BI Norwegian Business School - campus Oslo

GRA 19502

Master Thesis

Component of continuous assessment: Forprosjekt, Thesis MSc

Preliminary Master Thesis

Navn	Alexander Vestgård, Jørgen Hem Toftevaag
Start:	01.01.2018 09.00
Finish:	15.01.2018 12.00

Preliminary Master Thesis

BI Norwegian Business School

Economic Links and Predictable Returns: Boldly going where someone has gone before

(*Working title*)

GRA 19502 Master Thesis

Master of Science in Business, Major in Finance

Supervisor: Patrick Konermann

Start date:
01.01.2018

Finish date:
15.01.2018

BI Oslo

Table of contents

Abstract

List of people to thank

Introduction

"The illusion that we understand the past fosters overconfidence in our ability to predict the future."

Daniel Kahneman, 2011

This thesis is based on the work of Cohen and Frazzini (2008). In this article, the authors find evidence of return predictability across economically linked firms. Firms are linked together through customer-supplier relationships, where announcements and news in one firm will, or should, affect other linked firms. Cohen and Frazzini (2008) argue that investors display limited attention and therefore returns from one firm do not immediately reflect news about the related firm. Specifically, bad (good) news from the customer has a negative (positive) effect on the supplier and the supplier's return does not immediately incorporate this. Cohen and Frazzini (2008) mainly used a zero-cost customer momentum trading-strategy, which in short consists of buying the quintile of supplier firms stocks whose customers had the most positive returns and short-selling the quintile of supplier firms whose customers had the most negative returns. In an efficient market, one should earn no abnormal returns using this strategy as one uses information known to investors. They showed that this strategy generates abnormal returns of 1.55% per month or an annualized return of 18.6% per year. Further, they draw the conclusion that this is due to investor limited attention as investors are not able to incorporate all information and the effect of the related firm is therefore not immediately priced in the market.

The research question is whether the abnormal returns by using the Cohen and Frazzini (2008) method are still present today. Further, McLean and Pontiff (2016) finds that abnormal returns drastically decline post-publication. However, they did not include the Cohen and Frazzini customer-momentum-factor in their paper, therefore, it would be interesting to look at how it may have changed after the publication of "Economic Links and Predictable Returns" (2008). McLean and Pontiff (2016) argue that "If return predictability reflects mispricing and publication leads sophisticated investors to learn about and trade against the mispricing, then we expect the returns associated with a predictor should disappear or at least decay after the paper is published."

We create a new dataset using the same procedure as Cohen and Frazzini, this gives us the possibility to investigate post-publication abnormal returns. Further, we are introducing a new econometric method to the old and the new dataset. By using Vector Autoregression (VAR) and variance decomposition, we can then provide an additional argument for (against) investor limited attention, should the results prove significant (insignificant). To our knowledge neither using Vector Autoregressions in this setting or reviewing Cohen and Frazzini's results on new data has been done before. Additionally, Fama (1998) argues that long-term return anomalies tend to disappear with reasonable changes in methodology, and we, therefore, argue that it is a fitting extension to include VARs.

By investigating Cohen and Frazzini's method on a new dataset we might be able to argue that investor limited attention towards customer-supplier links have changed. To solidify the underlying assumption that we have reproduced Cohen and Frazzini's methodology, we must use our methodology on their dataset to provide evidence that we are able to incorporate it correctly. By checking the post-publication returns, we can conclude if there has been a statistically significant change after the paper was published. In addition, by using the additional VAR estimates, we can argue in a different manner whether investors immediately incorporate positive or negative shocks or if there is a lagged effect in the manner shown by Cohen and Frazzini's customer momentum.

The thesis is important as we review whether a proven trading strategy is still valid after it was published in a paper, which is essentially a check for market efficiency. This is relevant, as there has been an increased focus on abnormal return anomalies disappearing after cross-predictability of return articles have been published. Further, by extending Cohen and Frazzini (2008) we are putting our paper in the context of Fama (1998), who finds that anomalies tend to disappear when the methodology is changed, and McLean and Pontiff (2016), who find that published academic articles destroy stock return predictability.

Insert summary of results

The remainder of this thesis is organized as follows. Section I gives a brief overview of relevant literature. Section II describes the data, both where and how it can be found, and we present summary statistics for our sample. Further, section III provides the main theory and section IV presents the methodology of our thesis. in section V and VI we provide results and robustness checks and in section VII we conclude.

I Literature Review

Cohen and Frazzini (2008) show that stock returns of customers predict stock returns of their suppliers when investors do not immediately act to customer-supplier information. They were the first to do this, and there have been a wide variety of researchers who have used this methodology as a starting point of their papers. Pandit et al. (2011) find that suppliers experience information externalities at the time of their customer's earnings announcements, which could be used to predict returns. Further, Paatela et al. (2017) propose a trading approach in which one creates a portfolio of supplier stocks that exclusively consists of companies whose main customer’s quarterly sales evolve favourably, this portfolio is robust to a range of control variables and is uncorrelated with the market, attributes which Cohen and Frazzini found as well. Further, we have the limited attention hypothesis. Kahneman (1973) and Peng and Xiong (2006) concentrate on investors’ learning behaviour given limited attention, and they find that attention is a scarce cognitive resource and attention to one task necessarily requires a substitution of cognitive resources from other tasks. Peng (2005) shows that information capacity constraints can cause a delay in asset price responses to news, which is what we are utilizing in our paper. Pashler and Johnston (2008) have a summary of the literature on attention limitations. McLean and Pontiff (2016) show that cross-predictability of returns disappear post-publication. They argue that returns are often a result of mispricing and that sophisticated investors would swarm towards a possibility of arbitrage.

After Cohen and Frazzini (2008), researchers within the field of return prediction within economic links seems to heavily focus on industries, and not firms. Given the obvious economic links between firms within the same industry or along the industry supply chain, the current literature focus is probably not surprising. Aobdia et al. (2013), Menzly & Ozbas (2010) and Rapach et al. (2015) looks at inter-industry networks and cross-predictability of returns. Herskovic (2015) examines the asset pricing implications of input-output networks. Further, Müller (2017) argues that firms with the same stock-characteristics omit a higher cross-predictability than other firms, and he, therefore, moves away from the industry focus. Most of these papers use the same or similar dataset as we are using, and the main idea behind their research is the same and stems from Cohen and Frazzini, Menzly and Ozbas (2006) and Hong et al. (2007). One of the first who modelled the concept of investor inattention was Merton (1987). Merton's model is based on investors who obtain information, and trade, on a small number of stocks. Stocks with fewer traders sell at a discount stemming from the inability to share risks.

Not much is done on firm-specific customer-supplier return predictability after Cohen and Frazzini (2008). However, Cen et. al. (2017) shows that their speed measure helps investors generate "sharper" customer momentum strategies regarding slow information diffusion, nonetheless, this is not the focus of our thesis. As most of the literature is focused on industries, and not firms, there could be gaps in the firm-specific customer-supplier momentum literature. However, the focus on industries might be a consequence of a lack of firm-related topics. Additionally, Cohen and Frazzini presented a very thorough investigation in their paper, which might have relieved the topic for further research. We believe, nonetheless, that our new dataset and use of VAR fits a gap in the literature, as there is no one who has done this before, to the best of our knowledge. Further, by checking whether the momentum strategy by Cohen and Frazzini (2008) is still there, we elaborate on their study with more data and by introducing a different methodology of checking for effects, and at the same time, we are checking if returns are still there post-publication for a study not considered in McLean and Pontiff (2016).

Most of the industry focused papers are utilizing input-output networks, which is technically more advanced than what we are going to do. We are replicating the methodology used in Cohen and Frazzini (2008). They mainly focus on a zero-cost long-short strategy, which should in a correctly priced market earn zero abnormal returns, if this is not the case it would indicate mispricing. In the case of Cohen and Frazzini, abnormal return is due to investor limited attention. The zero-cost long-short strategy is an intuitively appealing methodology which is easy to grasp and explain, and in our view, correct. Our extension of implementing vector-autoregressions presents a different method of finding (not finding) evidence for the investor inattention. Bad news (bad returns) concerning the customer can be regarded as a negative shock to the supplier. By using VAR, we can then review if the shock is on average incorporated in the supplier's stock price when the information is accessible, or if the shock is incorporated at a later stage which would be in line with investor limited attention.

In addition to the customer-supplier literature, our paper also contributes to a broader literature on behavioural finance, asset pricing and market efficiency. We contribute to this literature by reviewing whether the abnormal returns are still there or not by using a new dataset, which is a check for if the market has removed the strategy's potential returns after it is published, or market efficiency (McLean and Pontiff, 2016; Fama, 1998).

II Data

We obtain the first range of data from Andrea Frazzini's homepage[1]. This data is from the period 1980-2004 and contains customer sales, total sales, date of relation, together with customer name and CRSP permno number and finally supplier CRSP permno. The remaining data needed for us to reproduce Cohen and Frazzini's results are stock returns, which can be found in the CRSP database and matched with the dataset through date and CRSP permno number. In accordance with their method, we will impose a 6-month gap between fiscal year end-dates and accounting information to ensure investors are aware of the customer-supplier relation.

We create the new dataset with the intention of mimicking Cohen and Frazzini's data collection process. As the goal is to review whether the abnormal returns are still present today, it is important that we do not change the data collection process relative to the dataset of which we are comparing results. If the datasets would somehow systematically differ from each other, it would weaken possible results as it could be a symptom of contrasting datasets. In accordance with the above-

[1] http://people.stern.nyu.edu/afrazzin/data_library.htm

mentioned arguments, we have therefore followed all restrictions Cohen and Frazzini used. Firstly, we extract customer-supplier relationships using the COMPUSTAT customer segment file for the time period between 2011-2017 (2011.06-2017.09[2]). This file contains information about customers who represents more than 10 % of a supplier's sales which suppliers report in their financial statements[3]. In practice, a firm can also voluntarily disclose customers that account for less than 10% of total revenues.

Furthermore, we match the company name with their CRSP permno number. For suppliers, the CUSIP number is also matched towards permno number as this is included in the COMPUSTAT customer segment file. Some customer names are written in slightly differing ways[4], we manually control all non-matches and conservatively adjust if segment, industry and name suggest a match is warranted. This procedure is similar to Cohen and Frazzini's hand-matching of pre-1998 relations, our matching is simply somewhat more tedious. An additional variable we extract from the COMPUSTAT database is book values, which we match according to name and date, to ensure the matched book value is matched at the correct time. Cohen and Frazzini focus exclusively on common stocks[5]. We assume that share code assignment has not changed since 2008, as we have found no indication that would suggest otherwise. Therefore, only common stocks with share codes 10 or 11 are included in our dataset as well, since this information is contained in the CRSP database we utilize the CRSP permno number together with the date of the relation to match our observation with the share code. Some companies change share code over time, therefore, we ensure that the correct share code is matched with the correct company at a given time by including the date of the customer-supplier relation. The next step is to obtain the stock prices for both supplier and customer firms in our sample from the CRSP monthly returns database. We match these stock prices with the firms from Compustat. Currently, the data accessible to us is limited to 2011-2016, however, after the annual update, around February, the 2017 data will become accessible and we will then assign monthly

---

[2] Will be extended to 2017.12

[3] Statement of Financial Accounting Standards (SFAS) No. 14 and No. 131 states that suppliers must report all customers with more than 10% of total sales.

[4] E.g. inc. compared to inc

[5] Share code 10 and 11

returns to the customers and suppliers and impose the 6-month gap mentioned earlier.

Insert summary statistics, available after the CRSP February update

Compare the summary statistics with Cohen and Frazzini's summary statistics

Firms are required to disclose the identity of any customer representing more than 10% of total reported sales, therefore we are likely to identify larger firms as customers since larger firms are more likely to be a major stakeholder. So far, our dataset contains 3,865 observations for which we can identify both customer and supplier as a traded COMPUSTAT/CRSP listed firm. The observations have non-missing book values at the fiscal year end and are common stocks. Currently, we have an average of 618 observations per year, compared to Cohen and Frazzini's average of 1,082. The difference in the number of observations are initially worryingly large, however, we will be able to compare the datasets more thoroughly after the database update in February and then be able to determine if we must revise our data collection process, or not. Further, one could argue that the difference in the datasets is a symptom of change in the market between the two collection processes. The interpretation is then that less customer-supplier relations are simply reported in period 2011-2017. In other words, there are fewer companies who have customers which represent more than 10% of total sales. There are arguably some good arguments for this case. Increased globalization will intuitively lead to a wider customer base. Additionally, even if customers across countries represent more than 10% they would have to be listed in the US according to the requirements presented earlier and this would likely not be the case for the majority, which would lead to such relations not being included.

On a clarifying note, we would like to have included data from the years 2005-2010, however, we, unfortunately, do not have access to the necessary customer-segment COMPUSTAT database for the years 2005-2010. As Cohen and Frazzini's dataset ends in 2004, there is an out-of-sample possibility from 2004 until Cohen and Frazzini published their paper in 2008, such as the one in McLean and Pontiff (2016). However, we are still able to review whether mispricing still exists in the market, as there are some missing years we do not have the opportunity to exactly determine when the mispricing possibly disappeared or declined, even though it would clearly be both important and interesting to check. Nonetheless, our dataset should be statistically large enough to give unbiased estimates.

III Theory

The research question is to check whether it is still possible to earn abnormal returns from a customer-supplier momentum strategy of economically linked firms, which Cohen and Frazzini (2008) showed yielded abnormal annualized returns of 18.6%. This strategy should not yield any, or have significantly decreased[6], abnormal returns in the years after they published their paper, as sophisticated investors should have swarmed towards this strategy. By checking the returns, we are putting this strategy in the setting of McLean and Pontiff, who finds that return predictability disappears or decreases after a paper is published, and Fama who postulates that reasonable new econometric techniques could remove abnormal return anomalies.

Consider an example of two economically linked firms, one supplier and one customer. The basic theory and intuition behind our thesis are based upon the fact that if these two firms are linked, actions or announcements in one firm, should affect the other, and in our case, the customers should affect their supplier. Further, Chan-Lau (2017) argue that networks emerge naturally from direct bilateral exposures between financial institutions and other market participants, which is the case between suppliers and customers. Statement of Financial Accounting Standards (SFAS) No. 14 and No. 131 states that suppliers must report all customers who represent more than 10% of total sales. It is these reports we collect our data from. Because of the 10% cut-off, our data has information about customers who can be regarded as major stakeholders. Therefore, if the customers share price or earnings per share (EPS) forecast drops or increases, so should the suppliers. The reason is that the supplier is somewhat dependent on the customer. Additionally, the decline in share price for the supplier should by intuition be based on the percentage of sales they have to that specific customer[7] and should happen effectively right after the customers share price drop if the markets are efficient.

[6] Costly arbitrage strategies or non-friction less trading could hinder the total disappearance of the strategy.
[7] Pandit et al (2011) finds evidence of this.

However, the practice is or has been, that this does not happen straight away, as the supplier's EPS lags behind the customers[8], which further brings us towards the Limited attention hypothesis. The limited attention hypothesis is based on underreaction in stock prices regarding firm-specific information that induces changes in the valuation of related firms, generating return predictability across assets. In particular, stock prices underreact to negative (positive) news involving related firms, and in turn generate negative (positive) subsequent price drift (Cohen and Frazzini, 2008). Hence, in the presence of investors subjected to attention constraints, stock prices do not rightly incorporate news about related firms, and thereby generate stock price predictability across assets.

The efficient market hypothesis states that it should not be possible to beat the market over time, and further that anomalies should disappear over time. Fama (1998) argues that "Consistent with the market efficiency hypothesis that the anomalies are chance results, apparent overreaction to information is about as common as underreaction, and post-event continuation of pre-event abnormal returns is about as frequent as post-event reversal". A strategy consisting of a zero-cost strategy with abnormal annualized return of 18.6% percent for the period between 1980 and 2004 is not in line with the description above.

Considering the slow information processing of investors (attention constraints), stock prices do not correctly incorporate announcements about linked firms, which translates into stock return predictability, as discussed earlier. However, we want to check if this relationship is there, by utilizing methods not used before. Our hypotheses are therefore, (i) that using vector autoregressions will yield a different result than the classical linear regression, in the sense that the VAR makes a simple network, which will, or should, consider the customer-supplier relationships in a better way. This puts our hypothesis (i) in the setting of Fama (1998), who argues that most anomalies should disappear with the introduction of new econometric techniques. Further, (ii) that the return predictability from Cohen and Frazzini (2008) has disappeared or has significantly (negatively) changed, considering the findings of McLean and Pontiff (2016).

[8] Figure 1: Coastcast-Callaway example.

IV Methodology

Based on the setting of this thesis, Cohen and Frazzini's methodology is justified and utilized in a handful of other top-journal papers, such as the ones mentioned in the literature review. To test if the returns from following the Cohen and Frazzini customer-momentum strategy still yields the same abnormal returns as before publication, we run a linear regression on the post-publication returns, as done in McLean & Pontiff (2016). This enables us to see if there is a significant change in returns post-publication. We follow the strategy Cohen and Frazzini primarily focused on, which is a zero-cost long-short strategy. Each month customers are assigned to five different groups according to last month's return in an ascending order. The groups are value weighted (equally) and we go long (short) in the respective group of suppliers for whose customers are in the group of most positive (negative) returns. As earlier discussed, a strategy such as this should earn zero abnormal returns in an efficient market as one is using information available to all investors (lagged returns and customer-supplier link). We believe that extending the classical linear regression (or AR-process) into a VAR, will better model the relationships between suppliers and customers. After running the VARs, we argue that doing Variance Decompositions, which models how much each variable contributes to the other variables (Brooks, 2015), enables us to make a simple network, which we can utilize to answer our hypotheses. Variance decomposition, or Forecast Error Variance Decomposition (Brooks, 2015), will show us to which degree customer returns affects supplier returns.

As there is likely a relationship between a firm's lagged returns or earnings and their stock price, an AR-process would be appropriate. Further, since we want to set up the regressions and results as a simple network, a vector autoregression which would take both the supplier and customers lagged variables into account, could be a better predictor for the relationship between the linked firms. Therefore, we argue that a VAR would be appropriate. Additionally, VAR, which controls for lagged returns and makes a simplified network, is the method we argue that is the most fitting, as we are able to run variance decompositions, which will give us the possibility to make a simple Variance Decomposition network similar to Diebold and Yilmaz (2014). Further, extending into a VARMA, a VAR with moving averages, could explain the data better than the VAR regarding the firm relationships and could be a valid extension. Mathematically and computationally difficult however, an input-output network, like the ones in Herskovic (2017) or Aobdia et al. (2013), would likely model the relationships in a more parsimonious way. However, this is not within the scope of our thesis.

An essential prerequisite for utilizing the customer-supplier relation is data on suppliers and their customers. Suppliers report customers who represent more than 10% of annual income in the financial reporting, as noted earlier. This data includes sales specified by date to specific customers. Further, to be able to reproduce the zero-cost strategy we need returns for both supplier and customer. Additionally, data on book values and share codes are needed to be in line with Cohen and Frazzini's use of common shares with a non-missing book and market value at fiscal year-end. In regard to control variables, a wide range can and should be used, including; analyst coverage (from I/B/E/S[9]), trading volume, difference in size between customer and supplier, as well as institutional ownership, all of which are discussed and tested in Cohen and Frazzini. These variables that are most often used in other empirical research regarding customers and suppliers. It is essential to include control variables to ensure that our results are not influenced by some other factors.

To find the optimal lag-length for our vector autoregression, we will use multivariate information criteria. The information criteria that we choose, most likely MAIC or MBIC[10], will choose the optimal lag-length for our VAR, or in other words, the model that is likely the most parsimonious. It is important to note that it is interesting to review how many lags are statistically significant as well, as to review when the shock is fully incorporated. Since the amount of data is substantial, we will program a loop which will run VAR's for a customer-supplier relation, return the data and run the VAR on the next relation and so on. The idea is to find when lagged returns are no longer statistically significant as this would indicate the time of full information incorporation in the stock price. Interesting tests to run on the data from this process include Granger causality test, variance decomposition and impulse response function.

[9] Explanation of IBES can be found at https://corporatefinanceinstitute.com/resources/data/bloomberg/ibes/

[10] Multivariate Akaike or (Schwarz)-Bayesian Information Criteria

Firstly, variance decomposition is a test for which variables affect the other variables in the VAR, how forecasting is affected and how shocks develop over time. We argue that the variance decomposition enables us to make a simple network for the customer-supplier relationships. Also, we argue that it will better model the relationships and it will help us to interpret the estimates of the vector autoregression. Additionally, given the presence of the 10% cutoff for the suppliers to report on their financial reporting, our sample has more information about customers who are major stakeholders and not the reverse. Thus, our main tests are in the direction of suppliers' stock price response to customers' shocks, as Cohen and Frazzini argue. However, by running a Granger Causality test (vargranger), we can test whether this is the case, as this test shows predictive causality, and which way it goes. Therefore, we will be able to confirm or reject this by running a vargranger. We can also run an impulse response function, as a mean to interpret the VAR. The impulse response function measures the persistence of shocks in a variable from other variables. At the very end, we would want to control our portfolios for Fama-French 3 factors, as well as the Fama-Macbeth factor, including momentum, as these are necessary steps to conclude if the portfolio yields abnormal returns.

As earlier discussed, our methodology mainly consists of two parts; a zero-cost strategy and vector autoregressions. We will now briefly touch upon possible results and the following interpretation of these. Concerning the results of the VAR-model, if we find estimates which result in a different answer than Cohen and Frazzini (2008) or our own zero-cost strategy, then this means that our hypothesis (i) is accepted. However, if the estimated parameters give the same answer as in Cohen and Frazzini and our own zero-cost strategy, then we reject the null. An interpretation of a rejection is straightforward, both methods point in the same direction and would strengthen the argument of investor limited attention. On the other hand, an acceptance of the null needs a more reflected interpretation. A possible explanation is that the VAR-model is able to pick up something in the data that the classical linear regressions are unable to, an argument one could say is related to Fama (1993) arguments of different econometric methods resulting in different answers. For our post-publication regression, the significance of the regressions implies if we can accept hypothesis (ii) or not. If the beta of the regression is statistically significant and the beta is negative, then this implies that we can keep our null hypothesis. On the other hand, if there is no significant change, then we will have to reject the null. If we reject the null hypothesis, then one of the interpretations might be that the Cohen and Frazzini strategy is not affected in the manner McLean and Pontiff argue, however, if we accept the null hypothesis, then we will be able to conclude that the abnormal returns are no longer present today, which would be in line with McLean and Pontiff discussion of how published papers destroy stock return predictability.

The Next Steps

The first step is to finish the dataset when the data for 2017 is released on CRSP (in February). Then we will be able to present the summary statistics for the dataset we have acquired for 2011-2017. Further, the next step is to correct our preliminary for the feedback we get. We will then continue to program the customer-supplier momentum portfolios, running this on both datasets, to ensure we have programmed it correctly and write about the results. Further, we will move on to the vector autoregressions and gather results. This is the main part of our analysis. At the end, we will do robustness-tests for a variety of factors, to check if the VAR is as robust as the classical linear regression. We plan to utilize "Variance Decomposition Networks: Potential Pitfalls and a Simple Solution" (Chan Lau, 2017) as an explanation and approach for how to make a VD network.

Bibliography

Aobdia, D., Caskey, J., & Ozel, N. (2013). Inter-Industry Network Structure, Information Transfers, and the Cross-Predictability of Earnings and Stock Returns. *SSRN Electronic Journal*. http://dx.doi.org/10.2139/ssrn.2196196

Brooks, C. (2015). *Introductory econometrics for finance*. Cambridge: Cambridge University Press.

Carhart, M. (1997). On Persistence in Mutual Fund Performance. *The Journal Of Finance*, *52*(1), 57. http://dx.doi.org/10.2307/2329556

Cen, L., Hertzel, M., & Schiller, C. (2017). Speed Matters: Limited Attention and Supply-Chain Information Diffusion. *SSRN Electronic Journal*. http://dx.doi.org/10.2139/ssrn.2925460

Chan-Lau, J. (2017). Variance Decomposition Networks: Potential Pitfalls and a Simple Solution. *IMF Working Papers*, *17*(107), 1. http://dx.doi.org/10.5089/9781475598407.001

Cohen, L., & Frazzini, A. (2008). Economic Links and Predictable Returns. *The Journal Of Finance*, *63*(4), 1977-2011. http://dx.doi.org/10.1111/j.1540-6261.2008.01379.x

Diebold, F., & Yilmaz, K. (2011). On the Network Topology of Variance Decompositions: Measuring the Connectedness of Financial Firms. *SSRN Electronic Journal*. http://dx.doi.org/10.2139/ssrn.1937613

Fama, E. (1991). Efficient Capital Markets: II. *The Journal Of Finance*, *46*(5), 1575. http://dx.doi.org/10.2307/2328565

Fama, E. (1997). Market Efficiency, Long-Term Returns, and Behavioral Finance. *SSRN Electronic Journal*. http://dx.doi.org/10.2139/ssrn.15108

Fama, E., & French, K. (1993). Common risk factors in the returns on stocks and bonds. *Journal Of Financial Economics*, *33*(1), 3-56. http://dx.doi.org/10.1016/0304-405x(93)90023-5

Herskovic, B. (2015). Networks in Production: Asset Pricing Implications. *SSRN Electronic Journal*. http://dx.doi.org/10.2139/ssrn.2615074

Hong, H., Torous, W., & Valkanov, R. (2007). Do industries lead stock markets?. *Journal Of Financial Economics*, *83*(2), 367-396. http://dx.doi.org/10.1016/j.jfineco.2005.09.010

Hou, K., & Moskowitz, T. (2005). Market Frictions, Price Delay, and the Cross-Section of Expected Returns. *Review Of Financial Studies*, *18*(3), 981-1020. http://dx.doi.org/10.1093/rfs/hhi023

Kahneman, D. (1973). *Attention and effort*. Englewood Cliffs, N.J.: Prentice-Hall.

Kahneman, D. (2011). *Thinking, fast and slow*. New York: Farrar, Straus and Giroux.

McLean, R., & Pontiff, J. (2016). Does Academic Research Destroy Stock Return Predictability?. *The Journal Of Finance*, *71*(1), 5-32. http://dx.doi.org/10.1111/jofi.12365

Menzly, L., & Ozbas, O. (2006). Cross-Industry Momentum. *SSRN Electronic Journal*. http://dx.doi.org/10.2139/ssrn.517822

Menzly, L., & Ozbas, O. (2010). Market Segmentation and Cross-predictability of Returns. *The Journal Of Finance*, *65*(4), 1555-1580. http://dx.doi.org/10.1111/j.1540-6261.2010.01578.x

Merton, R. (1987). A Simple Model of Capital Market Equilibrium with Incomplete Information. *The Journal Of Finance*, *42*(3), 483. http://dx.doi.org/10.2307/2328367

Müller, S. (2017). Economic Links and Cross-Predictability of Stock Returns: Evidence from Characteristic-Based “Styles”. *Review Of Finance*. http://dx.doi.org/10.1093/rof/rfx033

Paatela, A., Noschis, E., & Hameri, A. (2017). Abnormal Stock Returns Using Supply Chain Momentum and Operational Financials. *The Journal Of Portfolio Management*, *43*(2), 50-60. http://dx.doi.org/10.3905/jpm.2017.43.2.050

Pandit, S., Wasley, C., & Zach, T. (2010). Information Externalities Along the Supply Chain: The Economic Determinants of Suppliers' Stock Price Reaction to Their Customers' Earnings Announcements. *SSRN Electronic Journal*. http://dx.doi.org/10.2139/ssrn.1600394

Pashler, H., & Johnston, J. (1998). Attentional Limitations in Dual-task Performance. Retrieved from http://www.bryanburnham.net/wp-content/uploads/2014/01/Pashler-Johnston-1998-Attentional-Limitations-in-Dual-task-Performance.pdf

Peng, L. (2005). Learning with Information Capacity Constraints. *Journal Of Financial And Quantitative Analysis*, *40*(02), 307. http://dx.doi.org/10.1017/s0022109000002325

Peng, L., & Xiong, W. (2006). Investor attention, overconfidence and category learning. *Journal Of Financial Economics*, *80*(3), 563-602. http://dx.doi.org/10.1016/j.jfineco.2005.05.003

Rapach, D., Strauss, J., Tu, J., & Zhou, G. (2015). Industry Interdependencies and Cross-Industry Return Predictability. *SSRN Electronic Journal*. http://dx.doi.org/10.2139/ssrn.2566541

Schopohl, L. (2014). *Stata Guide to Accompany Introductory Econometrics for Finance*. Cambridge University Press. Retrieved from https://docmh.com/stata-guide-to-accompany-introductory-econometrics-for-finance-microsoft-excel

Appendices

Appendix A - Figures and Tables

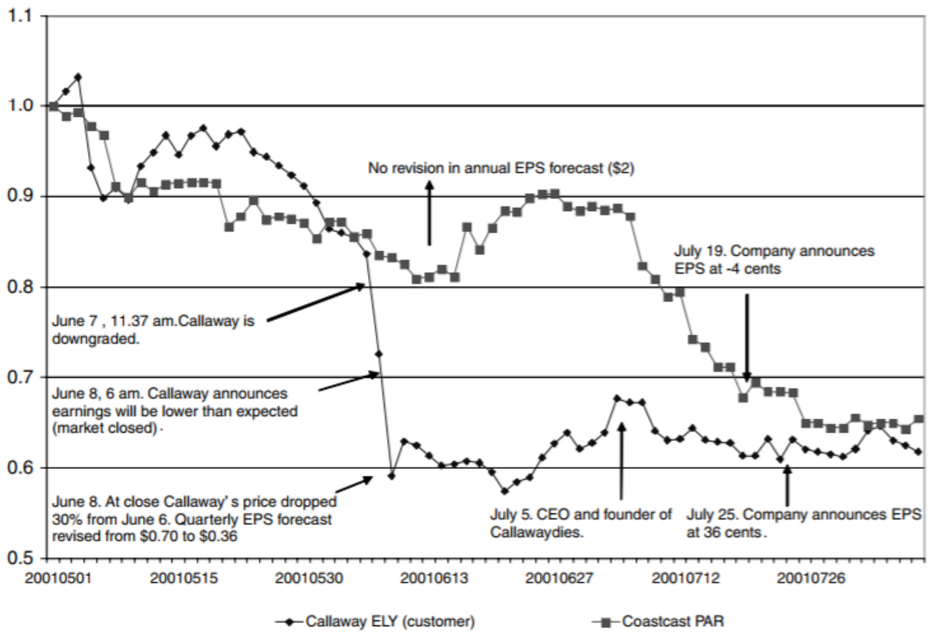

Figure 1. Coastcast Corporation and Callaway Golf Corporation. This figure plots the stock prices of Coastcast Corporation (ticker = PAR) and Callaway Golf Corporation (ticker = ELY) between May and August 2001. Prices are normalized (05/01/2001 = 1).

Appendix B - Code