# Research on Modern Implications of

# **Pairs Trading**

Mengyun Zhang
April 2012
zhang_amy@berkeley.edu
Advisor: Professor David Aldous

Department of Statistics
University of California, Berkeley
Berkeley, CA 94720

## I. Introduction

The world of stocks and securities is filled with uncertainties and risks. However, many investors see potential to make profits in this stock market with gathered information and strategies. While some investors buy stocks based on information regarding particular companies, others utilize strategies that try to profit based on understanding of the stock market and its behaviors in general. One of such strategies in stock trades is pairs trading.

In pairs trading, investors select two correlated stocks or other comparable equities and trade only those two stocks based on their performance relative to each other. The investors would find a mean ratio for the prices of the two stock selections, buy the stock that is calculated to be underpriced, and sell the stock that is overpriced. The idea is that in the long run, the price ratio between two stocks fluctuates less than stocks themselves. So when the price of a particular stock deviates too much from the calculated mean, there is an opportunity for profits as the price will eventually go back to the mean ratio.

**In the following research, we will look into pairs trading in more details, explore the modern implications of such strategy, and conduct a simulation study to explore possible relationship and connection between profit and different variables associated with stock selections in pairs trading.**

## II. Evolution of the Strategy

Pairs trading, sometimes referred to as *statistical arbitrage*, was first developed and used by Nunzio Tartaglia in the 1980s. He led a team of mathematicians, physicists, and computer scientists in Morgan Stanley that aimed to develop automated trading systems that could take advantage of mispricing in the financial market (Vidyamurthy, 2004). These arbitrage opportunities are often times hard to spot since the true values of securities are unknown. Pairs trading uses the relative values between stocks to find "mispricing" and successfully refrains from relying on true security values. This strategy was highly profitable when it was first developed and implemented. However, as knowledge spread, profit from pairs trading was said by some to have diminished while the risk involved shot up.

Contrary to common belief, there exists more than one pairs-trading methods. The three main ones include the distance method, the cointegration method, and the stochastic spread method (Do, 2006). However, we will focus on the widely-used distance method due to its straight-forward structure.

The distance method, as implied in its name, keeps track of the sum of squared differences (a.k.a. distance) between two normalized stock prices. Stock volatility is estimated from historical data and a trading trigger is determined, usually to be two standard deviations from the mean. When the thresholds are met, the stocks are assumed to be mispriced in terms of their relative value with each other. Investors buy the underpriced one and short the overpriced one. Investors will profit in the long run if the relative price levels of the stocks revert back to the mean.

The distance method is non-parametric and does not assume the stock prices to follow certain models. Therefore it is not subject to the assumptions of established parametrized models or potential errors resulting from stock behavior's deviation from such models. Also, the utilization of relative price in general eliminates the volatility from co-funding variables inherent to all stock prices including but not limited to inflation rate. However, it has a critical underlying assumption of a static price level distance between two stocks. This assumption renders the method vulnerable to price-level divergence, which could lead to sizeable losses to the investors, especially if a stop-loss method is not implemented. Another downsides of this non-parametric method is its lack of "forecasting ability regarding the convergence time or expected holding period", as summarized by Binh Do (2006).

One of the major concerns as pairs trading becomes more popular is its profitability. As mentioned earlier, many believe that pairs trading only allowed the first users to take advantage of the arbitrage opportunities or market inefficiencies that it could catch. They argue that as knowledge spread, these market inefficiencies would be eliminated in no time. Others had supporting evidence of the opposite and argue that "the public dissemination of the results has apparently not affected the general risk and return characteristics of the strategy, despite curiosity from the professional sector" (Gatev et al. 2006). In the following study, I will try to evaluate how the pairs trading method performs compares in recent years.

## III. Example:

I decided to simulate a series of trades to illustrate the pairs-trading distance method. Note that I manually selected my two stocks: Nordstrom (JWN) and GAP (GPS) as both of them are clothing retailers and competitors with each other. Ideally, pairs of stocks are better selected using automated programs which could spot the ones that tend to move closely together in terms of price ratio.

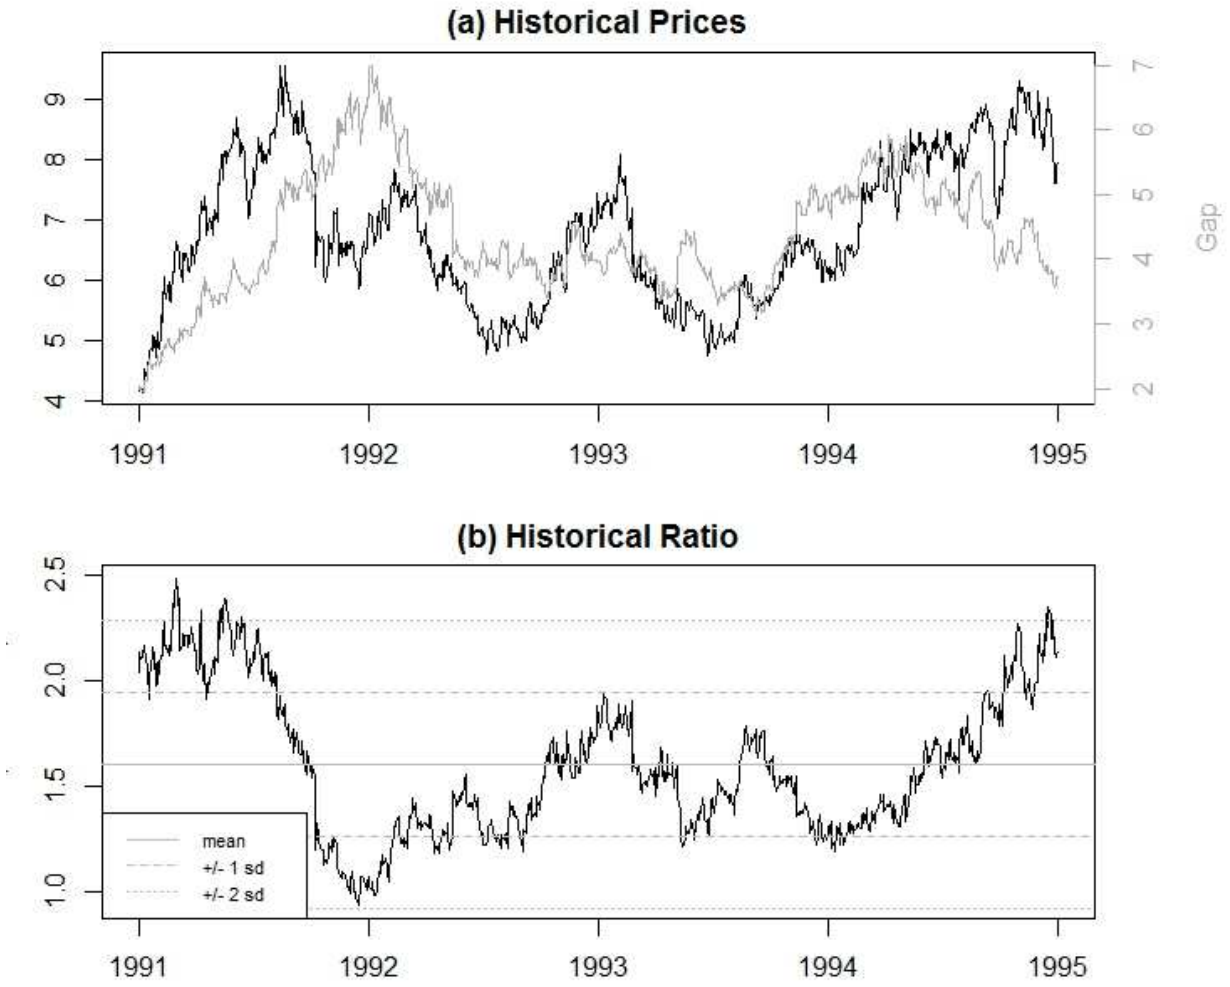


Figure 1. Historical prices and ratio for Nordstrom(JWN) and Gap(GPS).

(a) Historical Prices: the curve and y-axis in black corresponds to JWN;
the curve and y-axis in grey corresponds to GPS.

(b) Historical ratio: JWN/GPS

Panel(a) in Figure 1 above shows their historical prices from 1991 to end of 1994. The black curve models Nordstrom's historical prices while the grey curve represents that of Gap's. Notice how their prices are positively correlated. In Panel(b), the historical ratio of the two stocks are shown over the same time period, illustrating the deviation of the price ratio about the mean. When this ratio deviates from the mean by more than a few standard deviations, there is an opportunity to make money. The tradition is to set the threshold at two standard deviations about the mean. Notice the ratio curve crosses the mean a few times, allowing for closings of deals when the price ratio hits the mean.

Figure 2: Ratio of Nordstrom to Gap for 1995-present.
The grey dotted lines show 1 and 2 standard deviations away from the mean ratio.

Taking the above data from 1991-1994 as data of reference, we implement the pairs trading strategy to the data from 1995 to present (May 9th) and calculate our profit from the trades done during this period. Figure 2 above shows the ratio of Nordstrom to Gap with trading points circled. Setting the trading threshold at two standard deviations above or below the mean ratio, we open a position when the ratio first reaches a threshold. We buy the relatively "undervalued" stock and sell the relatively "overvalued" stock. We close each deal when the ratio goes back to the mean ratio, selling shares at hand. One rule is that we do not do any more trades before we close one.

Our first trade occurs when the ratio reaches -2 standard deviations, meaning Nordstrom is underperforming relative to Gap. Therefore we buy (1/price of Nordstrom) units of Nordstrom and sell (1/price of Gap) units of Gap. As shown on the figure above, from 1995 to present, we have a total of seven trade transactions in theory. However, since our last position is an opening position, we close our position on May 9th 2012 when calculating profit even though the price ratio has not hit the mean yet. If trading is completely free with zero transaction cost, the profit adds up to ~$4.00. If there is a transaction fee ($p$) of 0.001 every time we trade, then the profit is down to ~$3.98. It means that for every $1-trade we engaged in, **we would have profited ~$3.98 by the end of this time period** with a transaction cost of 0.001. Note that the $1-trade actually does not involve any net investment on our part. Instead, it means that we long $1 worth of one stock and short $1 worth of the other.

Definition: $1-trade = longing $1 worth of one stock and short $1 worth of the other

Although $3.98 seems like a decent return, the time span of the series of trade is very long-term (~17.5 years!). An investor would have made much more money if he simply longed the stocks. Inflation and economic development over the years are directly reflected on the stock market. The adjusted closing price of S&P 500 was $34.37 on January 3, 1995 and $135.74 on May 9, 2012. However, one characteristic of the pairs trading strategy is its relative immunity to macroeconomic changes. Thus, many investors may prefer pairs trading as opposed to simply longing or shorting stocks to hedge macroeconomic risks.

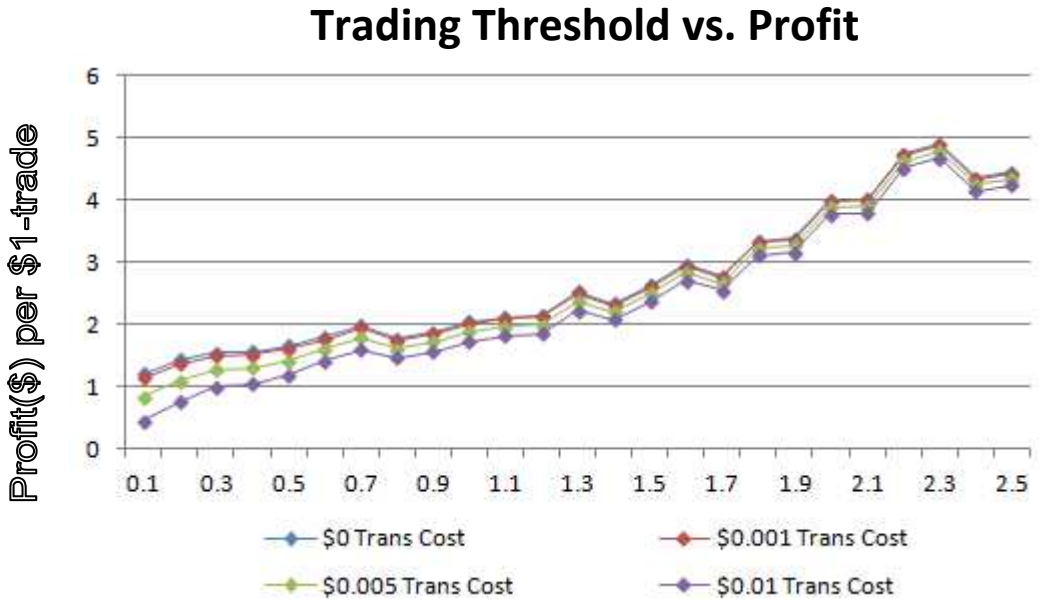


Figure 3. Profit (y-axis) as a function of trading thresholds (units: standard deviation)
Stock Selection: Nordstrom and Gap

However, at times the 2-standard-deviation threshold may seem arbitrary. The above graph sheds light on the effect of having different threshold selection on profit. Note that this graph can only be done in retrospect and cannot serve as a means to find an optimal threshold in real-time trades.

## IV. Comparison of Performance: 1980s vs. Present

In this section, I'm set out to evaluate the performance of pairs trading in modern days as compared to in the 1980s by doing another pairs trading simulation. I changed my selection of stocks in this part of the study since Gap and Nordstrom only went onto the market in the late 1980s. Instead, I will be using Coca-Cola (Symbol: KO) and Pepsi (Symbol: PEP) to conduct this comparison.

1980s: Assuming it is the beginning of 1986 and an investor decides to pairs trade with Coca-Cola and Pepsi. He uses data from the beginning of 1980 to the end of 1985 as reference data to determine a price ratio between the two. Similar to our example, Panel(a) in Figure 3 below shows the stocks' historical prices from 1980 to end of 1985. Notice how their prices are positively correlated just as one expects. In Panel(b), the historical ratio of the two stocks are shown over the same time period, illustrating the deviation of the price ratio about the mean.

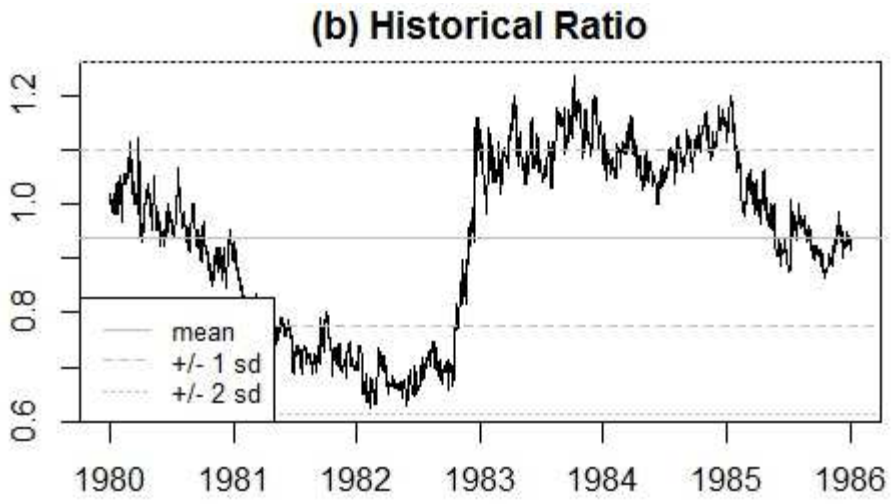


Figure 4. Historical prices and ratio for Coca-Cola(KO) and Pepsi(PEP).

(a) Historical Prices: the curve and y-axis in black corresponds to KO;
the curve and y-axis in grey corresponds to PEP.
(b) Historical ratio: KO/PEP

This time, suppose the investor sets the trading threshold at 1 standard deviation away from the price ratio mean. This is mainly because the price ratio of the two never reached the 2-standard-deviation threshold during the period. In real practice, one cannot predict as to whether a particular threshold can be reached beforehand. An investor who is determined to trade at 2-standard-deviation about the price-ratio mean would simply not trade at all with this pair of stocks. However, for study purposes, we will look at the trading activities and profit outcomes given various thresholds. The following Figure 4 shows the price ratio changes from 1986 to the end of 1990 with the gray circles indicating the investor's trading points.

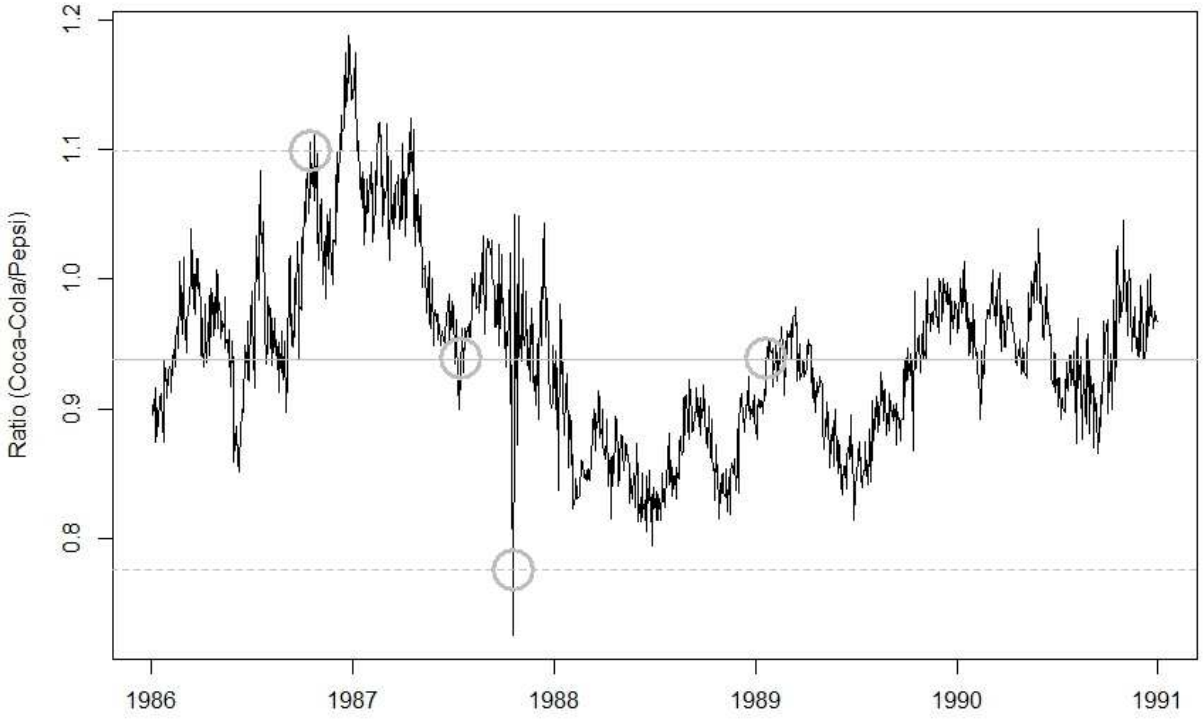


Figure 5: Ratio of Coca-Cola to Pepsi for 1986-1990.

The grey dotted lines show 1 and 2 standard deviations away from the mean ratio.

The investor's end-period profit totals to ~$0.62 per $1-trade as defined on page 7 assuming zero transaction fee. The profit dilutes to ~$0.61 assuming 0.001 transaction cost.

Again, let us examine the of threshold selection on the overall profit. The following graph (Figure 5.) illustrates the profit outcomes with different trading thresholds. The three curves reflect the different scenarios given by different transaction costs.

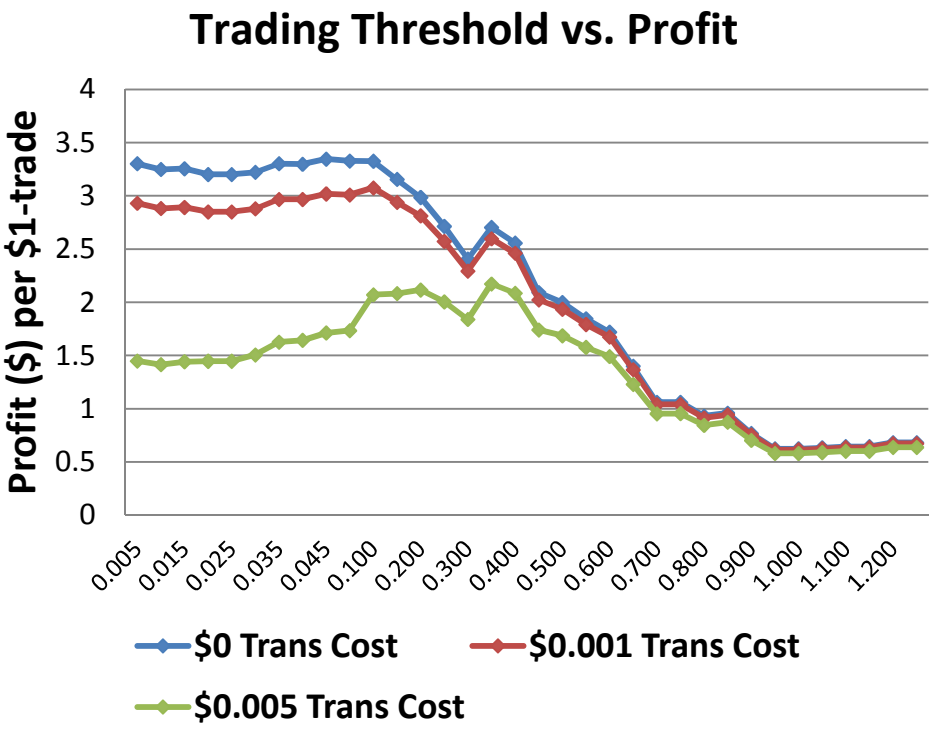


Figure 6. Profit (y-axis) as a function of trading thresholds (units: standard deviation)
Stock selection: Coca-Cola and Pepsi

When transaction cost is zero, it seems that the investor could have made the most profit by choosing a very small trading threshold in this case. However, with a transaction cost of $0.005, or 0.5%, setting minimal trading threshold is not economical. The three scenarios converge as trading threshold goes up. Finding out the optimal trading threshold beforehand using this graph is not realistic since the outcomes can only be analyzed in retrospect. However, an investor may be curious to learn if there is an optimal trading threshold to follow in general cases.

It is interesting to contrast Figure 6 with Figure 3. Contrary to what one may have expected, the Trading threshold vs. Profit graphs for the two simulations show completely opposite trends. While profit is maximized if the investor opened positions whenever the

price ratio hit 2.3 standard deviations above or below the mean in the Nordtrom/Gap example, profit is maximized setting the trading threshold at ~0.1 standard deviations in the Coca-Cola/Pepsi example here. Scrutinizing the price-ratio history of the two examples, we can see that the price ratio is more volatile in the Nordstrom/Gap case. On the other hand, the price ratio of Coca-Cola/Pepsi fluctuates more closely around the mean, allowing for more profiting opportunities if smaller trading thresholds are chosen.

Present-times:

Next, let us see what would have happened if we applied the pairs trading methodology used above to Coca-Cola and Pepsi again in the 2000s.

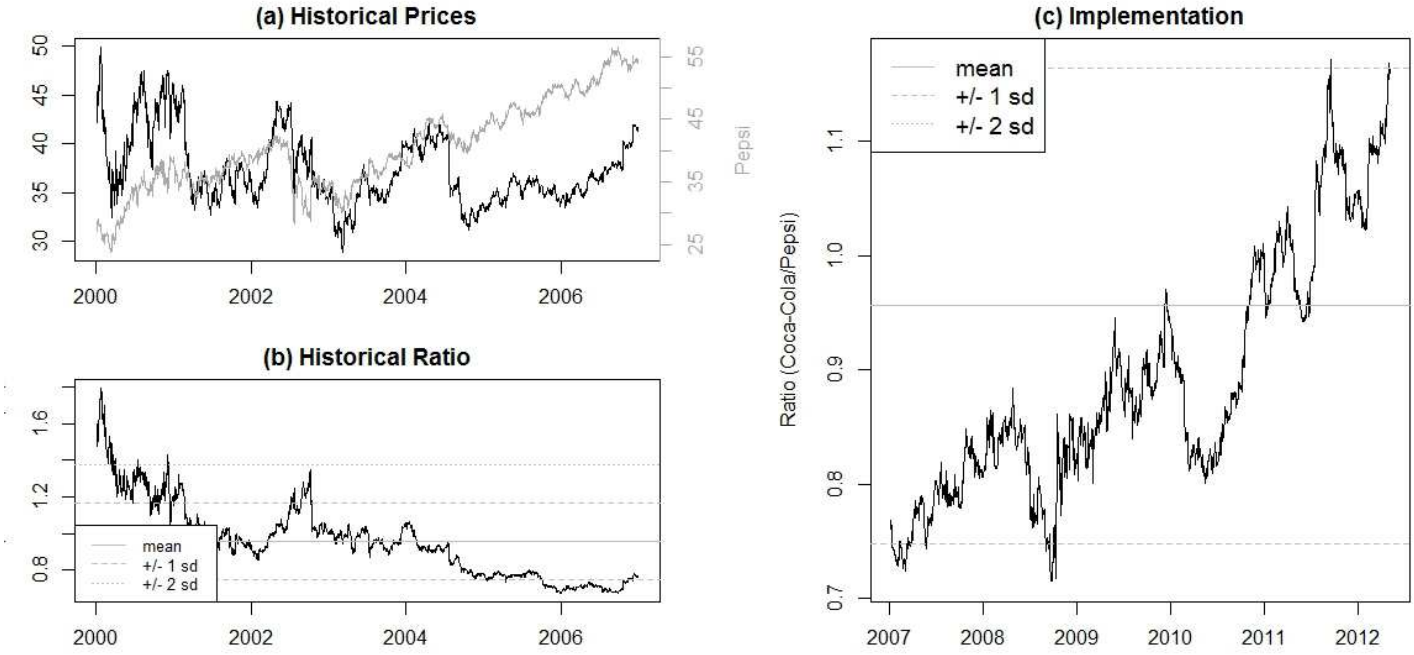


Figure 7. (a) Historical Prices: the curve and y-axis in black corresponds to KO; the curve and y-axis in grey corresponds to PEP.
(b) Historical ratio: KO/PEP
(c) Price ratio 2007- May 9 2012 for Pairs Trading implementation

If an investor is looking for pairs trading opportunities at the beginning of 2007 and examines the historical prices and ratio between Coca-Cola and Pepsi, he can probably spot a divergence trend and not select this pair of stocks. From 2000 to the end of 2006, the price of Pepsi increases relatively steadily against Coca-Cola. The historical ratio of Coca-Cola to Pepsi shows a decreasing trend that emits warnings to potential pairs-traders not to engage. A trended price ratio (as opposed to a constant one or random-noise one) in general signals the failure of a basic assumption that pairs trading relies on: a static price-ratio difference.

However, if we look at (c) in Figure 7, we can see that the price ratio of the two actually reverts back to the mean calculated using historical prices from 2000-2006. In retrospect, we see that the relative price of Coca-Cola hits a low-point during 2006 but began to increase afterward. About six years passed from the time the price-ratio dived below the mean to the time it reverted back at the end of 2009. Comparing to its behavior in the 1980s, the price ratio of Coca-Cola to Pepsi became much more volatile in the 2000s. If I were an investor, this pair is definitely not my top-pick for pairs trading.

It is important to note that depending on which years of data are used as references, the investors can draw very different conclusions about the stocks. Price-ratio mean and standard deviation may vary drastically as reference data change. Therefore, it is critical that an investor chooses appropriate number of years to reference, which can involve substantial research into the history of the two companies and judgment calls. For example, an investor should catch differences in targeted companies' directions, which may lead to divergence in their price-ratio and affect profitability of a pairs trading strategy.

## V. Simulation Study:

In this section of the study, I will examine the relationship between profit and stock selections by doing a simulation study. The simulation involves generating stock data with different properties and calculating the profits using simulated data. I used R to generate stock picks following a model called vector autoregression. The code to simulate stock behavior in R was written by Professor Kaufman in UC Berkeley. In short, this model allows us to generate stock picks with certain properties that are similar to those in the real time. This model incorporates the following variables about the virtual stock picks, or sequences:

- psi: a value between zero and one that controls how correlated the sequences are with each other.
- beta1: a slope that is assigned to each stock price sequence; it controls the change in stock prices over time.
- *rho: a value between zero and one that controls how correlated the sequences are in time, which is set to 0.99 throughout the simulation.*
- *beta0: sequence's base value, which is set to $100 throughout the simulation.*
- *sigma squared: controls the distribution spread of the generated sequences, which is set to 1 throughout the simulation.*
- *p: trading cost, which is set to 0.001 throughout the simulation*

Among the listed variables above, I only varied the first two (psi and beta1) values in generating the stock pick sequences. The rest that are italicized is held fixed throughout the simulation study. Figure 8 below demonstrates the effect of varying psi and beta1 values.

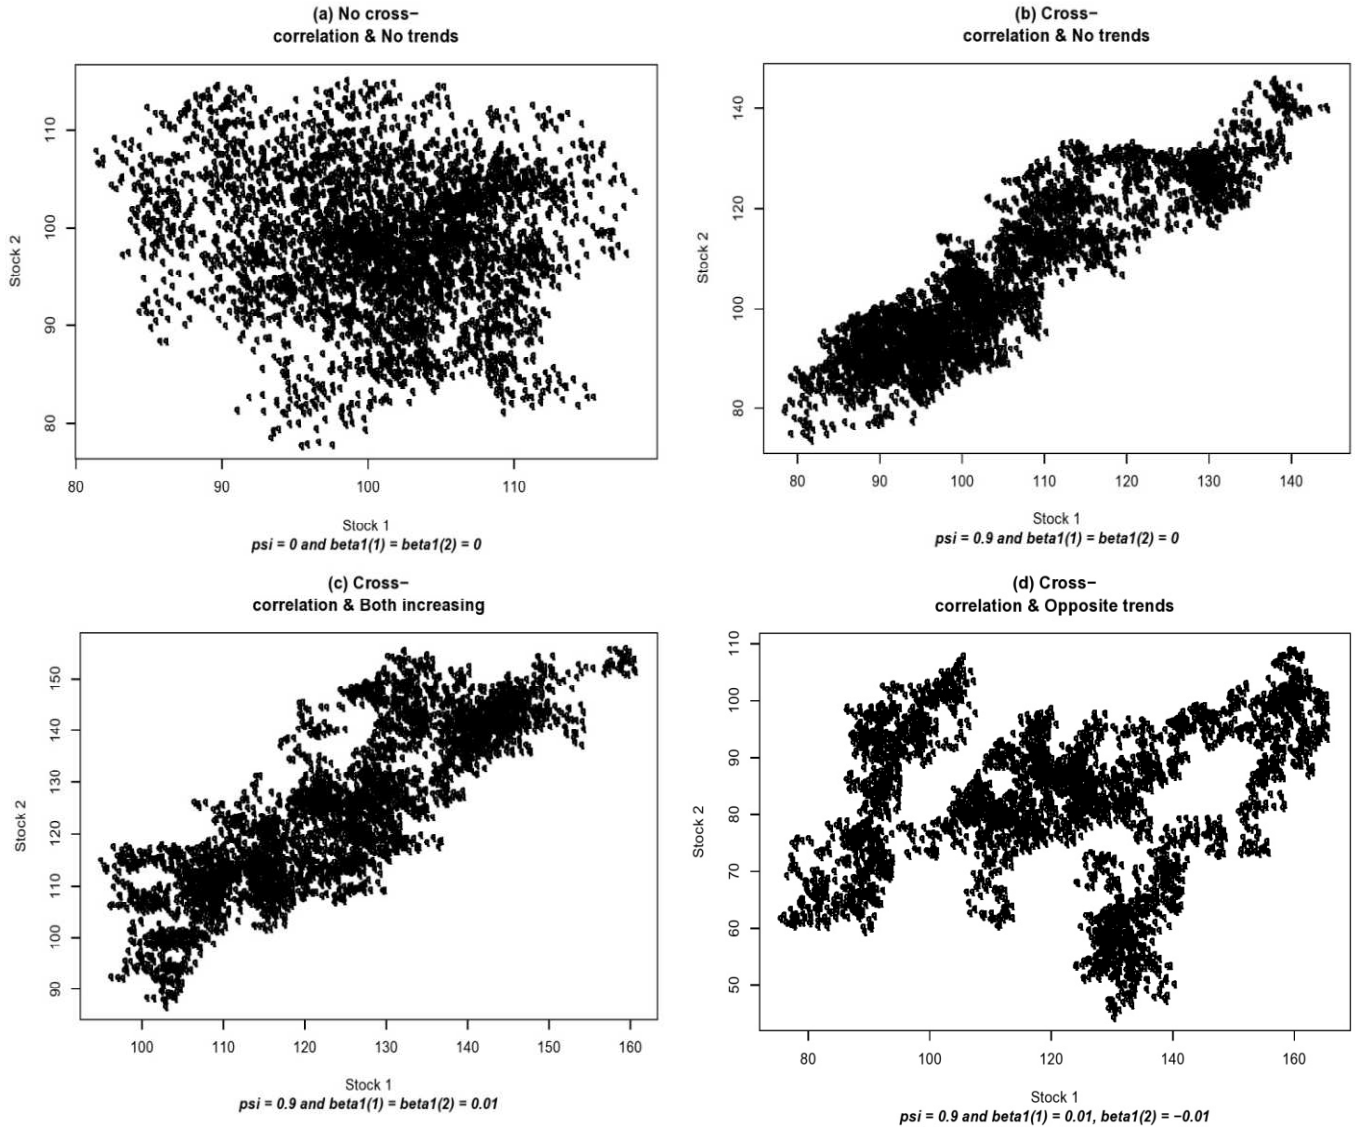


Figure 8. Effect of varying psi and beta1 values in 4000 days of data each.

As shown in the Figure 8(a) above, when there is no cross-correlation between the two stock picks and no trend for price change over time, the values of the two stocks when plotted together form random scatters. Figure 8(b) shows the price plot with the two stocks closely correlated (psi = 0.9), and no clear trend for price change over time. Figure 8(c) is when the two stock picks are closely correlated and have increasing price over time. Notice that the scatter plot in 8(b) is not as closely correlated as those in 8(c) as it

incorporates an increasing price trend over time. Figure 8(d) is when the stock picks are positively correlated but lead opposite price-change trends over time. Notice the clusters formed by scatters of data.

Next, I implemented the pairs trading strategy to different pairs of stock picks generated from the stock simulation model described above. In this part of pair trading, in addition to *psi* and *beta1* values, I varied another variable *k*, which is the number of standard deviation that sets the threshold for trading. Notice that the bigger the *k* value is, the fewer the expected number of transactions. However, a bigger *k* also means that the trade amount each time is greater.

To summarize, I varied the following variables during the pairs trading simulation of virtual sequence picks:

- psi: correlation between the two sequences, *variations of psi used: 0, 0.5, 0.9*
- beta1: slopes of the two sequences in relation to time,

  *Variations of beta1 used: (0,0), (0.01, 0.01), (0.01, -0.01), (-0.01, -0.01)*
- k: threshold value, *variations of k used: 1, 2, 3*

The total number of variations derived from the variations of the above variables acumulates to: 3 x 4 x 3 = 36 combinations. For each combination, 1000 simulations are conducted to calculate an average profit. Each simulation is composed of two stock picks (sequences) of 4000 days of data.

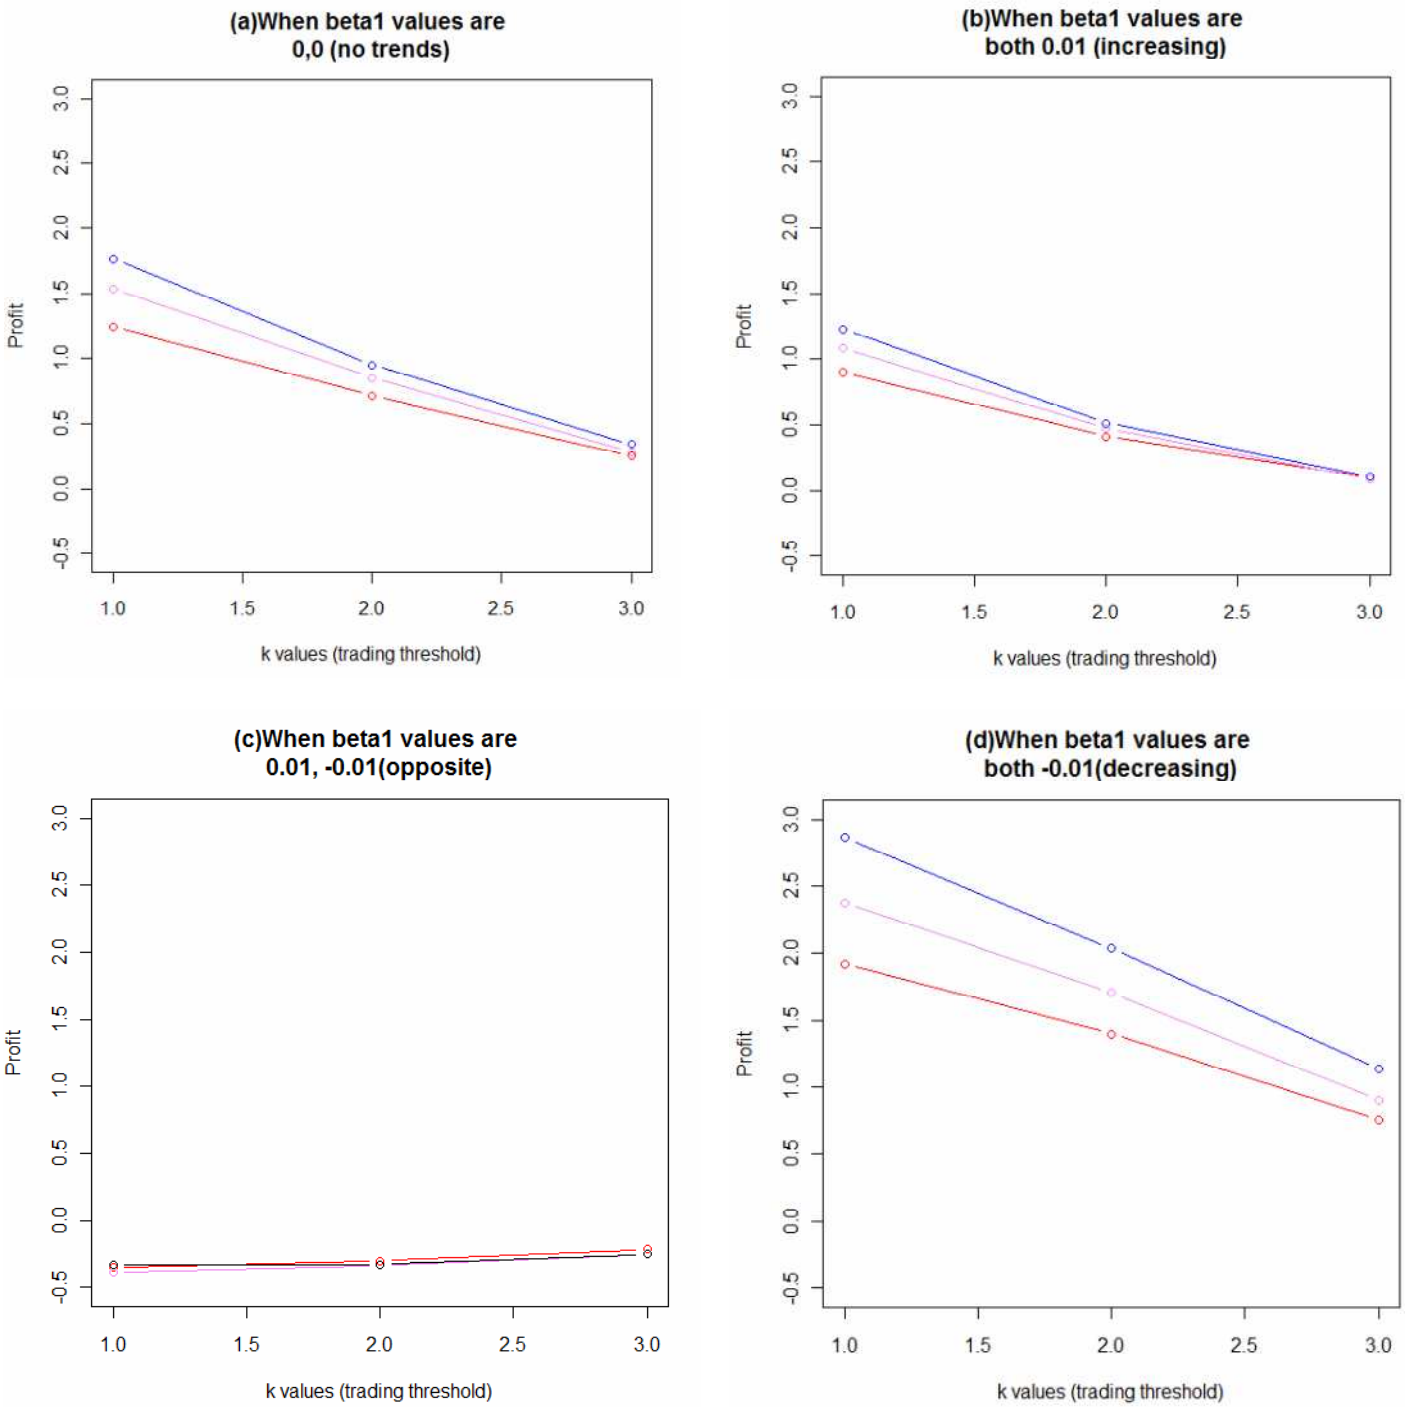


Figure 9. Average Profit vs. Combinations of different psi, beta1, and k values
Red curve: psi = 0, Violet curve: psi = 0.5, Blue curve: psi = 0.9

Figure 9 illustrates the average profits from varying the three factors (psi, beta1, and k). On each graph, the x-axis corresponds to the three k values: 1, 2, and 3, while the y-axis shows the average profits calculated from the simulation. There are three curves within each graph, with the red curve having psi = 0, violet curve having psi = 0.5, and blue curve having psi = 0.9.

Examining Figure 4, we can see the effects of the three variables on the average profit of pairs trading clearly:

Psi (=0. 0.5, 0.9): Pairs of stocks that are more closely-correlated make more money, since the blue curve (psi = 0.9) always yield the highest profit in each graph besides graph (c), which yields a negative profit because of its beta1 values. The red curve with no correlation between the two stocks yields the least profit in each case.

Beta1: Comparing the four graphs, (d) wins with average profits higher than in any other graph. It seems that pairs trading can yield the most profit when the two stocks both show decreasing (price) trend over time. Pairs trading for two stock selections that have opposite trend over time yield negative profit over time (lose money). To summarize, the more negative the slopes are, the more profitable pairs trading is when applied, given the two stock selections have the trend.

k (=1, 2, 3): Besides in panel(c), pairs trading is more profitable with a lower threshold value. In real time terms, it's better to set a lower threshold value and make many small trades than to do fewer large trades (higher threshold value). This concept does not apply when the two stock selections have opposite trends over time, since small trades seem only to lose more money. Note how our Nordstrom/Gap example counters this phenomenon.

## VI. Limitations to the Simulation Study

Since the study is not firm-focused, it overlooked significant changes within the firms that could potentially change their stock performances forever. These factors include but are not limited to: changes in infrastructure, technological advancement, mergers and acquisitions. Some other limitations of this simulation study lies with the fact that we made a lot of assumptions which aren't necessarily true in real life. For example, we held the variables constant during the period of time we take the samples. However, in real life, we cannot assume that two firms have the same determining variables over long periods of time. We also assumed that we can always buy and short stock shares and that we do not have a limit to our budget. In real life, we always have limits to our budget and cannot afford to face unlimited uncertainties.

In fact, our pairs trading examples using historical data of stocks already proved some parts of this simulation to be too ideal in real practice. For instance, setting a smaller trading threshold for the Nordstrom/Gap trade modeled above can only reduce our profit due to the volatility of their price-ratio.

## VII. Future study possibilities

Due to my limited knowledge on the different mathematical models that can also be used to analyze pairs trading, I overlooked many possibilities in my study above. As mentioned on page 3, cointegration provides a valid model for pairs traders as well.

The idea of cointegration and the term were first proposed by Engle and Granger, who won the Nobel prize in 2003. A key underlying notion is error correction. Cointegrated systems have a long-term equilibrium; "that is, the long-run mean of the linear combination

of the two time series", where the times series indicate the stock movement of the two stocks in the pairs-trading case (Vidyamurthy, 2004). It implies that if the linear combination of the two times series deviates from the long-run mean, then the series adjust themselves to restore the equilibrium eventually. This idea echoes with the underlying assumption we use in pairs trading, where we assume the price-ratio of the two stocks is static over time. This error correction process can be represented in the following equations:

$$y_t - y_{t-1} = \alpha_y\left(y_{t-1} - \gamma x_{t-1}\right) + \varepsilon_{y_t}$$
$$x_t - x_{t-1} = \alpha_x\left(y_{t-1} - \gamma x_{t-1}\right) + \varepsilon_{x_t}$$

Where *x* and *y* are time series. The first part on the right side of the equation is the error correction term, α is the error correction rate which decides the speed with which x and y correct themselves to go to equilibrium, and γ is the coefficient of cointegration. The error term, or any deviation from the long-run equilibrium, represents any white noise in the process. These white noises are the only source of deviation in this model.

Nevertheless, the cointegration method has its limitations in real practice. As Vidyamurthy put it, "it is unlikely that the stocks' factor exposures will be perfectly aligned" (Vidyamurthy, 2004). Imperfect alignment deviates from the ideal condition for cointegration. He also mentions the likely nonstationary common factor spread in real stocks, which again create challenges for the cointegrated models to work ideally.

Geometric Brownian motion is also one thing that can be studied together with pairs trading. In fact, it is the most commonly-used model to model stock behaviors in modern days. It assumes that the logarithmic stock returns over time is normal, and is a stoichastic process. However, I decided to focus more on the actual historical data to illustrate the

process of pairs trading. Nevertheless, I am aware of the limitations of using historical data too: my stock selections are critical to the outcomes. I was careful not to draw too much conclusions about pairs trading in general with my specific studies.

**References:**

Do, B. (2006). A new approach to modeling and estimation for pairs trading. Retrieved from http://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=3&ved=0CHMQFjAC&url=http%3A%2F%2Fwww.finanzaonline.com%2Fforum%2Fattachments%2Feconometria-e-modelli-di-trading-operativo%2F1048428d1238757908-spread-e-pair-trading-pairstradingbinhdo.pdf&ei=-NmqT7XdBoSy2QXOnIWnAg&usg=AFQjCNF6mGmHQq0Ubqa1AzHq1VWHR5Zv_g&sig2=GVF6tNGNRCXaNbOQQGRElA

Gatev, E. (2006). Pairs trading: Performance of a relative value arbitrage rule. Retrieved from http://rfs.oxfordjournals.org/content/early/2005/12/31/rfs.hhj020.full.pdf

Vidyamurthy, G. (2004). *Pairs trading quantitative methods and analysis*. New Jersey: John Wiley.