Precidian's Proposed ETF and the Possibility of Reverse Engineering

July 2015

Prepared by

Ricky Alyn Cooper, PhD (rcooper3@iit.edu)
IIT-Stuart School of Business
Chicago, IL

Purpose and Scope of this Research

The overall purpose of this research is to determine the ease or difficulty a person would encounter in attempting to reverse engineer the portfolio weights of the proposed Precidian actively managed exchange traded fund. This research consisted of:

1) Modeling the essential elements of the proposed Precidian ETF construction methodology;

2) Using this modeled construction methodology, and relevant statistical techniques, determine if the previous night's closing portfolio weights could be reconstructed (these are weights on which the intra-day quotes are based);

3) Using this modeled construction methodology, and relevant statistical techniques, determine if the time series of estimated portfolio weights could be used to forecast future changes in the weights.

This research did not attempt to model the effect of active trading on the predictability of returns. Also, this research used 1 second bar data. This means that the pricing information of all assets were synchronous, based on the last trade or quote within a second, as constructed from historical data. The real time quotes being released would be based on data sampled from within a second and would necessarily have asynchronous quotes. This source of noise would cause an unknown effect in size, but in direction certainly would make the discovery of the reverse engineering process more difficult.

Finally, this research also assumed that changes in the attractiveness rankings of stocks were not predictable. That is, we assumed that the "alpha model" of the active manager is unknown to the outside world.

Summary of the Results

The major findings of this research were as follows:

1) The system of price quotation proposed by Precidian does not allow one to successfully recover the previous night's weights from the time series of quotes and asset prices;

2) The lesser reason for the previous result are that Precidian will quote the ETF price using the mid-spread of the top of book bid and ask for each constituent security each second, as opposed to the last traded price;

3) The greater reason for the previous result is that Precidian is scaling the price of the ETF to an initial value of $20. This scaling dampens price volatility of the ETF and makes statistical techniques less effective at reverse engineering the constituents;

4) The changes in the portfolio weights from day to day may not be forecasted accurately for realistic portfolio construction techniques, making knowledge of the next day's target weights unobtainable;

5) The previous result is a function of the fact that given Precidian plans to use a proprietary model to change its portfolio holdings, meaning changes in holdings will appear random.

6) Additionally, the noise in the time series of forecasted weights is such that small predictabilities in the changes in portfolio weights are not discernible in the final series of estimated weights.

The conclusion drawn from these findings is that it is not possible to ascertain with any degree of accuracy the previous day's portfolio weights, nor is it likely that anyone will, with any great accuracy, be able to ascertain the changes in portfolio weights being executed by the portfolio manager.

Facts about Precidian ETF Modeled , Data Employed, and Methods Used

The Precidian ETF being considered is an actively managed ETF, with a proprietary stock selection system, and a non-algorithmic, opportunistic trading system. The ETF will offer pricing once per second, based on the previous night's portfolio weights, and the mid-spread price of each of its constituent securities. The price will be scaled to equal $20 at the beginning of trading. The ETF will divulge its actual holdings once per quarter with a delay of approximately two months.

In constructing this test certain simplifying assumptions were made to model the facts stated above. These assumptions are designed to make the experiment tractable, but also made in such a way that any bias in the test should be towards discovery of the underlying ETF portfolio. The following paragraphs document the specifics of the data and methods employed.

The universe of stocks from which the portfolio is chosen was fixed as being from the NASDAQ-100 universe as currently composed, this universe is smaller than any currently envisioned universe for the proposed ETF. Specifically the data analyzed consisted of 1-second spaced observations, from 9:30:01 to 16:00:00 for each of the 100 stocks of the current active top of book bid price, top of book ask price, and last traded price. This data was analyzed for 44 trading days from 05/01/2014 to 07/02/2014. Mid-spread prices were created for each security using the formula msp=(Bid price + Ask price)/2.

In order to construct portfolios a fake attractiveness, or "alpha", score was created for each stock each day. The alpha score was constructed to guarantee a certain correlation between the score of one day and the next. This construction guaranteed a certain level of continuity among the portfolio weights from one day to the next. Specifically, each day's stock forecast was generated using the formula

$$\alpha_t = \rho\alpha_{t-1} + \sqrt{1-\rho^2}N(0,1)$$

where rho is the desired level of daily auto-correlation and N(0,1) is a random variable. The first day's forecast was simply 100% of N(0,1). To test various scenarios alpha forecasts were created for six different correlations— 0.8, 0.85, 0.90, 0.95, 0.99, and 0.999. The random noise takes as given that the alpha process is unknowable to the outside observer at face value. This research did not attempt to model the actual alpha process and Precidian did not reveal any details of it in conjunction with this research project

The portfolio construction methodology was to simply hold the top 40 stocks at an equal weight of 2.5% each day and the other 60 at 0%. The discovery techniques were not given this information, since the number of actual assets held will vary through time. In practice, however the simple structure of this construction was likely easier to discover than the actual portfolios in a live ETF. Trading was not modeled. Since, the reported portfolio weight is based on the previous night's closing weight, the modeling of trading was not required. However, leaving out opportunistic trading also skewed the results towards discovery of predictability in the pattern of weights.

Statistics and Results

Table 1 gives turnover of each of the six alpha processes studied. All tests were run on all of the alpha models, though all results are only presented where turnover played a role in the analysis. Only alphas 5 and 6 represent realistic turnover numbers, though the others will be useful for illustrating a later result.

Table 1

Daily Turnover Implied by the Six Portfolio Alphas

	Alpha 1	Alpha 2	Alpha 3	Alpha 4	Alpha 5	Alpha 6
correlation	0.8	0.85	0.9	0.95	0.99	0.999
Turnover	0.28	0.22	0.18	0.12	0.06	0.01

The first battery of tests involved using restricted least squares estimation. In order to test the role of Precidian's method of reporting ETF prices in determining the ease of reverse engineering the underlying portfolio, Table 2 reports three different experiments and a naïve control. In the first test the dependent variable is the one second ETF price calculated as the portfolio weighted sum of the last traded price of each of the stocks. The independent variables are the last traded price of each of the 100 stocks. The regressions are run daily and the coefficients are restricted to be between 0 and .1, as well as to sum to 1 across all securities. The number .1 was chosen as an upper bound since it was considered safe to assume there would be some diversification in the portfolio. The accuracy of the regression is reported using the square root of the mean square error in the weights averaged across the 44 days. That is, the reported statistic is:

$$\sqrt{\frac{1}{100\times 44}\sum_{t=1}^{44}\sum_{i=1}^{100}\left(\hat{w}_{i,t}-w_{i,t}^{p}\right)^{2}}$$

where w-hat is the restricted least squares parameter estimate of stock i for day t and w-p is the effective portfolio weight for the day. The results were the same regardless of alpha model chosen, since the results do not depend on portfolio changes or turnover. However, these specific numbers are reported for the fifth alpha model.

Notice that in this first case the error is extremely low, meaning that with complete price information and high decimal precision, the ETF constituent weights are easy to reproduce.

Table 2
Average Errors of the Potfolio Weights

Least squares regressions on 44 days of pricing, weights restiructed between 0 and .1
Erorrs reported based on estimated obtained at end of day

Test	Ave. Error	Max. Error
I. Unrounded ETF price reported based on last traded price of stocks	0.0000	0.0000
II. Rounded ETF price reported based on mid-spread prices of stocks	0.0012	0.0188
III. Scaled ($20) and Round ETF price reported based on mid-spread prices of stock	0.0104	0.0764
IV. Control-- Naïve equal weighted index	0.0122	0.0150

The second regression test uses mid-spread prices to report the ETF prices, but still does not scale the ETF price. The price reported for the ETF is the portfolio weighted mid-spread price of the underlying stocks. Also, the final ETF price is rounded to the nearest penny. The independent variables are switched to the unrounded mid-spread price of the stocks to maximize the chance of the regression finding the correct answers. As the table shows, the results are worse than in the first case, but not terrible. The square root of mean square error of .0012 is about 12% of the magnitude of the average portfolio weight (.01). A 12% error is not a great amount, though it makes building a replicating portfolio difficult without more information.

The reason for this result is two-fold. Primarily, mid-spread prices do not convey as much information as last traded prices. Last traded prices bounce from bid to ask and these bounces correlate with small changes in the price of the ETF. However, the mid-spread prices do not bounce this way. Furthermore, any change in price of the ETF is now rounded so that any change less than half a penny goes unnoticed.

The third regression test is like the second except ETF price is scaled to begin the day at $20. This scaling greatly diminishes the precision of the regression. For example, the actual unscaled ETF price is approximately $100. This means a five cent move must occur in the unscaled ETF price to create a one cent move in the scaled price. When combined with rounding to the nearest penny this make reverse engineering very difficult. This is confirmed as the square root of mean square error is .0104. This is an average error of 100% of the average weight. Moreover, the maximum error found is an extremely large value of .0764.

This result is comparable to the error of a naïve tracking portfolio. For the purposes of this experiment an appropriate naïve tracking portfolio is the equal weighted NASDAQ-100 portfolio for which it is easily derived analytically the square root of mean square error in the weights is .0122. This implies that reverse engineering is likely to achieve the same success by the end of the day as a naïve index tracking strategy would achieve. Or, put another way, the naïve portfolio has as much information at the start of the day as the reversed engineered portfolio does at the end.

One caution is that some attention must be paid to the scaling of the ETF. The information dampening effect described above is tied to the average price of the underlying holdings versus the reported price of the ETF. A universe of stocks with a much lower average price than the NASDAQ-100 should have the EFT scaled to a lower price. This is not really an issue with most large capitalization US indices, especially given the other sources of noise inherent in the price reporting process, but it is a point to keep in mind if the ETF expands into many different universes.

Bayesian Analysis on the Data

Some consideration was given to applying a Bayesian regression methodology on this data set, and indeed a Bayesian, independent normal gamma regression using Gibb's sampling was run on several days with the result that the estimation did not improve. This is not unexpected because for the NASDAQ-100 universe the number of seconds per day creates a data matrix with enough degrees of freedom that the data overwhelms any prior of less than 100% certainty (which of course means one is ignoring the data because he or she already knows the answer). In fact, feeding the true portfolio weights at the beginning of the day as a prior to the Bayesian regression (with an estimated prior variance of only 0.0001) did not substantially help the estimation error obtained in the scaled ETF regression. This is because the error is inherently "cooked into" every quote provided. The inherent imprecision in the data creates uncertainty that overwhelms even a perfect Bayesian prior. The way to think about this result is that the reporting mechanism for prices in of the ETF makes it impossible to confirm statistically any intuition one has about the portfolio weights.

Time Series Properties of the Weight Estimates

The purpose of this section is to investigate the role of the ETF reporting mechanism in masking any auto-correlation in the portfolio changes. Given the nature of the portfolio construction used in this investigation one would expect there to be negative auto-correlation in the portfolio weights. Since the possible portfolio weights are only 2.5% or 0%, there is a non-zero probability of the weights reverting to their other possible value but no chance of the weight continuing upwards or downwards. That is, the simple portfolio weighting scheme built for this paper has negative autocorrelation similar to the bid-ask bounce in stock prices. There is no reason to expect a real portfolio construction to have auto-correlation but the purpose of this section is to examine to what extent the fuzzy estimation of portfolio weights presented in the last section would, or would not, mask the actual auto-correlation in weights of the portfolios.

Table 3 presents statistics for the five portfolios constructed for this test. The table reports the number of stocks out of 100 that had significant auto-correlations (at 5%) in weight changes for the 44 days tested, and the number of parameter estimates that had significant auto-correlations for the day tested. The first conclusion is that the number of stocks with significant auto-correlation decreases as turnover decreases. This is because, given the construction method of this paper, the more random fluctuation allowed in the alpha model, the more bouncing between 0% and 2.5% there will be in the portfolio. Only, the last alpha models have realistic turnover in them. For Alpha 6, even with the stylized portfolio construction, only 12 stocks out of 100 have significant auto-correlation.

Table 3

Number, out of 100 of Stocks with Significant Auto-correlation in Portfolio Weight Changes

Results computed for 44-days, six different alpha models, significance at 5%

	Alpha 1	Alpha 2	Alpha 3	Alpha 4	Alpha 5	Alpha 6
actual weights	55	44	37	35	22	12
estimated weights	27	28	22	19	14	6

However, the real focus of this table is the decline in significance in the number of stocks with significant auto-correlation in the weights when the parameter estimates are used. At every level of turnover the amount of significance in auto-correlated stocks diminishes tremendously. The effect is strong enough that, for the lowest turnover alpha model (Alpha 6), even with this portfolio construction methodology there would not appear to be a more than random number of significantly auto-correlated stocks (at 5% significance level one could expect up to 10 stocks to have statistical significance, without there actually being any significant auto-correlation).

The conclusion is that the fuzziness of the portfolio weight estimation will mask any small level of auto-correlation in weights that the portfolio construction might induce. The effect appears to be strong enough that any likely real world construction system will have random appearing time series properties in its portfolio weights.

Conclusion

This research studied the Precidian ETF construction methodology using a small universe of securities, a stylized portfolio construction and the price quote methodology they propose. The primary result is that given three components of their proposed reporting methodology—mid-spread pricing, rounding to the nearest penny, and scaling the price to $20 at the start—the underlying portfolio weights are not recoverable with any great degree of accuracy. The main driver of this result is the scaling.

The other result of this research is that it is highly unlikely that anyone will discern any pattern in the time series of portfolio weight changes for any likely real world portfolio construction because the fuzziness in the daily weight estimation masks all but very large auto-correlated weight changes.

The summary conclusion is that it seems rather unlikely that the Precidian ETF construction methodology will result in a product that can be reverse engineered for purposes of front running, or that can be tracked with an engineered portfolio better than a simple naïve index portfolio can track.

Ricky (Rick) Cooper

Assistant Professor of Finance

Stuart School of Business, Illinois Institute of Technology

Education:

Ph.D., Vanderbilt University, Finance

M.B.A., Vanderbilt University, Finance

B.S., University of Chicago, Mathematics (with General Honors)

Biography:

Ricky "Rick" Cooper is Assistant Professor of Finance at IIT Stuart School of Business. Throughout his academic career, Dr. Cooper has taught all aspects of financial theory, investments, and corporate finance at Wayne State University, Harvard University, and Vanderbilt University.

Dr. Cooper began his professional career with State Street Global Advisors in Boston, MA, where he quickly rose from Active International Portfolio Manager, to co-founder of the Enhanced Index Group, to co-founder and Associate Director of the Advanced Research Center.

He then returned to his hometown of Chicago, where he worked as Senior Partner and Director of Analytics for Harris Investment Management. In this role, he modernized the analytic systems, and led the revamping of the models with a commensurate uptick in investment performance. Dr. Cooper also spent several years as Owner and Chief Investment Officer of his own firm. He currently serves as a research consultant and Director of Risk Management for Xambala, Inc.'s proprietary high frequency trading systems.

Dr. Cooper's research has been published in The *Journal of Futures Markets*, *The Financial Analyst's Journal*, *The Journal of Financial Economics*, and several other books and journals. Dr. Cooper has been a speaker at numerous conferences, and has been quoted in both the *Wall Street Journal* and *Crain's Chicago Business*.

Affiliated Programs:

M.S. in Finance

Ph.D. in Management Science