

THE DYNAMICS OF LEVERAGED AND INVERSE EXCHANGE-TRADED FUNDS

Minder Cheng and Ananth Madhavan**

Leveraged and inverse Exchange-Traded Funds (ETFs) have attracted significant assets lately. Unlike traditional ETFs, these funds have “leverage” explicitly embedded as part of their product design. While these funds are primarily used by short-term traders, they are gaining popularity with individual investors placing leveraged bets or hedging their portfolios. The structure of these funds, however, creates both intended and unintended characteristics that are not seen in traditional ETFs. This note provides a unified framework to better understand the underlying dynamics of leveraged and inverse ETFs, their impact on market volatility and liquidity, unusual features of their product design, and questions of investor suitability. We show that the daily re-leveraging of these funds can exacerbate volatility towards the close. We also show that the gross return of a leveraged or inverse ETF has an embedded path-dependent option that under certain conditions can lead to value destruction for a buy-and-hold investor. The unsuitability of these products for longer-term investors is reinforced by the drag on returns from high transaction costs and tax inefficiency.¹



1 Introduction

Leveraged and inverse Exchange-Traded Funds (ETFs) provide leveraged long or short exposure to the daily return of various indexes, sectors, and asset classes. These funds have “leverage” explicitly embedded as part of their product design. The category has exploded since the first products were introduced in 2006, especially in volatile sectors such as Financials, Real Estate, and Energy. As of

end-January 2009, there are now over 106 leveraged and inverse ETFs in the US alone with Assets Under Management (AUM) of about \$22 billion.²

The space now comprises leveraged, inverse, and leveraged inverse ETFs offering $2\times$ or $3\times$ long exposure or short exposure of $-1\times$, $-2\times$, or $-3\times$ the underlying index returns. The most recent products authorized by the US Securities and Exchange Commission (SEC) offer the highest leverage factors. However, the bulk of AUM remains in $2\times$ leveraged products. Coverage has also expanded beyond equities and includes commodities, fixed

*Barclays Global Investors, N.A. 400 Howard Street, San Francisco, CA 94105.

income and foreign exchange. There is strong growth in this space outside the US as well. In addition, option contracts on leveraged ETFs have also gained in popularity. Leveraged and inverse mutual funds analogous to ETFs have also grown in popularity. Other than the fact that they offer investors liquidity at only one point in the day, the structure of these products is identical to leveraged and inverse ETFs and hence our analysis is fully applicable to these funds too.

Several factors explain the attraction of leveraged and inverse ETFs. First, these funds offer short-term traders and hedge funds a structured product to express their directional views regarding a wide variety of equity indexes and sectors. Second, as investors can obtain levered exposure within the product, they need not to rely on increasingly scarce outside capital or the use of derivatives, swaps, options, futures, or trading on margin. Third, individual investors—attracted by convenience and limited liability nature of these products—increasingly use them to place longer-term leveraged bets or to hedge their portfolios.

The structure of these funds, however, creates both *intended* and *unintended* characteristics. Indeed, despite their popularity, many of the features of these funds are not fully understood, even among professional asset managers and traders. This paper provides a unified framework to better understand some key aspects of these leveraged and inverse ETFs, including their underlying dynamics, unusual features of their product design, their impact on financial market microstructure, and questions of investor suitability.

Specifically, leveraged ETFs must re-balance their exposures on a daily basis to produce the promised leveraged returns. What may seem counterintuitive is that irrespective of whether the ETFs are leveraged, inverse or leveraged inverse, their re-balancing

activity is *always* in the same direction as the underlying index's daily performance. The hedging flows from equivalent long and short leveraged ETFs thus do not "offset" each other. The magnitude of the potential impact is proportional to the amount of assets gathered by these ETFs, the leveraged multiple promised, and the underlying index's daily returns. The impact is particularly significant for inverse ETFs. For example, a double-inverse ETF promising $-2\times$ the index return requires a hedge equal to $6\times$ the day's change in the fund's Net Asset Value (NAV), whereas a double-leveraged ETF requires only $2\times$ the day's change. This daily re-leveraging has profound microstructure effects, exacerbating the volatility of the underlying index and the securities comprising the index.

While a leveraged or inverse ETF replicates a multiple of the underlying index's return on a daily basis, the gross return of these funds over a finite time period can be shown to have an embedded path-dependent option on the underlying index. We show that leveraged and inverse ETFs are not suitable for buy-and-hold investors because under certain circumstances the long-run returns can be significantly below that of the appropriately levered underlying index. This is particularly true for volatile indexes and for inverse ETFs. The unsuitability of these products for longer-term investors is reinforced by tax inefficiency and the cumulative drag on returns from transaction costs related to daily re-balancing activity.

The paper proceeds as follows: Section 2 shows how leveraged and inverse ETF returns are related to those of the underlying index and provides an overview of the mechanics of the implied hedging demands resulting from the daily re-leveraging of these products. Section 3 explains the microstructure implications and resulting return drag from trading costs associated with hedging activity. Section 4 analyzes the longer-term return characteristics of these products and the value of the embedded

option within; and Section 5 summarizes our results and discusses their implications for public policy.

2 The Mechanics of leveraged returns

2.1 Producing leveraged returns

As leveraged returns cannot be created out of thin air, leveraged and inverse ETFs generally rely on the usage of total return swaps to produce returns that are a multiple of the underlying index returns. Futures contracts can also be used in addition to, or instead of, total return swaps. However, given their exchange-imposed standardized specification (to facilitate exchange-based trading and clearing), futures are not as customizable as total return swaps and are more limited in terms of index representation. In addition, basis risk is more significant with the futures than with total return swaps.³

Leveraged returns can also be produced by trading in physicals on margin. In other words, by borrowing the required capital in excess of its AUM, a leveraged ETF can invest in a properly levered position of the securities comprising the ETF's index benchmark. A negative implication of such an implementation strategy is that the financing cost will create a drag on the fund's performance with respect to its promised leveraged return. On the other hand, an inverse or leveraged inverse ETF can short the securities comprising the ETF's index benchmark and accrue interest income. Interestingly, a new breed of leveraged and inverse ETFs has recently emerged that are managed against customized index benchmarks. These benchmarks explicitly incorporate the financing cost (for leveraged ETFs) or accrued interest (for inverse and leveraged inverse ETFs) in index construction. Consequently, financing cost and accrued interest will not appear as a deviation against the funds' index benchmark.⁴ Throughout this paper, we will assume that leveraged and inverse ETFs rely on total return swaps to produce the promised leveraged

returns.⁵ Our findings remain unchanged regardless of how leveraged returns are produced, whether by trading in physicals on margin, equity linked notes, futures or other derivatives besides total return swaps.

Unlike traditional ETFs, leveraged and inverse ETFs can be viewed as pre-packaged margin products, albeit without any restrictions on margin eligibility. It is also worth noting that creations and redemptions for leveraged and inverse ETFs are in cash, while for traditional ETFs this is typically an "in-kind" or basket transfer.

2.2 Conceptual framework

Now we turn to the development of a unified conceptual framework to analyze inverse and leveraged ETFs. We will utilize a continuous time framework. All extant leveraged and inverse ETFs promise to deliver a multiple of its underlying benchmark's daily returns, so we will focus on the dynamics of the index and of the corresponding leveraged and inverse ETFs over a discrete number of trading days indexed by n where $n = 0, 1, 2, \dots, N$. Let t_n represent the calendar time of day n , measured as a real number (in years) from day 0. We assume $t_0 = 0$ initially, a convenient normalization. Note the frequency of n does not have to be daily. If there are leveraged or inverse ETFs designed to produce a multiple of the underlying benchmark's return over a different frequency (e.g., hourly, weekly, monthly, quarterly, etc.), then we can redefine n accordingly without any loss of generality.

Let S_t represent the index level which a leveraged or inverse ETF references as its underlying benchmark at calendar time t . Later, in Section 4, we will explicitly describe the continuous time process underlying the evolution of the index level, but for now let r_{t_{n-1}, t_n} represent the return of the underlying index from t_{n-1} to t_n , where

$$r_{t_{n-1}, t_n} = \frac{S_{t_n}}{S_{t_{n-1}}} - 1 \quad (1)$$

We will assume there are no dividends throughout to focus on the price and return dynamics without any loss of generality. Let x represent the leveraged multiple of a leveraged or inverse ETF. Therefore $x = -2, -1, 2$ and 3 correspond to double-inverse, inverse, double-leveraged and triple-leveraged ETFs.

2.3 Return divergence and path dependency

Later it will become clear that the exposures of total return swaps underpinning leveraged and inverse ETFs need to be re-balanced or re-set daily in order to produce the promised leveraged returns. In effect, these funds are designed to replicate a multiple of the underlying index's return on a *daily* basis. The compounding of these daily leveraged moves can result in longer-term returns, as expressed by:

$$\prod_{n=1}^N (1 + xr_{t_{n-1}, t_n}) \quad (2)$$

that have a very different relationship to the longer-term returns of the underlying index leveraged statically, as given by:

$$(1 + xr_{t_0, t_N}) \quad (3)$$

We can use a double-leveraged ETF ($x = 2$) with an initial NAV of \$100 as an example. It tracks an index that starts at 100, falls 10% on day 1 and

then goes up 10% on the subsequent day. Over the two-day period, the index declines by -1% (down to 90, and then climbing to 99). While an investor might expect the leveraged fund to decline by twice as much, or -2% , over the two-day period, it actually declines further by -4% . Why? Doubling the index's 10% fall on the first day pushes the fund's NAV to \$80. The next day, the fund's NAV climbs to \$96 upon doubling the index's 10% gain. This example illustrates the path dependency of leveraged ETF returns, a topic we return to more formally when we model the continuous time evolution of asset prices in Section 4.

2.3.1 Example: DUG and DIG

Real world examples of the effects noted above—and the confusion they cause among retail investors—are not difficult to find. The relation between short- and long-run performance of leveraged ETFs is well illustrated in the case of the $-2\times$ ProShares UltraShort Oil & Gas (DUG) and its $2\times$ long ProShares counterpart (DIG) that track the daily performance of the Dow Jones US Oil & Gas index. As shown in Figure 1, these funds are mirror images of each other over short periods of time, in this case a few trading days are in March. Over longer periods, however, the performance is materially different in the six month period as shown in Figure 2. Indeed between September 2008 and February 2009, *both* ETFs




Figure 1 DUG versus DIG (March 2–6, 2009)




Figure 2 DUG versus DIG (September 2008–March 2009).

were down substantially. These examples illustrate the path-dependency highlighted in the analysis.

2.4 Re-balancing and hedging demands

The re-balancing of inverse and leveraged funds implies certain hedging demands. Since extant funds promise a multiple of the day’s return, it makes sense to focus on end-of-day hedging demands. One benefit of modeling returns in continuous time, however, is that our analysis generalizes to any arbitrary re-balancing interval. Indeed, there has been recent discussion of new leveraged funds that track an index on a monthly versus daily basis.

Let A_{t_n} represent a leveraged or inverse ETF’s NAV at the close of day n or at time t_n . Corresponding to A_{t_n} , let L_{t_n} represent the notional amount of the total return swaps exposure that is required before the market opens on the next day to replicate the intended leveraged return of the index for the fund from calendar time t_n to time t_{n+1} . With the fund’s NAV at A_{t_n} at time t_n , the notional amount of the total return swaps required is given by:

$$L_{t_n} = x A_{t_n} \tag{4}$$

On day $n + 1$, the underlying index generates a return of $r_{t_n, t_{n+1}}$ and the exposure of the total return swaps, denoted by $E_{t_{n+1}}$, becomes:

$$E_{t_{n+1}} = L_{t_n}(1 + r_{t_n, t_{n+1}}) \tag{5}$$

$$= x A_{t_n}(1 + r_{t_n, t_{n+1}}) \tag{6}$$

At the same time, reflecting the gain or loss that is x times the index’s performance between t_n and t_{n+1} , the leveraged fund’s NAV at the close of day $n + 1$ becomes:

$$A_{t_{n+1}} = A_{t_n}(1 + x r_{t_n, t_{n+1}}) \tag{7}$$

which suggests that the notional amount of the total return swaps is required before the market opens on the next day to maintain constant exposure is:

$$L_{t_{n+1}} = x A_{t_{n+1}} \tag{8}$$

$$= x A_{t_n}(1 + x r_{t_n, t_{n+1}}) \tag{9}$$

The difference between Eqs. (6) and (9), denoted by $\Delta_{t_{n+1}}$, is the amount by which the exposure of the total return swaps that need to be adjusted or re-hedged at time t_{n+1} , as given by:

$$\Delta_{t_{n+1}} = L_{t_{n+1}} - E_{t_{n+1}} \tag{10}$$

$$= A_{t_n}(x^2 - x)r_{t_n, t_{n+1}} \tag{11}$$

2.4.1 Example: daily hedging demands

We can illustrate the above using an example that is built on the same case already discussed in Section 2.3. With an initial NAV of \$100 on day 0 for the double-leveraged ETF, the required notional amount of the total return swaps is \$200 (or 2 times \$100). As the index falls from 100 to 90 on day 1, the fund’s NAV drops to \$80 whereas the exposure of the total return swaps falls to \$180, reflecting a 10% drop of its value. Meanwhile the required

notional amount for the total return swaps for day 2 is \$160 (or 2 times \$80), which means the fund will need to reduce its exposure of total return swaps by \$20 (or \$180 minus \$160) at the end of day 1. And note $100 \times (2^2 - 2) \times 10\% = 20$. The dynamics of $S_{t_n}, A_{t_n}, L_{t_n}, E_{t_n}$ and Δ_{t_n} for $n = 0, 1$ and 2 are summarized in Table 1.

Tables 2 and 3 provide examples for an inverse ETF and a double-inverse ETF, respectively, with the same assumptions of the index's performance over two days. These examples highlight the critical role of the hedging term $(x^2 - x)$ in Eq. (11). This term is non-linear and asymmetric. For example, it takes the value 6 for triple-leveraged ($x = 3$) and double-inverse ($x = -2$) ETFs. As $(x^2 - x)$ is always positive (except for when $x = 1$ when

Table 1 Dynamics of a double-leveraged ETF ($x = 2, r_{t_0, t_1} = -10\%$ and $r_{t_1, t_2} = 10\%$).

	S_{t_n}	A_{t_n}	L_{t_n}	E_{t_n}	Δ_{t_n}
$n = 0$	100	100	200	—	—
$n = 1$	90	80	160	180	-20
$n = 2$	99	96	192	176	+16

Table 2 Dynamics of an inverse ETF ($x = -1, r_{t_0, t_1} = -10\%$ and $r_{t_1, t_2} = 10\%$).

	S_{t_n}	A_{t_n}	L_{t_n}	E_{t_n}	Δ_{t_n}
$n = 0$	100	100	-100	—	—
$n = 1$	90	110	-110	-90	-20
$n = 2$	99	99	-99	-121	+22

Table 3 Dynamics of a double-inverse ETF ($x = -2, r_{t_0, t_1} = -10\%$ and $r_{t_1, t_2} = 10\%$).

	S_{t_n}	A_{t_n}	L_{t_n}	E_{t_n}	Δ_{t_n}
$n = 0$	100	100	-200	—	—
$n = 1$	90	120	-240	-180	-60
$n = 2$	99	96	-192	-264	+72

the funds are not leveraged or inverse), the reset or re-balance flows are always in the same direction as the underlying index's performance. So, when the underlying index is up, additional total return swaps exposure must be added, but when the underlying index is down, the exposure of total return swaps needs to be reduced. This is intuitively clear for a leveraged long fund, is always true whether the ETFs are leveraged, inverse or leveraged inverse. Why is the effect the same for funds that are short the index? Intuitively, an inverse or leveraged inverse fund's NAV will increase if the index falls, which requires it to increase short exposure still further, generating selling pressure. In other words, there is no offset or "pairing off" of leveraged long and short ETFs on the same index, which is why the re-balance flows are in the same direction between Table 1 for a double-leveraged ETF and Tables 2 and 3 for an inverse and a double-inverse ETF on day 1 and day 2, respectively.

Note that the need for daily re-hedging is unique to leveraged and inverse ETFs due to their product design. Traditional ETFs that are not leveraged or inverse, whether they are holding physicals, total return swaps or other derivatives, have no need to re-balance daily. We discuss the implications of daily re-balances in the next section.

3 Market structure effects

3.1 Liquidity and volatility near market close

As leveraged and inverse ETFs gather more assets, the impact of their daily re-leveraging on public equity markets is raising concerns. Many commentators have cited leveraged and inverse ETF re-balancing activity as a factor behind increased volatility at the close.⁶ Of course, other factors might also account for heightened volatility at the close. For example, traders might choose to trade later in the day because there are more macroeconomic announcements earlier in the day or because

the price discovery process takes longer. Nevertheless, as we show in this section, there are good theoretical and empirical grounds to support the argument that daily re-leveraging by leveraged and inverse ETFs contributes to volatility.

In theory, re-balancing activity should be executed as near the market close as possible given the dependence of the re-balancing amount on the close-to-close return of the underlying index as per Eq. (11). Whether leveraged and inverse ETFs re-balance their exposure of total return swaps immediately before or after the market close, however, their counterparties with which they execute total return swaps will very likely put on or adjust their hedges while the market is still open to minimize the risk to their capital and position taking. So, as leveraged and inverse ETFs gain assets, there likely will be a heightened impact on the liquidity and volatility of the underlying index and the securities comprising the index during the closing period (e.g., the last hour or half-hour) of the day's trading session.

The magnitude of the potential impact from re-balancing activity is proportional to the amount of assets gathered by leveraged and inverse ETFs in aggregate, the leveraged multiple promised by these funds, and the underlying index's daily return. Such predictable and concentrated trading activity near the market close may also create an environment for "front running" and market manipulation.⁷ Merely substituting total return swaps with alternative instruments such as trading in the physicals on margin, futures, equity-linked notes or other derivatives will have the same economic impact as long as the product design is based on taking leveraged positions in these instruments.

3.2 Order flow and prices

Stochastic variation in hedging demands will influence volatility and return dynamics of the underlying index and the securities comprising the index.

To understand this issue, we can extend the above conceptual framework to model the stochastic dynamics of the market impact induced by the hedging demand or daily re-balancing flows.

In the canonical microstructure framework, the price movement over a trading interval is modeled as a function of net order flow, capturing both inventory and information effects. At the end of day n , the hedging demand of leveraged and inverse ETFs is almost fully predictable, and hence it is reasonable to assume that this has a same-day effect on the close. We assume a portion $0 \leq \phi \leq 1$ of the ETF hedging flow contributes to price impact on day n (at the close) and $1 - \phi$ spills over to the following day. So, if $\phi = 1$, the entire impact of the hedging demand is realized at the close; if $\phi = 0$, the price effect is fully realized only on the successive day.

We model the price movement as:

$$S_{t_{n+1}} - S_{t_n} = \lambda(q_{t_n, t_{n+1}} + \phi \Delta_{t_n} + (1 - \phi) \Delta_{t_{n-1}}) + w_{t_n, t_{n+1}} \quad (12)$$

Here $\lambda > 0$ is the price impact coefficient (which in turn can be thought of as an increasing function of volatility and decreasing function of average daily volume), and $q_{t_n, t_{n+1}}$ is the signed order flow from market participants, excluding the hedging demand from leveraged and inverse ETFs. The stochastic error term $w_{t_n, t_{n+1}}$ captures the effect of unpredictable news shocks or noise trading. The hedging demand itself depends on the day's price change:

$$\Delta_{t_{n+1}} = a_{t_n}(x^2 - x)(S_{t_{n+1}} - S_{t_n}) \quad (13)$$

where $a_{t_n} = A_{t_n}/S_{t_n}$. Substituting and rearranging, we get

$$\begin{aligned} S_{t_{n+1}} - S_{t_n} &= \frac{\lambda(q_{t_n, t_{n+1}} + (1 - \phi) \Delta_{t_{n-1}}) + w_{t_n, t_{n+1}}}{1 - \lambda \phi a_{t_n}(x^2 - x)} \end{aligned} \quad (14)$$

where we require that $\lambda \phi a_{t_n}(x^2 - x) < 1$ for equilibrium to be defined. This expression makes

clear the magnified impact of hedging demands on overall price movements by increasing the market impact coefficient for *all* flows, irrespective of their source. Intuitively, hedging demand provides an additional momentum effect to same-day returns that increases the price pressure effect of any signed order imbalance regardless of its source.

The price impact is greater with higher volatility, lower market liquidity, higher same-day effects, increased AUM, and higher leverage ratios. The presence of the lagged hedge term $\Delta_{t_{n-1}}$ (when $\phi < 1$) induces serial correlation in returns because the previous period's hedge is linearly related to the previous return. Thus, the analysis shows that hedging flows in leveraged and inverse ETFs can exacerbate same-day volatility, add momentum effects, and induce serial correlation in returns.

3.3 Aggregate hedging demand

Recent volatility highlights the effect of hedging demands at the close. How large are these effects? To simulate the aggregate impact, we examined 83 leveraged and inverse ETFs on various US equity indexes including S&P 500, 400, and 600 indexes, Nasdaq QQQ, Russell 1000, 2000, and various

sub-indexes/sectors including Financials, Oil & Gas, Real Estate, Materials, etc. This sample is largely the universe of domestic equity leveraged funds, excluding commodity and currency leveraged funds. Ranked by notional value (as of January 2009) the top three ETFs are the ProShares Ultra S&P 500 (SSO), UltraShort S&P 500 (SDS), and the ProShares Ultra Financials (UYG) with notional AUM of \$3.0 billion, \$2.9 billion, and \$1.7 billion, respectively. Also included are the new 3 \times and -3 \times funds from Direxion covering the Russell 1000 (BGU and BGZ) and other indexes.

Figure 3 shows the percentage of AUM for the universe of US equity leveraged and inverse ETFs, broken down by leverage factor ($x = -3\times, \dots, 3\times$) for end-January, 2009. As of end-January 2009, the sample comprises 84 equity leveraged and inverse ETFs with a total of \$19 billion in AUM, of which \$1.4 billion or 7.3% has $x = 3$ or -3 , a category that did not exist in 2007.

For the above same 84 US equity leveraged and inverse ETFs, Table 4 reports details on various statistics of interest by leverage factor ($x = -3\times, \dots, 3\times$) as of January 2009.




Figure 3 AUM in US Equity Leveraged and Inverse ETFs by Leverage Factor (x)—January 2009.

Table 4 Summary statistics on US equity leveraged and inverse ETFs (January, 2009).

Leverage factor (x)	-3	-2	-1	2	3
Bid-Ask Spread (basis points)	10.5	35.3	11.8	42.4	16.8
At/Within Quotes (%)	64.3	81.5	81.5	88.0	68.5
Depth (\$ '000s)	45.7	249.2	87.2	134.2	33.3
Order Size (\$ '000s)	49.3	67.8	23.5	17.8	23.7
Expense Ratio (basis points)	81.7	89.4	95.0	89.4	81.7
Holding Period (days)	1.2	6.7	8.6	15.2	1.7

Several points are worth noting. Although bid-ask spreads are reasonably small, their economic impact is large given the short holding periods. Most transactions occur at or within the quotes. Liquidity is generally good, especially in double-leveraged products, where the ratio of depth (i.e., mean quote size in dollars) to order size (in dollars) is high.

Management fees and other costs are large relative to most traditional ETFs and do not vary significantly by leverage. Average holding periods (calculated as the ratio of month-end shares outstanding to average daily share volume for the month) are shortest for the most levered products, but are still significant for double-leveraged ETFs (15.2 days) which account for the bulk of AUM. Some of the most recent and active leveraged funds (e.g., Direxion Daily Financials $-3\times$ Bear (FAZ)) have holding periods as short as 0.2 of a trading day. Note that these are averages over heterogeneous investor populations. It is quite likely that holding periods vary by investor category. For example, sophisticated hedge funds may employ leveraged ETFs for within day bets while retail investors may have longer holding periods.

For each ETF, indexed by $i = 1, \dots, K$, we compute the notional hedge as before and hence the aggregate value of hedging demand at the end of day n can be expressed as a function of the hypothetical return as $\Delta_{t_{n-1}}(r_{t_{n-1}, t_n})$ where:

$$\Delta_{t_{n-1}}(r_{t_{n-1}, t_n}) = \left(\sum_{i=1}^K A_{i, t_{n-1}} (x_i^2 - x_i) \right) r_{t_{n-1}, t_n} \quad (15)$$

where $A_{i, t_{n-1}}$ is the AUM of ETF i at the close of day $n - 1$ or at time t_{n-1} .

The Table 5 shows the function $\Delta_{t_{n-1}}(r_{t_{n-1}, t_n})$ (in millions of dollars) for returns to the corresponding stock indexes from 1% to +15% as of end February 2009. We assume that all sub-indexes had the same return, which is an obvious simplification but not unreasonable given the highly correlated nature of the current market. Also shown is the hedging demand as a percentage of median market-on-close volume. We use the median as volume is heavily skewed by end-of-month flows, quarterly rebalances, and option expiration cycles, but the result is very similar if the mean were used instead. Note that since we are essentially computing a weighted average, the market-on-close demand in value terms is just a linear function of return. However, the demand as a fraction of closing volume is a non-linear function of return because the re-balancing demand is added to the denominator too. So, for example, a broad 1% move in the US equity market would result in additional MOC demand of about 17%, while a 5% move is associated with demands equal to 50% of the close.

As AUM changes and more levered and ultra short products gain traction, we would expect this slope

Table 5 $\Delta_{t_{n-1}}(r_{t_{n-1}, t_n})$ in \$ million.

r_{t_{n-1}, t_n}	1%	5%	10%	15%
$\Delta_{t_{n-1}}(r_{t_{n-1}, t_n})$	787	3.934	7.873	11.811
% MOC volume	16.8%	50.2%	66.9%	75.2%

to steepen. Indeed, the amount of hedging has been increasing. Using each ETF's AUM for December 31, 2008 and December 31, 2007, we compute the end-of-day dollar flows for a 5% return. This yields end-of-day flows of \$3.5 and \$2.1 billion, respectively. Observe that the hedging demand corresponding to a 5% move has already risen to \$3.9 billion, almost double from year-end 2007. These are likely conservative estimates as we exclude re-balancing flows from leveraged and inverse mutual funds.

3.4 Regression analysis of closing period volatility

Now we turn to a more direct test of the prediction that re-balancing flows exacerbate return movements towards the close. As leveraged funds typically will re-balance towards the end of the day, we focus on returns in last hour of the trading day for large capitalization stocks. Specifically, we look at re-balancing flows arising from 2× and 3× leveraged and inverse funds (SSO, SDS, BGR, and BGRZ) on large capitalization indexes. The ETF aggregate re-balancing flows are then computed as above for each day. This is likely a conservative estimate because there is also large cap re-balancing flow from other leveraged and inverse ETFs (e.g., SKF) on narrower benchmark indexes that contain large cap stocks.

We model the closing period returns as a function of the day's return and signed re-balancing flows. Using daily data over the period January 1, 2006 to mid-April 2009, we then estimate the following regression equation:

$$r_{t^-,t_n} = \beta_0 + \beta_1 r_{t_{n-1},t_n} + \beta_2 \Delta_{t_n} \quad (16)$$

where t_n^- denotes the calendar time of 3 p.m. Eastern time on day n , $r_{t_n^-,t_n}$ is the intra-day S&P 500 return from 3 p.m. Eastern time to the close on day n , r_{t_{n-1},t_n} is the close to close S&P 500 return, and Δ_{t_n} is the signed aggregate re-balance flow (in billions of dollars). Equation (16) has a straightforward economic interpretation: β_1 captures the

contribution to the day's return of the closing period. The coefficient β_2 captures the additional impact of re-balancing flows on market movements in the closing period over and above the daily return movement.

The regression estimates yield $\beta_1 = 0.249$ and $\beta_2 = 0.541$ with t -statistics of 5.806 and 2.503, respectively, and $R^2 = 0.35$ with 825 degrees of freedom. The coefficient β_0 is economically and statistically insignificant (0.01) as we would expect in a model of signed returns without directional bias. The closing period is 15.4% of the total trading day, but is a higher fraction (about 25%) of the day's volume and hence of the day's return as implied by the estimates. The estimate of β_2 suggests that after controlling for the day's return movement, \$1 billion of re-balancing flow adds 0.54% to the closing period's return.

Since concerns about re-balancing flows are often framed in terms of volatility, we also examine the second moment of closing period returns. Specifically, we want to see if there is an empirical relation between absolute returns in the closing period and the magnitude of re-balancing flows. It is possible that more volatile days will see greater movements at the close, so we use the daily value of the VIX index as a control variable. Accordingly, we run the following regression using daily data from January 1, 2006 to mid-April 2009:

$$|r_{t_n^-,t_n}| = \beta_0 + \beta_1 V_{t_n} + \beta_2 |\Delta_{t_n}| \quad (17)$$

where on day n , $|r_{t_n^-,t_n}|$ is the absolute value of the closing period return (defined as the intra-day S&P 500 return from 3 p.m. Eastern time to the close), V_{t_n} is the value of VIX at the close of day n , and $|\Delta_{t_n}|$ is the absolute value of aggregate re-balance flow (in billions of dollars).

We estimate $\beta_1 = 0.023$ and $\beta_2 = 0.724$ with t -statistics of 12.186 and 6.422, respectively. The coefficient β_0 is small in value, -0.18 , with a t -statistics of -4.526 . Model fit is good; $R^2 = 0.457$

with 825 degrees of freedom. The coefficient on β_1 is no surprise; higher volatility means greater return movements in the closing period. Controlling for overall volatility, the significantly positive coefficient β_2 confirms that greater imbalances add to closing volatility. These results also hold with other measures of closing volatility or with the dependent variable expressed as a fraction of VIX. Clearly, the analysis above can be extended in many ways including the examination of other indexes and countries, but this lies outside the scope of this paper.

4 Return dynamics

Having discussed the mechanics of leveraged and inverse ETFs and market structure implications of daily re-balancing activity, we turn now to an analysis of returns over longer periods of time. Our objective is to understand return dynamics and the role of the embedded option within the leveraged and inverse product to address questions concerning the suitability of these products for individual investors, particularly those with longer holding periods. To do this, we now need to model explicitly the evolution of security prices in continuous time.

Specifically, the index level S_t is assumed to follow a geometric Brownian motion, as given by:

$$dS_t = \mu S_t dt + \sigma S_t dW_t \quad (18)$$

with a drift rate of μ and a volatility of σ . Here, W_t in Eq. (18) is a Wiener process with a mean of zero and a variance of t . Then $\ln(S_t)$ follows a generalized Wiener process, as given by:

$$d \ln(S_t) = \left(\mu - \frac{\sigma^2}{2} \right) dt + \sigma dW_t \quad (19)$$

We need to relate the dynamics of the NAV of a leveraged or inverse ETF to the dynamics of the ETF's underlying index level. From Eq. (7), it can be shown algebraically that:

$$\frac{A_{t_{n+1}} - A_{t_n}}{A_{t_n}} = x \left(\frac{S_{t_{n+1}} - S_{t_n}}{S_{t_n}} \right) \quad (20)$$

Since Eq. (20) holds for any period, when the time interval between t_n and t_{n+1} is sufficiently small, this expression becomes:

$$\frac{dA_t}{A_t} = x \frac{dS_t}{S_t} \quad (21)$$

where A_t and S_t represent the NAV of a leveraged or inverse ETF and the ETF's underlying index level, respectively, at time t . Note that as described earlier, a leveraged ETF hedges at discrete time intervals but always maintains economic exposure at $x \times$ the underlying index return, as represented in Eq. (21). With Eqs. (21) and (18), it can be shown that:

$$dA_t = x\mu A_t dt + x\sigma A_t dW_t \quad (22)$$

and then $\ln(A_t)$ follows a generalized Wiener process, as given by:

$$d \ln(A_t) = \left(x\mu - \frac{x^2\sigma^2}{2} \right) dt + x\sigma dW_t \quad (23)$$

suggesting that A_t follows a geometric Brownian motion with a drift rate of $x\mu$ and a volatility of $x\sigma$. In other words, the volatility of a leveraged or inverse ETF is just x times the volatility of its underlying index. This is handy for pricing options on leveraged and inverse ETFs. Such options are gaining popularity and volume in parallel with activity in their underlying ETFs, and there are now over 150 different option contracts traded on these instruments.

4.1 Leveraged ETFs and the underlying index

The various relationships derived in a continuous-time setting, i.e., Eqs. (21)–(23), now can be applied to our earlier analysis at the daily level. Applying Eq. (19) to S_{t_N} and S_0 gives us the following relationship of the index level on day N and day 0, or between time t_N and time 0:

$$S_{t_N} = S_0 \exp \left(\left[\mu - \frac{\sigma^2}{2} \right] t_N + \sigma \sqrt{t_N} z \right) \quad (24)$$

where $\exp(z) = e^z$ is the exponential function and z is a standard normal distribution with a mean of 0

and a standard deviation of 1. So, the return to the index over the period, $\ln(S_{t_N}) - \ln(S_0)$, is normally distributed with mean $[\mu - \frac{\sigma^2}{2}]t_N$ and standard deviation $\sigma\sqrt{t_N}$. Likewise, applying Eq. (23) to A_{t_n} and A_0 will give us the following relationship of the leveraged or inverse ETF's NAV on day N and day 0:

$$A_{t_N} = A_0 \exp\left(\left[x\mu - \frac{x^2\sigma^2}{2}\right]t_N + x\sigma\sqrt{t_N}z\right) \quad (25)$$

In other words, the return to the leveraged ETF ($\ln(A_{t_N}) - \ln(A_0)$) is normally distributed with mean $[x\mu - \frac{x^2\sigma^2}{2}]t_N$ and standard deviation $x\sigma\sqrt{t_N}$. These draws are not independent, however, because the same realization of the sample path, captured by z , is in both returns.

We can further simplify this as follows:

$$A_{t_N} = A_0 \exp\left(\left[x\mu - \frac{x\sigma^2}{2}\right]t_N + x\sigma\sqrt{t_N}z + \frac{(x - x^2)\sigma^2 t_N}{2}\right) \quad (26)$$

$$= A_0 \left(\frac{S_{t_N}}{S_0}\right)^x \exp\left(\frac{(x - x^2)\sigma^2 t_N}{2}\right) \quad (27)$$

So, the total return of a leveraged or inverse ETF over a N -day period (i.e., (A_{t_N}/A_0)) is the x th power of the gross return of the ETF's underlying index over the same time period (S_{t_N}/S_0), multiplied by a scalar $\exp(\frac{(x-x^2)\sigma^2 t_N}{2})$. As $(x - x^2) < 0$, the scalar is less than 1. As shown in Section 2, leveraged and inverse ETFs are designed to replicate a multiple of their underlying index's return on a daily basis. As it turns out, ignoring the scalar in Eq. (27) the compounding of these daily leveraged moves is economically equivalent to compounding the gross return of the index to the x th power. Note that the risk free rate does not enter into Eq. (27) because this is implicit in the costs of the total return swaps.

4.2 Embedded option

Equation (27) also suggests that the NAV of a leveraged or inverse ETF will always be positive.⁸ This is in contrast to the risk implied by Eq. (3) that a statically leveraged fund may see its NAV completely wiped out, for example, if over a N -day period its underlying index has doubled for a short fund, or the index has lost half of its value for a double-leveraged fund. The preservation of capital is made possible because, unlike a statically leveraged fund which establishes its target leveraged notional amount initially and keeps it unchanged, a leveraged or inverse ETF dynamically adjusts its leveraged notional amount of total return swaps on a daily basis depending on the performance of its underlying index. In effect, leveraged and inverse ETFs have an embedded option on the underlying index.⁹

Given the dynamic nature of the leveraged and inverse ETFs' daily re-balancing activity, the value of the embedded option is dependent on not only S_{t_N} but also the entire path of the underlying index prior to reaching S_{t_N} at time t_N . Table 6 provides a simple example. It is based on the same double-leveraged ETF as in Table 1. As in Table 1, the index ends at 99 on day 2, however, it takes a different path in Table 6 to get there: The index is flat on day 1 and down 1% on day 2. Note that the NAVs on day 2 (A_{t_2}) are different between Tables 6 and 1.

The effect of such volatility is captured in the scalar of $\exp(\frac{(x-x^2)\sigma^2 t_N}{2})$ in Eq. (27). As $(x - x^2)$ is negative for the leveraged and inverse ETFs, the scalar

Table 6 Dynamics of a double-leveraged ETF ($x = 2$, $r_{t_0,t_1} = 0\%$ and $r_{t_1,t_2} = -1\%$).

	S_{t_n}	A_{t_n}	L_{t_n}	E_{t_n}	Δ_{t_n}
$n = 0$	100	100	200	–	–
$n = 1$	100	100	200	200	0
$n = 2$	99	98	196	198	–2

is always less than 1 and an ideal path is one with lower volatility (σ) and shorter time horizon (t_N). For example, as the index has demonstrated a significantly lower volatility over the two-day period in Table 6 than in Table 1, note the NAV on day 2 for the double-leveraged ETF in Table 6 is higher than that in Table 1 (\$98 vs. \$96).

Of course, in a practical context, we need to consider additional factors. First, actual payoffs are going to depend on the realized sample paths of the random variables making up the N -period returns. Second, we need to distinguish between calendar and trade days. Third, these expressions assume the management and other fees for the leveraged and inverse ETFs is the same as their un-levered counterpart, whereas they are typically much higher. Further, as shown above, leveraged, inverse, and leveraged inverse ETFs are designed to replicate a leveraged exposure to the underlying index on any one day, but not necessarily over extended periods of time. A leveraged or inverse ETF's return beyond a day is determined by the index returns and volatility, leverage (financing) costs, expenses, and the time horizon. The inability of leveraged and inverse ETFs to track index returns over periods longer than a day is not widely understood by retail investors in these funds. Note: The leveraged ETF's manager must re-invest at the close of each day without actually observing the close, which induces further slippage.

4.3 Value destruction of buy-and-hold

This above analysis helps us to understand questions of investor suitability. While leveraged and inverse ETFs do offer limited liability, this comes at a cost. It was mentioned before that unlike a statically hedged fund, leveraged and inverse ETFs contain an embedded option embedded in the product design. Recall from Eq. (27) that the N -period realized gross return of a leveraged or inverse ETF is just the gross return of its underlying index over that period

raised to the x th power, then multiplied by the scalar $\exp(\frac{(x-x^2)\sigma^2 t_N}{2})$. This scalar gives rise to the value of the embedded option in the first place, but is also the source—under certain conditions—of long-run erosion in value.

To see this, note that the scalar decreases with t_N and asymptotically approaches zero as t_N approaches infinity. Along any sample path of returns, the gross return of the leveraged or inverse ETF (relative to the underlying index return) will decline with the duration of holding and the volatility of the underlying index. This problem is exacerbated for higher leveraged (i.e., x) ETFs. In other words, the longer the buy-and-hold period of a leveraged or inverse ETF, the lower the ETF's long-term return relative to x times the underlying index return. For example, holding a triple-leveraged ($x = 3$) or double-inverse ($x = -2$) ETF for five years ($t_N = 5$) when the volatility of the index returns is 50% ($\sigma = 50\%$), the scalar becomes 2.3%.

Of course, the total return to a leveraged or inverse ETF along any sample path could greatly exceed the underlying index return even if the holding period is relatively long. Note that the first term in Eq. (27) reflects the realized sample path of index returns. If the mean return is large and positive and $x > 0$, then this term can imply substantial payoffs to the leveraged ETF holder over many sample realizations. Consequently, it is of interest to examine the expected value of a longer-term investment. This expression can be explicitly derived using the moment generating function of the lognormal distribution. Specifically, if Z is a random variable that is distributed normally with mean m and standard deviation s , it can be shown that the expected value of $\exp(kZ)$ is

$$E[\exp(kZ)] = \exp\left(km + \frac{k^2 s^2}{2}\right). \quad (28)$$

Recall the N -period index gross return $\ln(S_{t_N}) - \ln(S_0)$ has a mean $[\mu - \frac{\sigma^2}{2}]t_N$ while the

corresponding x -times leveraged or inverse ETF has a mean $[\mu x - \frac{x^2\sigma^2}{2}]t_N$. Then, using Eq. (28) with m equal to the corresponding mean, the expected gross return of the index at time t_N is $\exp([\mu - \frac{\sigma^2}{2}]t_N + \frac{\sigma^2 t_N}{2}) = \exp(\mu t_N)$, while the expected value of the leveraged or inverse ETF is $\exp(x\mu t_N)$. Intuitively, this result is a reflection of Jensen's inequality and the convexity of the exponential function. It follows immediately that if $x\mu < 0$ (i.e., an inverse ETF with positive index drift or a leveraged ETF with negative index drift) that the expected value of a leveraged or inverse ETF will tend to zero, i.e., a scenario of long-term value destruction.

It is important to understand that the expected value derived above is a mean value; it is affected by the small probabilistic weight placed on extreme sample paths. Because the lognormal distribution is positively skewed, the simple expectation is not likely the typical experience of the average investor. Consider a simple numerical example: suppose index returns are equally likely to be 4% or -2%. Over a successive two-day period, the four possible cumulative returns are 8.16%, 1.92%, 1.92%, and -3.96%. The average two-day cumulative return is thus 2.01%, but only one sample path (i.e., successive daily 4% returns) of the four exceeds this figure. Since the skewness of the lognormal distribution increases with variance, the mean value is especially unrepresentative for volatile assets. Consequently, it makes sense to focus on the typical or median investor's value change over an interval of time. Again, from the properties of the lognormal distribution, if Z is distributed normally with mean m and standard deviation s , the median value of $\exp(Z)$ is $\exp(m)$, which does not depend on the variance.

Formally, substituting $m = (x\mu - \frac{x^2\sigma^2}{2})t_N$ for the mean return in a given interval of time t_N , we can now explicitly characterize situations in which leveraged or inverse ETFs have a negative median

return:

1. $x < 0$ (e.g., inverse and leveraged inverse ETFs) and $\mu > 0$, i.e., the index drift is positive;
2. $x > 0$ (e.g., leveraged long ETFs) and $0 < \mu < \frac{x\sigma^2}{2}$; and
3. $x > 0$ and $\mu < 0$, i.e., negative drift of index returns.

In case 1, the returns to inverse ETF holders (i.e., $x < 0$) over long periods of time should be negative because we expect $\mu > 0$ in equilibrium so the first term in Eq. (27) will also tend to zero as t_N increases. In other words, the long-run inverse and leveraged inverse ETF value is zero. In case 2, the variance term dominates μ so the leveraged ETF has a negative drift. An interesting special case is where the index return has positive median returns but the corresponding leveraged ETF is negative. Then, for $x > 0$, if μ satisfies: $\frac{\sigma^2}{2} < \mu < x\sigma^2$, the long-run return of the leveraged or inverse ETF is negative despite a positive drift in the index. Case 3—the polar opposite to case 1—seems unlikely in a long-run equilibrium, but it is included for completeness.¹⁰

In summary, theory shows that leveraged and inverse ETFs are generally not suitable for buy-and-hold investors.¹¹ Indeed, it is possible for the median investor to face a negative return drift in a leveraged or inverse ETF although the underlying index has a positive drift. But while levered products can yield less than x times the index return, their volatility is exactly x times that of the index. From a long-run equilibrium perspective, the performance drag is most problematic for inverse ETFs and products based on volatile underlying indexes, as the return scalar also decreases with σ .

4.4 Frictions

From a practical standpoint, the cumulative impact of transaction costs can dramatically affect the

returns to a long-run investor in these ETFs. It is useful to distinguish between explicit costs (e.g., management fees) and implicit costs (e.g., dealer hedging costs and the impact costs arising from daily re-balancing activity) which represent a further drag on longer-term performance.

Explicit costs are easy to account for as they are a constant fraction of the fund's AUM. Total expense ratios for leveraged and inverse ETFs range from 75–95 basis points, which are high relative to traditional ETFs. But while explicit costs are transparent, the hidden performance drag caused by implicit costs is not readily quantified and is not well understood by investors.

Our previous discussion in Section 2.4 and examples provided in Tables 1, 2 and 3 did not factor transaction cost. But as leveraged and inverse ETFs re-balance their exposure of total return swaps on a daily basis, they will incur transaction costs and the cumulative effect of the transaction costs on the ETFs' performance cannot be under-estimated given the daily nature of the funds' re-balancing activity. In particular, these leveraged and inverse ETFs should expect to pay fully or at least partially the bid-ask spread of adjusting their total return swaps exposures. In addition, there may be market impact cost if the re-balancing amount on a given day demands more liquidity than the market is able to absorb in an orderly manner near the close. As mentioned in Section 3.1, the leveraged and inverse ETFs are always buying when the underlying index is up, and selling when the index is down. They will not only incur higher transaction cost because they almost always demand liquidity when they re-balance daily, but also likely create a performance drag in the long run.

We can model the impact of implicit transaction costs within the framework already developed. Let us begin with the costs of establishing the swap position. These hedging costs are visible to the fund (e.g., dealer spreads, etc.) but not to its investors.

At the beginning of each day, the fund pays a dealer a total transaction cost C_{t_n} proportional to the dollar amount of the total notional swap exposure. Let θ denote the hedging cost (dealer's fee) as a constant fraction of the notional amount of total return swaps required, where $0 \leq \theta < 1$, so that $C_{t_n} = \theta E_{t_n}$. Net of costs, the exposure in Eq. (6) is just scaled by $(1 - \theta)$, i.e., $E_{t_{n+1}} = (1 - \theta)x A_{t_n}(1 + r_{t_n, t_{n+1}})$. Note that the risk free rate is implicit in θ , although it is negligible at the daily level. The total notional swap exposure is directly reduced by transaction costs (much like a tax) so that the gain or loss is $(1 - \theta)x$ times the index's performance between t_n and t_{n+1} . So, Eq. (7) becomes

$$A_{t_{n+1}} = A_{t_n}(1 + (1 - \theta)xr_{t_n, t_{n+1}}) \quad (29)$$

If we define $x' = (1 - \theta)x$, we see that the analysis of Section 4 goes through fully by merely substituting x' for x throughout. Since $|x'| < |x|$, dealer transaction costs act as a performance drag on the fund, reducing the promised returns.¹²

Daily fund re-balancing also imposes a further performance shortfall from what is promised, as given by Eq. (2), due to market impact costs. We assume the market impact cost is proportional to the size of the re-balancing trade each period. The expected market impact cost, denoted by M_{t_n} , is modeled as proportional to the absolute hedging demand $|\Delta_{t_n}|$:

$$M_{t_n} = E[\lambda(x^2 - x)|r_{t_{n-1}, t_n}|] \quad (30)$$

where $\lambda > 0$ is the market impact coefficient. Define by $\tau = t_n - t_{n-1}$ the fixed re-balancing interval, typically a day. From our previous results, the daily return r_{t_{n-1}, t_n} is distributed normally with standard deviation $\sigma\sqrt{\tau}$. Assuming for convenience that the daily mean return is negligible and using the results for the folded normal distribution (i.e., the distribution of the absolute value of a normal variate), the expected daily impact

cost is

$$M_{t_N} = \lambda(x^2 - x)\sigma\sqrt{\frac{2\tau}{\pi}}. \quad (31)$$

This cost will be most evident over many days as the cumulative impact of these costs is compounded and is manifested in the form of slippage from the benchmark index return.

From Eq. (31), we see that market impact costs increase with volatility, σ , the market impact parameter, λ , the leverage factor, x , and the time between re-balances, τ . We can use this analysis to understand the effects of changing the re-balancing frequency (e.g., to twice a day or once a week) on the underlying economics. Interestingly, moving to less frequent re-balancing can mitigate market impact costs. To see this, suppose we increase the re-balancing interval to $\tau' = 2\tau$, i.e., every two trading days. Then, using Eq. (31), the expected re-balancing impact cost for a two-day re-balancing interval is 44% higher than the daily expected cost, not double the daily value. Intuitively, impact costs are proportional to the standard deviation of returns, which in turn increases with $\sqrt{\tau}$, the square-root of the calendar re-balancing interval. So less frequent re-balancing implies less absolute hedging and hence lower costs for a given calendar period. It is worth emphasizing though that while a change in re-balancing frequency can affect implicit costs, it does not alter the economics of the longer-run performance, as can be seen if t_N is interpreted not as days but as periods with μ and σ scaled appropriately.

Taxes represent another source of friction. Leveraged and inverse ETFs are not designed with respect to tax efficiency and an investor's after-tax performance might be significantly lower than a strategy using leverage with less turnover. At year-end 2008, for example, the major inverse ETF providers (Rydex and ProShares) paid out substantial capital gains tax distributions ranging from 4% to 86%. The impact of taxes, however, will vary depending on each individual investor's particular circumstances.

4.4.1 Example: UltraShort financials ETF

Empirically we observe leveraged and inverse ETF returns and index returns, but not the implicit costs of trading. We should, however, see a deterioration in longer-term leveraged and inverse ETF returns relative to the index returns if these costs are significant. As an example, consider SKF, the ProShares UltraShort Financials, a double-inverse ETF based on the Dow Jones US Financial Services index, that has been the subject of considerable controversy. Using two years of daily data from February 1, 2007 to February 1, 2009, we regressed the daily return of SKF on the daily return on IYF, the iShares Dow Jones Financial Services Index ETF. Figure 4 shows the plot of the 502 close-to-close returns of SKF against IYF.¹³ In theory, we have $x = -2$ but the regression has a slope of -1.94 with a very high $R^2 = 0.98$. As a further check, we find that the first order autocorrelation of both products is similar (-0.07 for SKF and -0.09 for IYF), indicative of small temporary price reversals. With $x' = 1.94$, we estimate $\theta = 0.03$ or 3% as our measure of transaction cost and impact. This relatively high cost estimate is perhaps not surprising given the high volatility of the financial sector in this period. At the monthly level over the period November 2008 to February 2009, we see a sharp deterioration in tracking ability. The slope coefficient falls to -1.75 and the goodness of fit drops with $R^2 = 0.76$. This deterioration is consistent with the cumulative effect of market impact costs as noted earlier. Indeed, over the period November 2008 to the end of March 2009, *both* SKF and IYF declined by -25.1% and -36.35% , respectively, as shown in Figure 5.

5 Conclusions and policy implications

As noted in the introduction, the class of leveraged and inverse ETFs continues to expand rapidly. The growth is especially rapid in more leveraged and leveraged inverse products. From the above analysis,




Figure 4 Daily SKF and IYF returns (February 2007–2009).




Figure 5 SKF and IYF cumulative returns (November 2008–March 2009).

however, it is precisely these products that pose the greatest challenges for public policy from both the perspective of individual investor protection and the larger goal of ensuring efficient and orderly markets.

Specifically, leveraged and inverse ETFs generally rely on the usage of total return swaps to produce returns that are a multiple of the ETFs' underlying index returns. As these ETFs are in effect designed to replicate a multiple of the underlying index's return on a daily basis, their exposures of total return swaps need to be re-balanced daily in

order to produce the leveraged returns promised. Whether the ETFs are leveraged, inverse or leveraged inverse, their re-balancing activity is always in the same direction as the underlying index's daily performance: When the underlying index is up, the additional exposure of total return swaps needs to be added; when the underlying index is down, the exposure of total return swaps needs to be reduced. The cumulative effect of the transaction costs on the leveraged and inverse ETFs' performance cannot be underestimated given the daily nature of the funds' re-balancing activity. The re-balance flows are executed as near the market close as possible given

their dependence on the close-to-close return of the underlying index. The need for daily re-balancing is unique to the leveraged and inverse ETFs due to their product design. Traditional ETFs that are not leveraged or inverse, whether they are holding physicals, total return swaps or other derivatives, have no such need for daily re-hedging.

From a market structure perspective, many broker-dealers believe the end-of-day flows needed to maintain the leverage promised by leveraged and inverse ETFs is having a heightened impact on liquidity and volatility of the underlying index and the securities comprising the index at the close. Inverse ETFs in particular have a magnified hedging demand relative to their long counterparts (e.g., as shown above, a $2\times$ inverse ETF has the same hedging multiplier as a $3\times$ leveraged long ETF).¹⁴ The above analysis shows leveraged and inverse ETFs amplify the market impact of all flows, irrespective of source. Essentially, these products require managers to “short gamma” by trading in the same direction as the market. There is a close analogy to the role played by portfolio insurance in the crash of 1987.¹⁵ The magnitude of the potential impact is proportional to the amount of assets gathered by these ETFs, the leveraged multiple promised, and the underlying index’s daily returns. Such predictable and concentrated trading activity near the market closes may also create an environment for front running and gaming.

These issues suggest regulators need to carefully consider the impact of point-in-time liquidity demands when authorizing leveraged products, especially leveraged inverse products or products on relatively narrow, volatile indexes. Regulators, moreover, cannot easily track the implicit exposure in a leveraged or inverse ETF (this is not the case if the underlying ETF itself is levered), making it difficult to assess the market structure impact of these products as their AUM changes over time. Developing a mechanism to track notional exposure through total

return swaps would alleviate some of the concerns arising from the short gamma nature of the hedging strategy. Better disclosure might also help mitigate broader counterparty credit risks. Fund companies are not required to disclose the details of their swap agreements, making it difficult to determine counterparty credit quality. If counterparties are not willing or unable to provide total return swaps, the leveraged ETF manager must replicate these returns using derivatives, which can be costly and induce further tracking error.

Alternatives to existing products could be based on changing the re-balancing frequency. However, shorter re-balancing (e.g., twice a day) will magnify intra-day price movements while extending the re-balancing frequency (e.g., to weekly) increases the risk of the fund being wiped out (i.e., as the fund becomes gradually similar to a statically hedged structure). As long as product design is based on taking leveraged positions, the underlying economics of these instruments is not altered.

Our results also raise questions of investor suitability. Some investors—despite very clear language in the prospectuses—do not understand the point clearly that leveraged and inverse ETFs will not necessarily replicate the leveraged index return over periods longer than a day.¹⁶ Generally, the greater the holding period and the higher the daily volatility, the greater the deviation between the leveraged ETF’s return and a statically levered position in the same index. Also note that as leverage is embedded in the product, individual investors can gain additional leverage by buying these products on margin. Better education, margin restrictions, and tighter requirements on investor eligibility are possible options regulators could consider. For example, some retail investors may be holding leveraged funds in accounts that are not approved for margin, either because of lack of documentation or in accounts (e.g., retirement) that explicitly exclude leverage.

Our analysis of return dynamics provides conditions under which buy-and-hold investors experience eventual value destruction in leveraged or inverse ETFs, a point not well understood even by experts. The gross return of a leveraged or inverse ETF over a finite time period can be shown algebraically to be simply the gross return of the ETF's underlying index over the same period raised to the power of the leveraged multiple of the ETF, multiplied by a scalar that is less than one. The NAV of a leveraged or inverse ETF is non-negative (in contrast to the risk that a statically leveraged fund may see its NAV completely wiped out) because a leveraged or inverse ETF dynamically adjusts its leveraged notional amount of total return swaps on a daily basis depending on the performance of its underlying index. In effect, the leveraged and inverse ETFs have an embedded path-dependent option on the underlying index. The value of the option is dependent on the entire path of the underlying index. Specifically, if volatility is sufficiently high, the median investor will experience a long-run erosion in value in a leveraged or inverse ETF.

Other considerations also might have a material impact on longer-term investors. Transaction costs in leveraged, inverse, and leveraged inverse ETFs may be higher than those of replicating leveraged long or short exposure directly through swaps, options, futures, or trading on margin. In addition, the cumulative impact of transaction costs arising from daily re-balancing activity will only reduce these funds' NAV further. There are also important tax consequences, particularly with inverse ETFs. In essence, the leveraged and inverse ETFs are not suitable for buy-and-hold investors. Neither these products are designed to deliver long-term performance for volatile indexes.

Notes

¹ The views expressed here are those of the authors alone. The strategies discussed are strictly for illustrative and

educational purposes and should not be construed as a recommendation to purchase or sell, or an offer to sell or a solicitation of an offer to buy any security. We thank Mark Coppejans, Matt Goff, Allan Lane, Hayne Leland, J. Parsons, Heather Pelant, Ira Shapiro, Jim Smith, Mike Sobel, Richard Tsai and an anonymous referee for their helpful comments.

- ² This figure does not include leveraged mutual funds that have the same characteristics as leveraged ETFs except for the fact that they are not traded intraday.
- ³ Basis risk refers to the risk associated with imperfect hedging, possibly arising from the differences in price, or a mismatch in sale and expiration dates, between the asset to be hedged and the corresponding derivative.
- ⁴ See, for example, Dow Jones STOXX Index Guide (2009). This structure is more common in Europe.
- ⁵ In addition to swaps, leveraged funds' assets typically include a pool of futures contracts to manage liquidity demands and reduce transaction costs. Some long leveraged funds listed in the US are guided by a no-action letter from the SEC that stipulates a minimum (e.g., 80% or more) of assets should be in the underlying. However, this does not apply to the inverse funds.
- ⁶ See, e.g., Lauricella, Pulliam, and Gullapalli (2008), who note: "As the market grew more volatile in September, Wall Street proprietary trading desks began piling onto the back of the trade knowing that the end-of-day ETF-related buying or selling was on its way. If the market was falling, they would buy a short ETF and short the stocks or the market some other way. If the market was rallying, they would buy a bull fund and go long."
- ⁷ Illegal front running refers to a broker executing orders for its own account before filling pending client orders that are likely to affect the security's price. Here, we use the term more colloquially to refer to trading ahead of the flows that are highly predictable given public information on index returns and fund AUM that need not be illegal.
- ⁸ See also Despande, Mallick, and Bhatia (2009).
- ⁹ This distinction also applies to levered active products and hedge funds. These funds rely on static hedges that do not require daily re-levering. The daily re-balancing activity of these funds is driven mainly by alpha changes and risk management, and their trades do not have to be concentrated near the market close.
- ¹⁰ In a general equilibrium framework, it is possible for a risky asset to have a long-run expected return below the risk free rate if it acts as a hedge (e.g., with an embedded put), but this would seem unlikely for an equity index.

- ¹¹ See also Lu, Wang, and Zhang (2009) who reach a similar conclusion from an empirical analysis of longer term leveraged ETF returns.
- ¹² See, for example, Leland (1985) for an analysis of the impact of volatility on expected transaction costs in the context of option pricing and replication.
- ¹³ Source: Bloomberg, with closing prices and returns adjusted for corporate actions.
- ¹⁴ These arguments could also extend to the Over-the-Counter (OTC) counterparts of leveraged ETFs and to levered non-exchange traded funds.
- ¹⁵ See, e.g., Grossman and Vila (1989). The working paper underlying this article circulated prior to the Crash of October 1987.
- ¹⁶ See, e.g., Zweig (2009), who provides several examples of leveraged funds failing to track the underlying index. Zweig notes, however, that: “Still, many financial advisors believe these funds are a good long-term hedge against falling markets.”

Dow Jones STOXX Index Guide (2009). “*Dow Jones STOXX Index Guide*.” STOXX and Dow Jones, Inc.

Grossman, S. and Vila, J.-L. (1989). “Portfolio Insurance in Complete Markets: A Note.” *Journal of Business* 62(4), 473–476.

Lauricella, T., Pulliam, S. and Gullapalli, D. (2008). “Are ETFs Driving Late-Day Turns? Leveraged Vehicles Seen Magnifying Other Bets; Last-Hour Volume Surge.” *Wall Street Journal*, December 15, C1.

Leland, H. (1985). “Option Pricing and Replication with Transactions Costs.” *Journal of Finance* 40(5), 1283–1301.

Lu, L., Wang, J. and Zhang, G. (2009). “Long Term Performance of Leveraged ETFs.” Working paper, Baruch College.

Zweig, J. (2009). “How Managing Risk with ETFs Can Backfire.” *Wall Street Journal*, February 28, B1.

Keywords: Leveraged and inverse ETFs; market microstructure

References

- Despande, M., Mallick, D. and Bhatia, R. (2009). “Understanding Ultrashort ETFs.” Barclays Capital Special Report.