s|^2)]\le C(|x|^2+|zs|^2)+Ce^{C(T-t)}E\left[\int_t^T(|x|^2+|zs|^2+S_u^2+1)du\right]$$

As a result,

$$E\left[\int_t^T(\varphi_u^*)^2S_u^2du\right]=E\left[\int_t^T Y_u^2du\right]<\infty,$$

so that $\varphi^*\in\mathcal{A}_1$. Q.E.D.

[Proof of Lemma 1] For notational convenience, in the following we reverse time and show that the solution to the time-reversed system, still denoted as $\mathbf{v_u}(t)$, exists for $t\in(0,T]$.

We shall prove the result by contradiction. Assume the system blows up at time $\tau\in(0,T]$.

By Proposition 2,

$$F(\mathbf{v_u}(t))(s,x,z)=V_1(T-t,s,x,z;F(\mathbf{u}-\mathbf{u}_0))\ge 0,\quad t\in[0,\tau).$$

Therefore, for any $t\in[0,\tau)$, we have $a(t)\ge 0$ and $c_2(t)\ge 0$; otherwise, the function U_1 goes to $-\infty$ as $x\to\infty$ and $s\to\infty$. In addition, $\sup_{t\in[0,\tau)}a(t)<\infty$ thanks to the (time reverse) ODE of a, because the drift is at most linear.

Next, the (time-reversed) system for b_1 and c_2 is given by

$$\begin{aligned}\frac{\partial b_1}{\partial t}&=-\frac{2}{\Lambda}b_1c_2+2(\mu-r)a+(\mu+r)b_1\\ \frac{\partial c_2}{\partial t}&=\sigma^2a+(\sigma^2+2\mu)c_2+(\sigma^2+\mu-r)b_1-\frac{2}{\Lambda}c_2^2.\end{aligned}$$

Using the same idea as on page 11 of Reid (1972), consider an alternative linear system

$$\frac{d}{dt}Y(t)=\begin{pmatrix}0 & 0 & \frac{2}{\Lambda}\\ 2(\mu-r)a(t) & \mu+r & 0\\ \sigma^2a(t) & \sigma^2+\mu-r & \sigma^2+2\mu\end{pmatrix}Y(t),$$

with $Y(t)=(Y_1(t),Y_2(t),Y_3(t))^T$ and $Y(0)=(1,b_1(0),c_2(0))^T$. Assume that $Y_1(\eta-)$ reaches 0 at $\eta\le\tau$ for the first time. Note that $(b_1(t),c_2(t))=(Y_2(t)/Y_1(t),Y_3(t)/Y_1(t))$ gives a solution to the system of b_1 and c_2 for $t\in[0,\eta)$. Therefore, $Y_3(t)\ge 0$ for $t<\eta$, because $c_2(t)\ge 0$ and $Y_1(t)>0$ (as Y_1 starts from a positive number). As a result, the equation for Y_1 implies that Y_1, starting from 1, is nondecreasing and hence positive for $t<\eta$, contradicting the assumption that $Y_1(\eta-)=0$. This implies that b_1 and c_2 are not blowing up at τ.

Given a,b_1,c_2, the equations for b_0,c_1,d,e_1 in the ODE system are linear. Since a,b_1,c_2, as time-dependent coefficients in this linear equation, do not blow up at τ, we have b_0,c_1,d,e_1 do not blow up at τ as well. Given these seven components, the remaining equation for c_0,e_0,f are again linear,

meaning that they are also not blowing up at τ. This contradicts with the assumption that the system blows up at $\tau \le T$. Q.E.D.

[Proof of Proposition 3] Consider the strategy $\hat{\varphi}_u = -\frac{z}{T}\chi_{u\in[0,T]}$ which liquidates all the stock position on the first day. It is straightforward to see $\hat{\varphi} \in \mathcal{A}$. Using this strategy, the expected cost on the first day is finite (although suboptimal); from the second day, there is no market friction cost, and the aggregated deviation cost from day 2 to N is

$$E\left[\sum_{i=2}^{N} e^{-\rho(i-1)T}\frac{1}{2}\beta^2\left(\frac{S_{iT}}{S_{(i-1)T}}-1\right)^2\right] \le E\left[\sum_{i=2}^{\infty} e^{-\rho(i-1)T}\frac{1}{2}\beta^2\left(\frac{S_{iT}}{S_{(i-1)T}}-1\right)^2\right] < \infty.$$

This cost is finite for any $\rho > 0$ since the squared simple returns of S are i.i.d. Therefore,

$$\begin{aligned}
&J_N(t,s,x,z;0,\hat{\varphi},\bar{s},\bar{x}) \\
&\quad = E\left[\frac{1}{2}\left(\beta\left(\frac{S_1}{s}-1\right)-\left(\frac{\hat{X}_1}{x}-1\right)\right)^2 + \sum_{i=2}^{N}\frac{1}{2}e^{-\rho(i-1)T}\left(\beta\left(\frac{S_{iT}}{S_{(i-1)T}}-1\right)\right)^2\right. \\
&\quad \left. + \frac{1}{2}\Lambda\frac{z^2}{T^2}\int_0^T S_u^2 du\right] \\
&\quad \le J_\infty(t,s,x,z;0,\hat{\varphi},\bar{s},\bar{x}) < \infty.
\end{aligned}$$

Now, for $\rho > 0$, recall that $\mathbf{u}^{(N)} = (\mathcal{T}^\rho)^N \mathbf{u}_0$. By iterating the one-period for $N-1$ times, we have

$$F(\mathbf{u}^{(N)})(s,x,z) = e^{-\rho T}V_N(0,s,x,z;0) + \frac{1}{2}(\beta(s-1)-(x-1))^2. \tag{B.9}$$

Note that the right-hand side is nondecreasing as $N \to \infty$, since both costs for deviation and price impact are nonnegative. Also, note that since $\hat{\varphi}$ is admissible,

$$V_N(t,s,x,z;0) \le J_N(t,s,x,z;0,\hat{\varphi}).$$

Therefore, for any $(s,x,z) \in \mathbb{R}^+ \times \mathbb{R}^2$, $(F(\mathbf{u}^{(N)})(s,x,z))_{N\ge1}$ are nondecreasing, have a finite upper bound, and therefore are convergent as $N \to \infty$. Since $F(\mathbf{u}^{(N)})(s,x,z)$ is a polynomial in (s,x,z), this further implies the convergence of $(\mathbf{u}^{(N)})_{N\ge1}$. Next, recall that by definition, $\mathcal{T}^\rho \mathbf{u} = \mathbf{u}_0 + e^{-\rho T}\mathbf{G}(\mathbf{v}_\mathbf{u}(0))$. Thanks to the continuous dependence of solution of ODE on the initial data (cf. Theorem 8.3 in Amann (1990)), and the continuity of $\mathbf{G}$, $\mathcal{T}^\rho$ is continuous. As a result, sending $N \to \infty$ in $\mathbf{u}^{(N+1)} = \mathcal{T}^\rho \mathbf{u}^{(N)}$ shows that $\mathbf{u}^{(\infty)}$ is a fixed point of $\mathcal{T}^\rho$. Q.E.D.

[Proof of Proposition 4] We follow the notation in Proposition 3. For 1), for any $\varphi \in \mathcal{A}$, (B.9) implies

$$F(\mathbf{u}^{(N)})(s,x,z) \le e^{-\rho T}J_N(0,s,x,z;0) + \frac{1}{2}(\beta(s-1)-(x-1))^2. \tag{B.10}$$

Sending $N \to \infty$, by monotone convergence theorem and the definition of $\mathbf{u}^*$, we have

$$F(\mathbf{u}^*)(s,x,z) \leq e^{-\rho T} J_\infty(0,s,x,z;0,\varphi) + \frac{1}{2}(\beta(s-1)-(x-1))^2,$$

therefore,

$$U(0,s,x,z) \leq J_\infty(0,s,x,z;0,\varphi) = J(0,s,x,z;\varphi).$$

Part 2) follows from repeating the above proof and replacing the inequality (B.10) by the equality (B.9). For 3), it is straightforward to verify that φ^* defined in (A.20) satisfies (B.6). Therefore, since $V_N(0,s,x,z) \leq J_N(0,s,x,z;\hat{\varphi})$, we have

$$\begin{aligned}
&E\left[\sum_{i=1}^{N} e^{-\rho(i-1)T}\int_{(i-1)T}^{iT}\frac{1}{2}\Lambda(\varphi_u^*)^2 S_u^2 du\right] \\
&\leq E\left[\int_0^T \frac{1}{2}\Lambda\hat{\varphi}_u^2 S_u^2 du + \frac{1}{2}\left(\beta(S_T-1)-(\hat{X}_T-1)\right)^2 + \sum_{i=2}^{N} e^{-\rho(i-1)T}\frac{1}{2}\beta^2\left(\frac{S_{iT}}{S_{(i-1)T}}-1\right)^2\right] < \infty.
\end{aligned}$$

Q.E.D.

C. Evaluating Cost Functionals for Suboptimal Strategies without Market Overnight Closure

In this section, we discuss how to calculate $J(0,s,x,z;\varphi)$ for $\varphi \in \mathcal{A}$ that is not necessarily optimal. It turns out that as long as φ has the same form as (12) with different coefficients $\tilde{b}_1, \tilde{c}_1, \tilde{c}_2, \tilde{e}_1$, J can still be calculated via an iterative procedure, in terms of a system of ODEs

$$\begin{aligned}
\frac{\partial a}{\partial t} &= \frac{1}{\tilde{\Lambda}}\tilde{b}_1 b_1 - \frac{\Lambda}{2\tilde{\Lambda}^2}\tilde{b}_1^2 - 2ra \\
\frac{\partial b_0}{\partial t} &= \frac{1}{\tilde{\Lambda}}(\tilde{b}_1 c_1 + \tilde{c}_1 b_1) - \frac{\Lambda}{\tilde{\Lambda}^2}\tilde{b}_1\tilde{c}_1 - (\mu+r)b_0 \\
\frac{\partial b_1}{\partial t} &= \frac{2}{\tilde{\Lambda}}(\tilde{b}_1 c_2 + \tilde{c}_2 b_1) - \frac{2\Lambda}{\tilde{\Lambda}^2}\tilde{b}_1\tilde{c}_2 - 2(\mu-r)a - (\mu+r)b_1 \\
\frac{\partial c_0}{\partial t} &= -(\sigma^2+2\mu)c_0 + \frac{1}{\tilde{\Lambda}}\tilde{c}_1 c_1 - \frac{\Lambda}{2\tilde{\Lambda}^2}\tilde{c}_1^2 \\
\frac{\partial c_1}{\partial t} &= -(\sigma^2+2\mu)c_1 - (\sigma^2+\mu-r)b_0 + \frac{2}{\tilde{\Lambda}}(\tilde{c}_1 c_2 + \tilde{c}_2 c_1) - \frac{2\Lambda}{\tilde{\Lambda}^2}\tilde{c}_1\tilde{c}_2 \\
\frac{\partial c_2}{\partial t} &= -\sigma^2 a - (\sigma^2+2\mu)c_2 - (\sigma^2+\mu-r)b_1 + \frac{4}{\tilde{\Lambda}}\tilde{c}_2 c_2 - \frac{2\Lambda}{\tilde{\Lambda}^2}\tilde{c}_2^2 \\
\frac{\partial d}{\partial t} &= \frac{1}{\Lambda}(\tilde{b}_1 e_1 + \tilde{e}_1 b_1) - \frac{\Lambda}{\tilde{\Lambda}^2}\tilde{b}_1\tilde{e}_1 - rd \\
\frac{\partial e_0}{\partial t} &= -\mu e_0 + \frac{1}{\tilde{\Lambda}}(\tilde{c}_1 e_1 + \tilde{e}_1 c_1) - \frac{\Lambda}{\tilde{\Lambda}^2}\tilde{c}_1\tilde{e}_1 \\
\frac{\partial e_1}{\partial t} &= -(\mu-r)d - \mu e_1 + \frac{2}{\Lambda}(\tilde{c}_2 e_1 + \tilde{e}_1 c_2) - \frac{2\Lambda}{\tilde{\Lambda}^2}\tilde{c}_2\tilde{e}_1 \\
\frac{\partial f}{\partial t} &= \frac{1}{\tilde{\Lambda}}\tilde{e}_1 e_1 - \frac{\Lambda}{2\tilde{\Lambda}^2}\tilde{e}_1^2.
\end{aligned}$$

Specifically, for $\mathbf{u} \in B^+$, define $\tilde{\mathbf{v}}_{\mathbf{u}}$ as the solution to the above system given terminal condition $\tilde{\mathbf{v}}_{\mathbf{u}}(T) = \mathbf{u}$. Furthermore, denote $\tilde{\mathbf{v}}^{(1)} = \tilde{\mathbf{v}}_{\mathbf{u}_0}$, and $\mathbf{u}_{k+1} = \mathbf{u}_0 + e^{-\rho T}\mathbf{G}(\tilde{\mathbf{v}}^{(k)}(0))$, $\tilde{\mathbf{v}}^{(k+1)} = \tilde{\mathbf{v}}_{\mathbf{u}_k}$, for all $k \geq 1$.

PROPOSITION 5. *For any*

$$\tilde{\varphi}(t,s,x,z;\bar{s},\bar{x}) = -\frac{1}{\tilde{\Lambda}}\left(\tilde{b}_1(t)\frac{x}{s} + \tilde{c}_1(t)\frac{\bar{x}}{\bar{s}} + 2\tilde{c}_2(t)z + \tilde{e}_1(t)\frac{\bar{x}}{s}\right),$$

where $(\tilde{b}_1(t), \tilde{c}_1(t), \tilde{c}_2(t), \tilde{e}_1(t)) \in \mathbb{R}^4$ *and continuous with respect to* t*, and* $\Lambda > 0$*, we have*

$$J(t,s,x,z,\bar{s},\bar{x};\tilde{\varphi}) = \bar{x}^2 J(t, \frac{s}{\bar{s}}, \frac{x}{\bar{x}}, \frac{\bar{s}}{\bar{x}}z, 1, 1; \tilde{\varphi}).$$

Also, $\tilde{\mathbf{v}}^k$ *is well-defined for* $k \geq 1$*. Furthermore,* $(F(\tilde{\mathbf{v}}^k))_{k\geq 1}$ *is nondecreasing, and*

$$\lim_{k\to\infty} F(\tilde{\mathbf{v}}^k(t)) = J(t,s,x,z,1,1;\tilde{\varphi}). \tag{C.1}$$

[Proof] The homotheticity of J follows the same argument for V. The PDE corresponding to J is (A.4) with φ chosen as the feedback form in this proposition instead of taking infimum. Then it is straightforward to verify that, by the ansatz (11), the HJB equation reduces to the above ODE system.

Note that this ODE system is linear with continuous coefficients, so the existence of solution $\mathbf{v}_{\mathbf{u}^{(0)}}$ in $t \in [0,T]$ of the one-period problem is guaranteed. On the other hand, a direct application of Ito's formula shows that $F(\mathbf{v}_{\mathbf{u}^{(0)}}(t))$ indeed equals the one-period cost functional $J_1(t,s,x,z;\tilde{\varphi})$. Therefore, $\mathbf{u}^{(1)} \in B^+$ from the positivity of J_1. By repeating this argument $N-1$ times, (11) with coefficients given by $F(\mathbf{v}_{\mathbf{u}^{(N-1)}}(t))$ equals the N-period cost functional $J_N(t,s,x,z;\tilde{\varphi})$.

Since the daily costs from deviation and price impact are nonnegative, J_N is nondecreasing in N, and so is $F(\mathbf{v}_{\mathbf{u}^{(N-1)}}(t))$ as a result. Therefore, equation (C.1) holds. Q.E.D.

D. The Case with Overnight Price Changes

In the general case with overnight price jump, the main results can be derived by adapting the proofs in the case without an overnight price jump. Here, we only point out the main modification in the proofs.

First, note that

$$S_{t_{2i+2}} = S_{t_{2i+1}} \cdot (1+\zeta_i), \quad i \geq 0,$$

where $\{\zeta_i\}_{i\geq 0}$ is a collection of i.i.d. random variables with the same distribution as

$$\zeta = e^{(\mu_n - \sigma_n^2/2)\delta T + \sigma_n\sqrt{\delta T}Z} - 1,$$

where Z follows a standard normal distribution. Also,

$$X_{t_{2i+2}} = e^{r\delta T}(X_{t_{2i+1}} - \theta_{t_{2i+1}} S_{t_{2i+1}}) + \theta_{t_{2i+1}} S_{t_{2i+1}}(1+\zeta_i),$$

It is straightforward to calculate that

$$\begin{aligned} E(1+\zeta) &= e^{\mu_n \delta T} \\ E(1+\zeta)^2 &= e^{(2\mu_n+\sigma_n^2)\delta T} \\ E\zeta(1+\zeta) &= e^{(2\mu_n+\sigma_n^2)\delta T} - e^{\mu_n \delta T} \\ E\zeta^2 &= e^{(2\mu_n+\sigma_n^2)\delta T} - 2e^{\mu_n \delta T} + 1. \end{aligned}$$

At each market opening t_{2i}, $i \geq 0$, the manager is faced with an infinite-horizon problem where the state variables have the initial value $(S_{t_{2i}}, X_{t_{2i}})$ while the reference values are $(S_{t_{2i-1}}, X_{t_{2i-1}})$. Therefore, the value function at the market opening is

$$\begin{aligned} &V(t_{2i}, S_{t_{2i}}, X_{t_{2i}}, \theta_{t_{2i}}; S_{t_{2i-1}}, X_{t_{2i-1}}) \\ &= V(t_{2i}, S_{t_{2i-1}}(1+\zeta), X_{t_{2i-1}} e^{r\delta T} + \theta_{t_{2i-1}} S_{t_{2i-1}}(1 - e^{r\delta T} + \zeta), \theta_{t_{2i-1}}; S_{t_{2i-1}}, X_{t_{2i-1}}), \end{aligned}$$

and the value function at the previous market closing t_{2i-1} is

$$\begin{aligned} &V(t_{2i-1}, S_{t_{2i-1}}, X_{t_{2i-1}}, \theta_{t_{2i-1}}; S_{t_{2i-1}}, X_{t_{2i-1}}) \\ &\quad = E_{t_{2i-1}}[V(t_{2i}, S_{t_{2i}}, X_{t_{2i}}, \theta_{t_{2i-1}}; S_{t_{2i-1}}, X_{t_{2i-1}})] \\ &\quad = E[V(t_{2i}, S_{t_{2i-1}}(1+\zeta), X_{t_{2i-1}} e^{r\delta T} + \theta_{t_{2i-1}} S_{t_{2i-1}}(1 - e^{r\delta T} + \zeta), \theta_{t_{2i-1}}; S_{t_{2i-1}}, X_{t_{2i-1}})], \end{aligned}$$

where the last expectation is with respect to ζ.

Note that if V satisfies the quadratic form as in (11) at t_{2i+1}, then V also has this form at t_{2i} by the same reasoning as in the case without market closure. Next, we show that if V has this form at t_{2i}, it also has this form at t_{2i-1}, from which we see that V has this form for all $t \geq 0$. To this

end, note that the homotheticity of V as in Proposition 1 still holds. Therefore,

$$\begin{aligned}
&V(t_{2i}, S_{t_{2i-1}}(1+\zeta), X_{t_{2i-1}}e^{r\delta T} + \theta_{t_{2i-1}}S_{t_{2i-1}}(1-e^{r\delta T}+\zeta), \theta_{t_{2i-1}}; S_{t_{2i-1}}, X_{t_{2i-1}}) \\
&\quad = X_{t_{2i-1}}^2 V(0, 1+\zeta, e^{r\delta T} + \theta_{t_{2i-1}}S_{t_{2i-1}}(1-e^{r\delta T}+\zeta)/X_{t_{2i-1}}, \theta_{t_{2i-1}}S_{t_{2i-1}}/X_{t_{2i-1}}; 1, 1) \\
&\quad = X_{t_{2i-1}}^2 \Bigg[a(e^{r\delta T} + \theta_{t_{2i-1}}S_{t_{2i-1}}(1-e^{r\delta T}+\zeta)/X_{t_{2i-1}})^2 \\
&\qquad + (b_0 + b_1\theta_{t_{2i-1}}S_{t_{2i-1}}/X_{t_{2i-1}})(1+\zeta)(e^{r\delta T} + \theta_{t_{2i-1}}S_{t_{2i-1}}(1-e^{r\delta T}+\zeta)/X_{t_{2i-1}}) \\
&\qquad + (c_0 + c_1\theta_{t_{2i-1}}S_{t_{2i-1}}/X_{t_{2i-1}} + c_2(\theta_{t_{2i-1}}S_{t_{2i-1}}/X_{t_{2i-1}})^2)(1+\zeta)^2 \\
&\qquad + d(e^{r\delta T} + \theta_{t_{2i-1}}S_{t_{2i-1}}(1-e^{r\delta T}+\zeta)/X_{t_{2i-1}}) + (e_0 + e_1\theta_{t_{2i-1}}S_{t_{2i-1}}/X_{t_{2i-1}})(1+\zeta) + f \Bigg] \\
&\quad = a(X_{t_{2i-1}}e^{r\delta T} + \theta_{t_{2i-1}}S_{t_{2i-1}}(1-e^{r\delta T}+\zeta))^2 \\
&\qquad + (b_0 X_{t_{2i-1}} + b_1\theta_{t_{2i-1}}S_{t_{2i-1}})(1+\zeta)(e^{r\delta T}X_{t_{2i-1}} + \theta_{t_{2i-1}}S_{t_{2i-1}}(1-e^{r\delta T}+\zeta)) \\
&\qquad + (c_0 X_{t_{2i-1}}^2 + c_1\theta_{t_{2i-1}}S_{t_{2i-1}}X_{t_{2i-1}} + c_2\theta_{t_{2i-1}}^2 S_{t_{2i-1}}^2)(1+\zeta)^2 \\
&\qquad + d(e^{r\delta T}X_{t_{2i-1}}^2 + \theta_{t_{2i-1}}S_{t_{2i-1}}X_{t_{2i-1}}(1-e^{r\delta T}+\zeta)) \\
&\qquad + (e_0 X_{t_{2i-1}}^2 + e_1\theta_{t_{2i-1}}S_{t_{2i-1}}X_{t_{2i-1}})(1+\zeta) + fX_{t_{2i-1}}^2.
\end{aligned}$$

As a result, $V(t_{2i}, s, x, z; s, x)$ still has the same quadratic form. Since

$$\begin{aligned}
V(t_{2i-1}-, s, x, z; 1, 1) = &\frac{1}{2}\left(\beta(s-1) - (x-1)\right)^2 \\
&+ e^{-\rho(1+\delta)T}(1-\gamma)^2 x^2 V\left(t_{2i-1}, 1, 1, \frac{sz}{(1-\gamma)x}; 1, 1\right),
\end{aligned}$$

the connection condition (A.19) for the quadratic ODE system has the following modification:

$$\begin{aligned}
a(T) &= \frac{1}{2} + e^{-\rho(1+\delta)T}(1-\gamma)^2\big(a(0)e^{2r\delta T} + b_0(0)e^{r\delta T}E(1+\zeta) + c_0(0)E(1+\zeta)^2 \\
&\quad + d(0)e^{r\delta T} + e_0(0)E(1+\zeta) + f(0)\big) \\
b_0(T) &= -\beta \\
b_1(T) &= e^{-\rho(1+\delta)T}(1-\gamma)\Big(2a(0)e^{r\delta T}(1-e^{r\delta T}+E\zeta) + b_0(0)[E(\zeta(1+\zeta)) + (1-e^{r\delta T})E(1+\zeta)] \\
&\quad + b_1(0)E(1+\zeta)e^{r\delta T} + c_1(0)E(1+\zeta)^2 + d(0)(1-e^{r\delta T}+E\zeta) + e_1(0)E(1+\zeta)\Big) \\
c_0(T) &= \frac{1}{2}\beta^2, \quad c_1(T) = 0, \\
c_2(T) &= e^{-\rho(1+\delta)T}(a(0)E(1-e^{r\delta T}+\zeta)^2 + b_1(0)E((1-e^{r\delta T}+\zeta)(1+\zeta)) + c_2(0)E(1+\zeta)^2), \\
d(T) &= \beta - 1, \quad e_0(T) = \beta(1-\beta), \quad e_1(T) = 0, \quad f(T) = \frac{1}{2}(\beta-1)^2.
\end{aligned} \tag{D.1}$$

Note that when there is no overnight jump (i.e., $\zeta = 0$), we recover the solution in the case without overnight price jump.

By adapting the proofs to account for the overnight jump of S, one can then prove the analog of Proposition 3 and Proposition 4. Specifically, the strategy $\hat{\varphi}$ still provides an upper bound for the value function, and its corresponding expected total cost is still finite, since the daily return of S is still i.i.d. after incorporating overnight jumps.

E. Expected Aggregated Costs with Overnight Price Changes

This section compares the three strategies, Periodic-DN, Periodic-D and One-Day, in terms of the value function. Using Proposition 5 and including the overnight jumps by incorporating the changes to the terminal condition described in Section D, we calculate the expected aggregated costs from return deviation and market frictions, i.e., the first and second term in (5): $E\left[\frac{1}{2}\sum_{i=0}^{\infty} e^{-\rho t_{2i}} X_{t_{2i-1}}^2 D_i^2\right]$ and $E\left[\frac{1}{2}\sum_{i=0}^{\infty} e^{-\rho t_{2i}} \int_{t_{2i}}^{t_{2i+1}} \Lambda S_u^2 \varphi_u^2 du\right]$.

The results, as reported in Table 6, is consistent with the observations in Section 7.2 and 6.1. First, the Periodic-DN strategy has the lowest cost from deviation, even though the cost is now measured using *squared* slippage. Second, as explained in 6.1, although the Periodic-D and OD strategies can have a lower cost from market frictions as compared to the Periodic-DN strategy, it is at the expense of increasing the cost from deviation. Indeed, they rebalance less since they do not account for the overnight risk and do not prepare for tomorrow, respectively, both of which lead to a suboptimal position level and larger slippage. Panel C shows that, by balancing the costs from deviation and market frictions, Periodic-DN strategy indeed has the lowest total cost.

Table 6 Expected Aggregated Cost from Deviation and Market Frictions

	Periodic-DN			Periodic-D			OD		
Λ	(2x)	(-1x)	(-2x)	(2x)	(-1x)	(-2x)	(2x)	(-1x)	(-2x)
Panel A: Aggregated Cost from Deviation									
10^{-6}	0.7737	0.1642	0.6995	1.4779	0.2408	1.3380	4.4323	0.5630	4.0233
10^{-5}	1.7147	0.2664	1.5515	2.1442	0.3133	1.9416	5.6511	0.6966	5.1369
10^{-4}	4.2085	0.5373	3.8097	4.4799	0.5673	4.0592	15.9274	1.8365	14.6284
10^{-3}	15.0365	1.7105	13.5939	15.2108	1.7323	13.7777	121.3283	14.8168	122.0370
Panel B: Aggregated Cost from Market Frictions									
10^{-6}	0.4347	0.0473	0.3940	0.3323	0.0362	0.3017	0.5485	0.0598	0.4986
10^{-5}	1.2888	0.1405	1.1698	1.1812	0.1288	1.0729	1.7147	0.1872	1.5603
10^{-4}	5.0419	0.5522	4.5887	4.8773	0.5340	4.4374	3.4557	0.3814	3.1763
10^{-3}	18.4982	2.0622	17.0329	18.3481	2.0431	16.8720	3.6694	0.4461	3.6931
Panel C: Total Costs									
10^{-6}	1.2084	0.2115	1.0935	1.8102	0.2770	1.6397	4.9808	0.6228	4.5219
10^{-5}	3.0035	0.4069	2.7213	3.3254	0.4421	3.0145	7.3658	0.8838	6.6972
10^{-4}	9.2504	1.0895	8.3984	9.3572	1.1013	8.4966	19.3831	2.2179	17.8047
10^{-3}	33.5347	3.7727	30.6268	33.5589	3.7754	30.6497	124.9977	15.2629	125.7301

This table reports the aggregated cost from deviation and market frictions, respectively (i.e., the first and second terms in (5), respectively), as well as the total cost (the whole (5)), for different β and Λ. All numbers are reported as 10^{-5} squared dollars. The total costs for the Periodic-DN, PD and OD strategies are calculated via the original cost functional (10) under these three strategies, respectively (see Section C). The costs from deviation for these three strategies are calculated via the cost functional (10) with $\Lambda = 0$ under these three strategies, respectively. The costs from market friction are then calculated as the difference between the total costs and the costs from deviation.

F. The Case with Overnight Price Changes but without Market Frictions

The proof of Theorem 1 is similar to that of Theorem 3. Here we mainly point out the difference.

The HJB equation is

$$\frac{\partial U}{\partial t}+\min_{\theta}\left\{\frac{1}{2}\sigma^2\theta^2s^2\frac{\partial^2U}{\partial x^2}+\frac{1}{2}\sigma^2s^2\frac{\partial^2U}{\partial s^2}+\theta\sigma^2s^2\frac{\partial^2U}{\partial x\partial s}+((\mu-r)\theta s+rx)\frac{\partial U}{\partial x}+\mu s\frac{\partial U}{\partial S}\right\}=0, \quad \text{(F.1)}$$

with the connection condition

$$\begin{aligned}U(t_1-,s,x)&=\frac{1}{2}(x-\beta s+(\beta-1))^2+e^{-\rho t_1}E[U(t_1,s,x(1-\gamma);s,x(1-\gamma))]\\&=\frac{1}{2}(x-\beta s+(\beta-1))^2+e^{-\rho t_1}\cdot\\&\quad\min_{\theta_T}E\left[U\left(0,s(1+\zeta_1),x(1-\gamma)e^{r\delta T}+\theta_Ts(1-e^{r\delta T}+\zeta_T);s,x(1-\gamma)\right)\right].\end{aligned}$$

Again, we make a quadratic Ansatz:

$$U(t,s,x)=a(t)x^2+b(t)xs+c(t)s^2+d(t)x+e(t)s+f(t),$$

where $a(t)>0$. Plugging this Ansatz into the HJB equation and compute the first order condition, we get the optimizer

$$\theta^*=-\frac{1}{2\sigma^2a(t)}\left[b(t)\sigma^2+(\mu-r)\left(b(t)+d(t)\frac{1}{s}+2a(t)\frac{x}{s}\right)\right]. \quad \text{(F.2)}$$

Therefore, the HJB equation becomes

$$0=\frac{\partial a}{\partial t}+\left(2r-\frac{(r-\mu)^2}{\sigma^2}\right)a \quad \text{(F.3)}$$

$$0=\frac{\partial b}{\partial t}+\left(2r-\frac{(r-\mu)^2}{\sigma^2}\right)b \quad \text{(F.4)}$$

$$0=\frac{\partial c}{\partial t}+(2\mu+\sigma^2)c-\frac{(\mu-r+\sigma^2)^2b^2}{4\sigma^2a} \quad \text{(F.5)}$$

$$0=\frac{\partial d}{\partial t}+\left(r-\frac{(r-\mu)^2}{\sigma^2}\right)d \quad \text{(F.6)}$$

$$0=\frac{\partial e}{\partial t}+\mu e-\frac{(\mu-r)(\mu-r+\sigma^2)bd}{2\sigma^2a} \quad \text{(F.7)}$$

$$0=\frac{\partial f}{\partial t}-\frac{(r-\mu)^2d^2}{4\sigma^2a}. \quad \text{(F.8)}$$

Plugging this Ansatz into the connection condition,

$$\begin{aligned}
&U(T,s,x(1-\gamma);s,x(1-\gamma))\\
&=x^2(1-\gamma)^2\min_{\theta_1}E\Bigg[a(0)(e^{r\delta T}+\theta_1\frac{s}{x(1-\gamma)}(1-e^{r\delta T}+\zeta_1))^2\\
&\quad+b(0)(e^{r\delta T}+\theta_1\frac{s}{x(1-\gamma)}(1-e^{r\delta T}+\zeta_1))(1+\zeta_1)+c(0)(1+\zeta_1)^2\\
&+d(0)(e^{r\delta T}+\theta_1\frac{s}{x(1-\gamma)}(1-e^{r\delta T}+\zeta_1))+e(0)(1+\zeta_1)+f(0)\Bigg]\\
&=-\frac{x^2(1-\gamma)^2}{4a(0)E_4}\Bigg[E_1^2(E_3^2-E_4)\left((2a(0)+b(0))^2-b(0)^2\right)+\left(E_3d(0)+(E_1E_3+E_4)b(0)\right)^2\\
&\quad+4a(0)\left(d(0)E_1(E_3^2-E_4)-e(0)E_4(E_1+E_3)-c(0)E_4(E_1^2+2E_1E_3+E_4)-f(0)E_4\right)\Bigg]\\
&:=x^2(1-\gamma)^2K,
\end{aligned}$$

where

$$\begin{aligned}
&E_1=e^{r\delta T},\quad E_3=E(1+\zeta_1-e^{r\delta T})=e^{\mu_n\delta T}-e^{r\delta T},\\
&E_4=E(1+\zeta_1-e^{r\delta T})^2=\left(e^{\mu_n\delta T}-e^{r\delta T}\right)^2+e^{2\mu_n\delta T}(e^{\sigma_n^2\delta T}-1).
\end{aligned}$$

The optimizer is

$$\theta_T^*=-\frac{d(0)}{2a(0)}\frac{E_3}{E_4}-\left(1+\frac{b(0)}{2a(0)}\right)\frac{E_1E_3}{E_4}-\frac{b(0)}{2a(0)}. \tag{F.9}$$

So the terminal condition for the ODE system is

$$\begin{aligned}
&a(T)=\frac{1}{2}+e^{-\rho t_1}(1-\gamma)^2K,\quad b(T)=-\beta,\quad c(T)=\frac{\beta^2}{2}\\
&d(T)=\beta-1,\quad e(T)=\beta(1-\beta),\quad f(T)=\frac{1}{2}(1-\beta)^2.
\end{aligned} \tag{F.10}$$

Note that K in the terminal condition depends on the time-0 value of the solution, which will be determined via a fixed point iteration.

The existence and uniqueness of the solution to the one-period ODE problem can be proved straightforwardly. The existence and uniqueness of a fixed point for the infinite-horizon problem can be established using an iterative method similar to Theorem 3, and the verification result similar to Proposition 4.

In the case where $r=0$, it can be directly verified that the unique solution to the ODE system satisfies $M=0$ in the connection condition. Also, by solving the system explicitly and plugging the solution into (F.2), one gets

$$\theta^*=\beta\frac{\bar{x}}{\bar{s}}-\frac{\mu}{\sigma^2}\frac{\bar{x}}{s}\left[\frac{x-\bar{x}}{\bar{x}}-\beta\frac{s-\bar{s}}{\bar{s}}\right].$$

Table 7 Statistics for Slippage with Identical Daytime and Nighttime Volatilities

		Periodic-DN			Periodic-D			One-Day		
Λ	Statistics	(2x)	(-1x)	(-2x)	(2x)	(-1x)	(-2x)	(2x)	(-1x)	(-2x)
10^{-6}	Mean	0.5019	0.8119	1.4090	0.9691	1.1842	2.9134	1.7618	1.9088	5.2990
	Std dev	0.2918	0.3188	1.0024	0.7393	1.2960	4.5102	2.9638	2.4634	6.5547
10^{-5}	Mean	0.6749	0.8987	1.8636	1.0137	1.2181	3.0336	1.8714	2.0150	5.6400
	Std dev	0.5321	0.6729	1.7077	1.6467	0.8698	3.5561	3.0800	4.2939	7.1519
10^{-4}	Mean	1.0160	1.1585	2.8645	1.2194	1.3665	3.5591	3.3318	3.4765	10.2935
	Std dev	1.4879	1.1545	4.9152	1.3753	2.1473	3.4725	9.4185	4.7990	30.4733
10^{-3}	Mean	1.9430	2.0350	5.7328	2.0132	2.1090	5.9587	9.5162	11.0682	40.8286
	Std dev	3.3632	3.2548	7.1357	3.7753	2.2599	5.6642	36.3550	54.0954	843.5617

The slippage measure is defined as $D_i = |R_i^X - \beta R_i^S|$ and reported in bps. The statistics for simulation results are calculated over the daily slippage for 2520×20000 day (20000 sample paths, 2520 days per sample path), and the estimator is constructed as the 1% trimmed mean of the daily absolute deviations. The strategies are defined in Definition 2.

Plugging this into the dynamics of X, we have

$$d\left(\frac{X_t-\bar{x}}{\bar{x}}-\beta\frac{S_t-\bar{s}}{\bar{s}}\right)=-\frac{\mu}{\sigma^2}\left(\frac{X_t-\bar{x}}{\bar{x}}-\beta\frac{S_t-\bar{s}}{\bar{s}}\right)(\mu dt+\sigma dW_t).$$

Therefore,

$$\frac{X_t-\bar{x}}{\bar{x}}-\beta\frac{S_t-\bar{s}}{\bar{s}}\equiv 0,\quad \theta^*=\beta\frac{\bar{x}}{\bar{s}}.$$

G. The Case with Identical Daytime and Nighttime Volatilities

In Section 6, we compared the slippage performance for three strategies, where we assume the same daytime and nighttime expected returns but different daytime and nighttime volatilities. However, one may ask whether the performance is caused by the change of market return-risk characteristics overnight. To this end, we perform an alternative test, where we keep both the expected return and volatility the same during daytime and nighttime.

The result is shown in Table 7, showing that the Periodic-DN strategy still outperforms the Periodic-D strategy; more precisely, accounting for the overnight jump risk is still beneficial.

H. Return Deviation during Multi-Day Periods

Throughout the paper, we have focused on the slippage, the daily deviation of the fund return from the target return. There is a body of literature studying the return deviation in multi-day periods; see, e.g., Cheng and Madhavan (2009), Avellaneda and Zhang (2010), and Haugh (2011). For example, Avellaneda and Zhang (2010) derived an exact formula relating the multi-day fund return with the multi-day underlying return in a continuous-time approximation, which implies

$$\tilde{D}_T := \log\left(\frac{X_T}{X_0}\right)-\beta\log\left(\frac{S_T}{S_0}\right)=\int_0^T \eta_u du+\frac{\beta-\beta^2}{2}\int_0^T \sigma_u^2 du,$$

where η_u depends on the interest rate, borrowing cost, and fund management fees. In other words, over the multi-day period $[0,T]$, the difference $\tilde{D}_T$ between the fund's log return and β times the underlying log return increases in the realized variance, with a proportionality constant of $\frac{\beta-\beta^2}{2}$.

Table 8 Log Return Deviation $\tilde{D}_{60}$ under Optimal Strategy over 60-day Periods

σ	Statistics	(2x)	(-1x)	(-2x)
0.05	Mean	−0.5424	0.1729	0.2873
	Std dev	0.1028	0.1299	0.3335
0.10	Mean	−0.7260	−0.0036	−0.2501
	Std dev	0.1820	0.4458	1.1652
0.15	Mean	−1.0283	−0.3049	−1.1573
	Std dev	0.7030	0.5900	3.1124
0.20	Mean	−1.4519	−0.7284	−2.4302
	Std dev	1.6999	1.8502	4.7017

The 60-day log return deviation $\tilde{D}_{60}$ is defined as $\tilde{D}_{60} = \log\left(\frac{X_{60}}{X_0}\right) - \beta \log\left(\frac{S_{60}}{S_0}\right)$. Both the mean and standard deviation are reported in percentage. The statistics for simulation results are calculated over 2520×20000 day (20000 sample paths, 2520 days per sample path), and the estimator is constructed as the 1% trimmed mean of the daily $\tilde{D}_{60}$. The rebalancing is via the Periodic-DN strategy that considers both market closure and infinite horizon.

Since this proportionality constant is negative for $\beta = 2, -1, -2$, they predict that the LETF will under-perform the leveraged return on the index when volatility is high.

Although our model differs from this strand of literature as we focus on the daily slippage and consider more factors (e.g., interest rate, management fees, and market closure) in the optimal strategy, we still expect these results to hold approximately. We also have the following two additional empirical predictions:

(i) Multi-day log return deviation $\tilde{D}_T$ from our strategy should decrease in the (squared) volatility.

(ii) The rate of such decrease should be similar for $\beta = 2$ and $\beta = -1$ (for which $\frac{\beta - \beta^2}{2} = -1$) and higher for $\beta = -2$ (for which $\frac{\beta - \beta^2}{2} = -3$).

To investigate these two predictions, we consider the 60-day log return deviation $\tilde{D}_{60}$. We perform a simulation test similar to that in Panel B of Table 1 by calculating $\tilde{D}_{60}$ under the Periodic-DN strategy, for various values of volatility σ and multiple β.

The empirical result depicted in Table 8 meets our expectation. First, for any given β, $\tilde{D}_{60}$ decreases in σ. For instance, for (2x) LETFs, the deviation is -0.5424% for $\sigma = 0.05$ and increases to -1.4519% for $\sigma = 0.20$. This is consistent with our first prediction. Second, for (2x) fund, the difference between the deviation for $\sigma = 0.2$ and $\sigma = 0.05$ equals $-1.4519\% - (-0.5424\%) = -0.9095\%$, which is very close to -0.9013% for the (-1x) fund. In comparison, the difference for (-2x) fund is much higher at -2.7175%. This is consistent with our second prediction. More interestingly, -2.7175% is approximately three times -0.9095% or -0.9013%, which is consistent with $\frac{\beta - \beta^2}{2}$ for $\beta = -2$ is three times the value for $\beta = 2$ or -1.