Overlapping ETF: Pair trading between two gold stocks

Bell, Peter N and Lui, Brian and Brekke, Alex

University of Victoria

1 April 2012

Overlapping ETF: Pair trading between two gold stocks

Brian Lui, Alex Brekke, Peter Bell[1]

University of Victoria Economics Term Paper[2]

1. Introduction

For our purposes, an exchange-traded fund (ETF) is an investment fund that tracks an index and can be traded at high frequency. When multiple ETF are based on one index then the percentage changes in net asset value for each ETF should be equal; we refer to these ETF as 'overlapping ETF'. For example, two ETF based on the SP500 index (SPY and IVV) trade very closely to each other except for structural difference in treatment of dividends. This paper proposes a type of mispricing for overlapping ETF and calculates the profit and loss based on trading this mispricing-strategy.

The purpose of this paper is to propose a trading strategy for overlapping ETF and calculate the profitability using real price data. For two overlapping ETF that are designed to provide the same intraday percentage change, the difference in percentage changes is a measure of mispricing. This mispricing is the central focus of the paper. The premise of the paper is that mispricing can take large positive or negative values, but it will always come back to zero. This assumption reflects our view that ETF are generally priced correctly but will occasionally deviate.

It is possible to discuss this trading in terms of the Efficient Market Hypothesis, in particular the 'impossibility paper' by Grossman and Stiglitz (1980) that argues markets cannot be efficient unless someone is monitoring prices and incurring costs to make them efficient. It is our view that overlapping ETF can be priced well or poorly; if they are priced well, there will be few opportunities for the trading described here. If they are priced poorly, then there will be more opportunities. Our results show that some small mispricing persists at all times. This mispricing is barely tradeable, approximately the size of the bid/ask spread, which suggests that the ETF are priced well. However, we also show that it is possible to achieve reliable profit by trading only against large mispricing. This shows that large mispricings do occur periodically and there is an incentive for professional traders to trade against inefficient prices, which is the logic of the impossibility paper in action.

[1] - Corresponding author, pnbell@uvic.ca

[2] - Econ 435 Financial Economics, April 2012, Dr Paul Schure

2. Introduction of Data

In this section we introduce the data we are working with and demonstrate the concept of mispricing. The overlapping ETF that we chose to work with are two gold ETF: IQQ SPDR Gold Trust (GLD) and iShares Gold Trust (IAU). Both ETF track the performance of an index of gold futures prices. We use these stocks because they are heavily traded and provide a simple example of the overlapping ETF concept. Since both are heavily traded, we assume that we could trade 100-share lots without disrupting price.

We use price data from February 13, 2012 recorded from IQ Feed (http://www.iqfeed.net/). The tick data is manipulated to provide a new observation on bid/ask prices every 10 seconds while the market is open. To simplify the explanation we refer to 'price' as the midpoint price for each stock (average of bid and ask). At each time step, we calculate the percentage change of the new price relative to the opening price for each stock. This is an important measure because the percentage changes for overlapping ETF should be equal. We show the percentage changes for the first few minutes of the trading day in Figure 1.

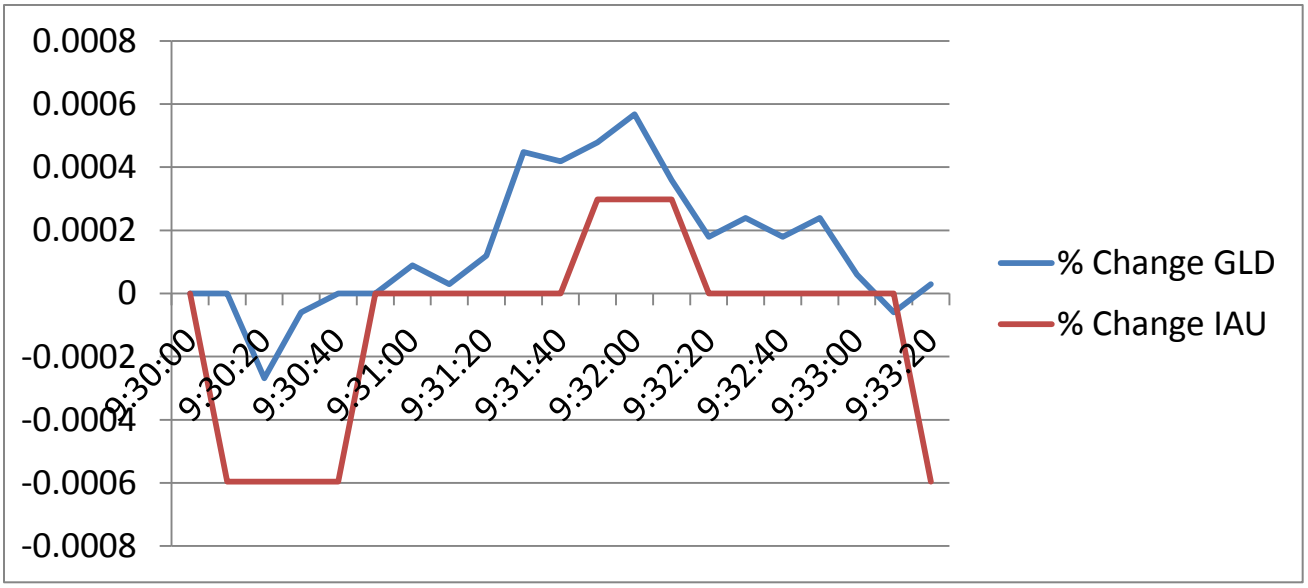

Figure 1: Percentage change for two overlapping ETF

The percentage changes of the ETF in Figure 1 show that we are on the right track. The ETF are moving roughly in the same direction, at the same time. But there are some small differences in the movements of each ETF. The purpose of our research is to determine if these small differences are sufficiently large that they can be traded on for profit.

We define the mispricing as: percentage change in GLD – percentage change in IAU. Based on this convention, we calculate the mispricing of the two ETF and report the results in Figure 2.

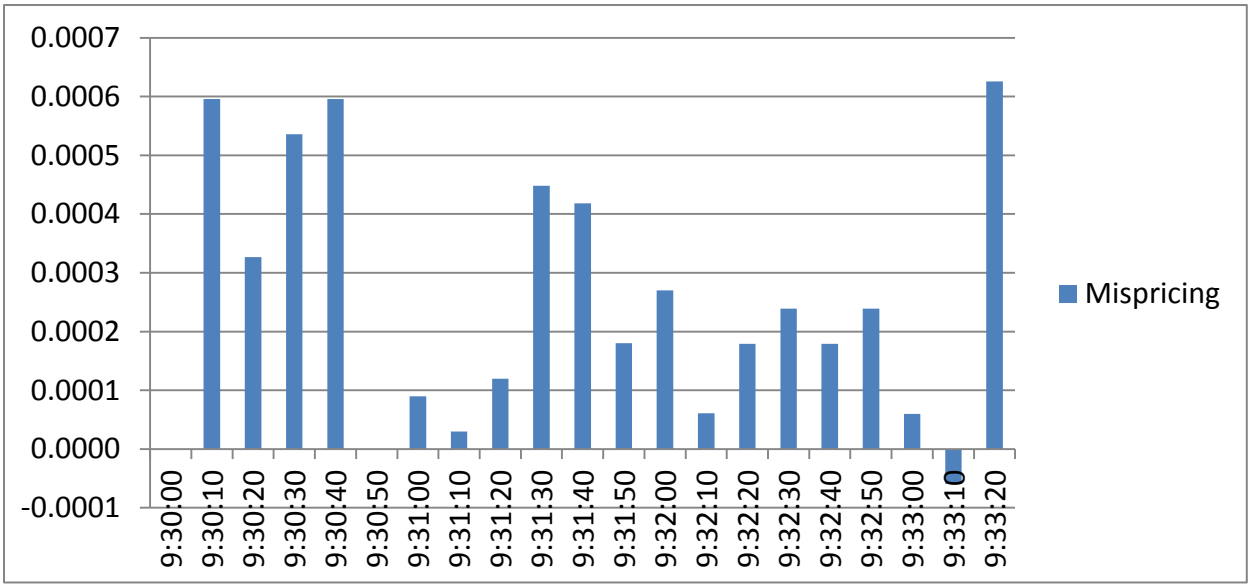

Figure 2: Mispricing between two overlapping ETF

The results shown in Figure 2 are exciting because it shows that mispricing has special behaviour: sometimes the mispricing takes large values, at other times it settles close to zero. This is the type of behaviour that suggests trading will be profitable. Our trading will attempt to enter into trades when the mispricing is large and exit when the mispricing is small. Since the mispricing is always very small, it seems that the ETF is being priced well. However, it seems that there may be money still on the table.

3. Trade Set Up

To illustrate an arbitrage opportunity when there is a mispricing, suppose that, in the first hour of trading, the gold commodity index rises 5%, the price of IAU rises by 5%, and the price of GLD rises by 8%. Suppose that the opening prices for the day were the 'correct' prices and they provide a good reference point to calculate percentage changes. It follows that GLD is overpriced, which would cause a large positive mispricing in our framework. One way to trade

this mispricing is to short GLD and buy IAU in equal dollar amounts and hold until the mispricing had 'corrected'. For example, in the next hour of trading if IAU fell by 1% but GLD fell by 4% then this trade would earn profit. The two ETFs are designed to give the same return on intraday time scales; when they do not, something is wrong. Based on the preliminary data analysis in Section 2, we have justification to expect that this mispricing will converge to zero.

Although the numbers in this example are extreme, the mechanics are general. The trade has two legs: a long and short. The long is the ETF that is relatively underpriced, and the short is the one that is overpriced. One ETF will be our long trade and the other will be the short – which is which depends on the sign of the mispricing. This is stated in Table 1.

Table 1: Trade strategies based on sign of mispricing

	Positive Mispricing	Negative Mispricing
Trade Action	*Entry: short GLD, buy IAU* *Exit: buy GLD, sell IAU*	*Entry: buy GLD, short IAU* *Exit: sell GLD, buy IAU*

A trading strategy has to specify the quantity, the direction, and the timing. Table 1 specifies the direction and we assume that we trade 1 share of GLD and 10 shares of IAU. Since the stock price of GLD is roughly ten times IAU, this achieves our goal of having equal dollar value on each stock. In a perfect world, the stocks have equal percentage changes; if we put equal dollar value in both stocks then the portfolio would have no change in value due to mispricing. It is important to have equal dollar value in each stock because this allows us to isolate the mispricing.

To specify the timing for trades in this framework, we use the magnitude of the mispricing. We want to enter when the mispricing is large and exit when it is small. We define an entry threshold *X,* so that when *MISPRICING>X* we enter a trade. In a similar fashion, we define the exit threshold *Y,* so that when *MISPRICING<Y* we exit the trade. When the entry signal is triggered, we will short GLD (bid price) and buy IAU (ask price). For these results, we only consider trades where GLD is overpriced, where *X* is a positive value. This is half of all possible trades. When the exit is triggered, we will cover GLD (ask price) and IAU (ask price). For the results provided below, we calculate the entry/exit prices using bid/ask rather than midpoint. This means the results we present will include the cost of the spread.

4. Results

The key result is the profitability of the trading rule. To initiate a trade when GLD is overpriced (*MISPRICING>X*), we short 1 share of GLD (bid price) and long 10 shares of IAU (ask price). This proportion of shares is meant to ensure that we have equal dollar value on each stock, which causes initial net cash flow to be zero. When the mispricing has 'corrected' (*MISPRICING<Y*), we cover GLD (ask price) and IAU (bid price). The resulting net cash flow is the profit or loss for this single trade. Table 2 displays the total net cash flows using this trading strategy as a dollar value and, in brackets, the number of 'round trip' trades required to achieve this profit. The columns denote the entry threshold X and the rows denote the exit threshold.

Table 2: Profit or loss for trading strategy using price data from February 13 2012

	0	0.0004	0.000475	0.0005	0.0006
-0.0005	*$0.15 (10)*	*$0.32 (10)*	*$0.38 (10)*	*$0.41 (10)*	*$0.43 (8)*
-0.0003	*-$0.22 (46)*	*$0.58 (38)*	*$0.83 (36)*	*$0.84 (32)*	*$0.78 (22)*
0	*-$12.25 (344)*	*-$0.93 (122)*	*$0.04 (76)*	*$0.26 (62)*	*$0.47 (30)*
0.0003	-	*-$3.79 (166)*	*-$1.40 (86)*	*-$0.92 (68)*	*-$0.08 (32)*
0.0005	-	-	-	*-$1.72 (78)*	*-$0.18 (34)*

Note: column denotes entry threshold (*X)* and row denotes exit threshold *(Y)*.
*Values in parenthesis () represent the total number of transactions

Table 2 shows that this trading strategy is profitable for several combinations of entry and exit thresholds. When we set the entry threshold at a mispricing of 0.0005 and the exit threshold at –0.0003, our trading strategy yields a profit of $0.84. To put into perspective, if we trade 100 shares of GLD and 1000 shares of IAU, we can expect a return of $84 by executing the strategy 32 times. This may be a considerable yield for professional day traders who are able to trade in large volumes and incur small transaction fees as compared to the average investor. When this trading strategy is applied at a higher entry threshold of 0.0006, the same 100 shares of GLD and 1000 shares of IAU can return $78. This return is less than that of using the lower

entry threshold of 0.0005, which is inconsistent with the idea that the higher the difference between the entry and exit thresholds, the higher the yield. However, the total yield alone does not provide sufficient information on the significance of the profit. It is important to also look at how many transactions were made. Standardizing the yields by calculating the ratio between total profit and total number of transactions can help determine the relative significance of each yield. This is shown in Table 3.

Table 3: Total yield to total number of transactions ratio ($/transaction)

	0	0.0004	0.000475	0.0005	0.0006
-0.0005	*0.0150*	*0.0320*	*0.0380*	*0.0410*	*0.0538*
-0.0003	*-0.0048*	*0.0153*	*0.0231*	*0.0263*	*0.0355*
0	*-0.0356*	*-0.0076*	*0.0005*	*0.0042*	*0.0157*
0.0003	-	*-0.0228*	*-0.0163*	*-0.0135*	*-0.0025*
0.0005	-	-	-	*-0.0221*	*-0.0053*

Note: column denotes entry threshold (*X)* and row denotes exit threshold *(Y)*.

By assessing the total yields relative to the number of transactions made, the results of Table 3 show a consistent trend. The yield per transaction increases as the difference between the entry and exit thresholds becomes larger, which is intuitive! Table 2 shows that when we increase the entry threshold from 0.0005 to 0.0006 and with the same exit threshold of –0.0003, the total yield decreases from $0.84 to $0.78. However, Table 3 reveals that the $0.78 profit is relatively more significant; it returns $0.0355/transaction while the $0.84 profit is equivalent to returning $0.0263/transaction. To put into perspective, if we invest in 100 shares of GLD and 1000 shares of IAU, using the entry thresholds 0.0005 and 0.0006 will return a yield per transaction of $2.63 and $3.55, respectively. Because there are fewer transactions as the difference between the entry and exit thresholds increases, an investor may incur fewer transaction fees which may further add to the significance of the profit. These results are intuitive because they show that when the entry/exit points are farther apart we find fewer trades that make larger profit per trade.

5. Extensions

To test whether this strategy has value for professional trader, we could apply it to other days to see if the pattern found in the results are consistent across time. We could also examine individual trades in more detail look to see whether each trade is profitable after transaction fees are included. We are in contact with a local proprietary trading office to explore these ideas.

Furthermore, we can apply this strategy to different ETFs that track the same index. For example, we can apply this strategy to bull/bear leveraged ETF. The basis of our trading strategy is to establish appropriate trade positions for each ETF based on the direction of mispricing, to invest equal dollar value in each stock, and to execute trades at specific levels of mispricing. Such strategy is meant to isolate and capitalize on the mispricing. In general, it is difficult to trade the exact number of shares required to invest equal dollar values in each stock, which introduces error into the implementation of our trading strategy. It may be the case that the error is larger than the mispricing we are trying to trade. Further, it may be possible to use options on ETF implement this strategy. This idea has been studied before and holds much promise (Zhang, 2010).

In our model, we defined mispricing as: percentage change in GLD – percentage change in IAU; where the opening price of each stock is used as a benchmark to calculate the percentage change. We can extend our model by including multiple benchmarks. What if we calculated percentage change relative to 10AM, 10:15AM, 10:30AM, and so on – would mispricing calculated using these percentage changes give us the same signal for our trading strategy? Because the two overlapping ETF are designed to provide the same intraday percentage change, the mispricing should be zero regardless of what the benchmark is. However, using multiple benchmarks will likely result in varied measures of mispricing at each time step. This could be extended into a graphic to help traders use the strategy in real time. Table 4 displays a simple example of how the calculation of mispricing changes with the addition of another benchmark.

Table 4: Comparison of calculated mispricing using one or two benchmarks

	Indicator 1	Indicator 2
	10:00AM Benchmark	10:15AM Benchmark
	GLD: *$167.12*, IAU:*$16.77*	GLD: *$166.99*, IAU: *$16.76*

	10:30AM	11:00AM	10:30AM	11:00AM
GLD price	$167.55	$167.80	$167.55	$167.80
IAU price	$16.81	$16.85	$16.81	$16.85
GLD % change	0.002573	0.004069	0.003324	0.004821
IAU % change	0.002386	0.004772	0.002984	0.005372
Mispricing	0.000187	-0.000703	0.000339	-0.000551

Both indicators show that there is positive mispricing at 10:30, and negative mispricing at 11:00. However, the degree of mispricing at each time step varies between the two indicators as a result of using two different benchmarks. At 10:30, indicator 1 calculates a mispricing of 0.000187, suggesting that GLD is overvalued. Our trading strategy returns a profit as long as the trade is initiated when the mispricing is greater than 0.000475 and exited when the mispricing is less than zero. So according to indicator 1, no trade will take place because the mispricing is too small. At this same time step, indicator 2 shows a mispricing of 0.000339. This level of mispricing is also below our strategy's entry threshold and will result in no trade.

For the purpose of demonstrating the effect of multiple benchmarks, suppose we set our entry threshold at 0.0003 and our exit threshold at 0. At 10:30, if we follow indicator 1, there will be no trade; if we follow indicator 2, a trade will be executed. This example suggests that calculating mispricing against multiple benchmarks can give rise to conflicting trade signals. One way to deal with this would be to wait for all benchmarks to give the same signal. However, it is unlikely that the signals will all be the same as the number of benchmarks used increases. Another way to deal with these conflicting signals would be to trade whenever the mispricing for any benchmark goes above an entry threshold (X), and cover when the same benchmark goes below an exit threshold (Y). Since it is not clear which way is better, this is an area that deserves further research.

6. Conclusion

This paper proposes a trading strategy for ETF that represent similar indices and calculates the profitability using real price data for two gold ETF, GLD and IAU. The GLD and IAU

trades very closely to each other, which means the opportunity for this type of trading is minimal. However, our results show that mispricing persists and large mispricings occur periodically; it may be possible to trade these mispricings profitably using sophisticated market access.

The results tell us that the market is highly efficient because mispricing between the two ETF is small. However, the market is efficient because sophisticated players monitor the prices of the ETF closely and trade to minimize the mispricing. The extent to which the market is inefficient (i.e. the degree of mispricing) determines the extent to which profits can be made through arbitrage. Although it is possible to calculate risk adjusted returns for this strategy to compare it with benchmark returns, that is the topic of another paper. The purpose of this paper is to demonstrate the concept and provide preliminary calculations to motivate interest in this topic.

The central idea behind our trading strategy is that mispricing between overlapping ETF converges to zero. The direction of the mispricing determines the trade position we take in each stock and the magnitude determines when a trade will be entered and closed. By defining specific combinations of entry and exit mispricing thresholds, this trading strategy may be able to yield significant profits for professional day traders, as they are able to trade in high volumes and at low transaction costs. In conclusion, our trading strategy provides evidence that it is possible to achieve profit trading on mispricings.

References

Grossman, S. J., & Stiglitz, J. E. (1980). *On the Impossibility of Informationally Efficient Markets*. The American Economic Review , 393-408.

Zhang, J. (2010). *Path-Dependence Properties of Leveraged Exchange-Traded Funds: Compounding, Volatility and Option Pricing.* New York University.