Using Partial Least Squares to Forecast Market Returns

Adam Stivers
DeGroote School of Business
McMaster University
Hamilton, ON L8S 4M4
Canada
stiverad@mcmaster.ca
phone: (905) 525-9140 x27019

Abstract

This paper forecasts market returns using various Fama-French portfolio returns and a partial least squares (PLS) regression approach. The results are an improvement over the Kelly and Pruitt (2013) results, which uses a three-pass regression filter (3PRF) and portfolio book-to-market ratios, based on both in-sample and out-of-sample R-squared values. This paper also contributes to the literature by examining the portfolio weights produced by partial least squares. These weights are found to be more stable, less extreme, and more reliable than OLS betas. Small size portfolios are found to have higher weights than large size portfolios when forecasting, and several industry portfolios are found to have consistently high weights. Finally, this paper finds that PLS and the 3PRF both capture a single risk factor: market risk.

This version: January 15, 2015

Introduction

Kelly and Pruitt (2013) use a three-pass regression filter introduced in Kelly and Pruitt (2012) to forecast market return and cash flow growth. They note that their three-pass regression filter (3PRF) is a special case of partial least squares. This paper expands on the work of Kelly and Pruitt (2013) by directly using partial least squares (PLS) to forecast, rather than using their 3PRF. While only rarely considered in Finance and Economics applications, PLS is *a priori* a strong candidate as a forecast method because it estimates jointly the factors and factor sensitivities that in-sample provide the highest covariance with the forecast variable.

By using PLS and a different set of predictors, this paper improves on the out-of-sample (OOS) R^2 of Kelly and Pruitt for market returns. PLS also allows for proper examination of portfolio weights when forecasting. This paper examines which portfolios matter more than others in explaining future market returns. Thus, this paper makes two contributions to the literature: improved forecast power for market returns and examination of which portfolios best explain future market returns.

PLS provides both a more accurate and more stable forecast compared to OLS. It is shown that OLS coefficients are more volatile over time, and the OOS performance of OLS is inferior to both the historical mean and PLS. OLS also produces extreme weights due to over-fitting. This is especially troublesome when there is a large set of predictor variables. If an investor were to use OLS weights to construct a portfolio or hedge against market risk, it would require constant rebalancing and short-selling. PLS produces positive, diffuse, and stable weights that would be more useful in practice (in addition to superior forecasts). This paper examines the weights on the size and value sorted portfolios as well as various industry portfolios that constitute the forecast factor.

The results show that small size-sorted portfolios explain more of future market returns than large size-sorted portfolios. Higher book-to-market portfolios tend to have higher weights compared to lower book-to-market portfolios, but the effect is not as clear as it is with size (especially when looking at the highest and lowest book-to-market sorted portfolios only). This analysis also allows for examination of which industries matter for future returns. Although this paper looks at several different Fama-French industry classifications, a clear pattern develops as to which industries have a large weight and which have a low weight. For example, the energy industry (mining, coal, oil and gas) and the telecommunications industry tend to have low weight when forecasting market return. At the other end of the spectrum, the entertainment industry and finance industry tend to have a larger weight than other industries.

While the portfolio weight information is valuable, Novy-Marx (2014) raises an issue that this paper attempts to address. Novy-Marx finds that random variables such as political party of US president, weather, and positioning of planets have significant predictive power for various anomalies. Thus, it could be suggested from this that out-of-sample predictive power may be spurious. The question could be asked as to whether the methodology used in this paper provides any useful information in forecasting, regardless of performance. This is addressed by performing principal-components analysis to create factors which are used as predictors instead of raw portfolio returns. It is expected that any relevant information from these portfolios would be captured in the first few factors (depending on how many risk factors there are). The results support this hypothesis: when forecasting with PLS and using the principal-component factors as predictors, only the first factor has significant weight. This is in stark contrast to the results on raw portfolio returns where the weights are very diffuse. With only the first factor having significant weight, it also suggests only one risk factor is at play. An AR(1) model of excess market returns is used to forecast, with predictive power similar to the results in this paper, and superior to the results of Kelly and Pruitt (2013). This suggests that the results from both PLS and the 3PRF are due to successfully capturing the market risk factor.

The sample period in this study is January 1930 – December 2010, with forecasting starting January 1980, as in Kelly and Pruitt (2013). However, Kelly and Pruitt use portfolio book-to-market ratios to forecast. This paper uses portfolio returns, which improves on the OOS-R^2 and allows weights and factors to be easily implemented in practice. Monthly excess log returns are used, as in Kelly and Pruitt. The CRSP value-weighted index is used for market returns, and the yield on 3-month Treasury bills is used for the risk-free rate. This paper looks at the 6, 25, and 100 size-value portfolios that Kelly and Pruitt use, as well as various Fama-French industry portfolios. The results are compared to OLS and the Kelly and Pruitt results.

The paper proceeds as follows. Section 2 reviews the Kelly and Pruitt (2012) and (2013) papers and other relevant forecasting literature. Section 3 introduces the partial least squares regression approach. Section 4 reviews the methodology and data. Section 5 provides the forecast results, a comparison to OLS, and examines the portfolio weights. Section 6 concludes the paper.

2. Literature Review

Kelly and Pruitt (2012) create a three-pass regression filter, which can be used to forecast. They note that it is essentially a special case of partial least squares. Kelly and Pruitt (2013) is the first (to this author's knowledge) to apply PLS to finance. First, this section will review the 3PRF of Kelly and Pruitt (2012). Then, the methodology and results of using the 3PRF to forecast market returns in Kelly and Pruitt (2013) will be reviewed.

The first pass of the 3PRF is a time series regression where each predictor variable is the dependent variable and the proxies (or the variables to be forecasted themselves) are the regressors. The first pass coefficients are stored and become the regressors in the second pass cross-sectional regression where the predictors are again the dependent variable. The second pass coefficients are then used as predictors in a third pass predictive regression. The three passes can be summarized as follows:

1. $x_{i,t} = \phi_{0,i} + z_t^{'}\phi_i + \varepsilon_{i,t}$

2. $x_{i,t} = \phi_{0,t} + \hat{\phi}_i^{'} F_t + \varepsilon_{i,t}$

3. $y_{t+1} = \beta_o + \hat{F}_t^{'}\beta + \eta_{t+1}$

where x are the predictors, z is a proxy to be selected by the forecaster, and y is the variable to be forecasted. Kelly and Pruitt (2012) find that this approach is especially useful when using a large set of predictors. They forecast macroeconomic variables and conduct Monte Carlo simulations to show that their method is superior to OLS and several variants of principal-components analysis (PCA).

Kelly and Pruitt (2013) apply the 3PRF to forecast market returns and cash flow growth on both a monthly and annual basis using portfolio book-to-market ratios. They do so using the 6, 25, and 100 size and book-to-market sorted Fama-French portfolios. They also do not use a proxy for market returns in the first pass; instead, they use market returns as the regressor in the first pass (y in place of z in 1 above). They obtain significantly positive in-sample and out-of-sample R-squared values (OOS calculated based on the historical mean). The results hold for various robustness tests, and for both US and international data.

The literature on predictability of market returns is vast and diverse. For a summary of forecasting the equity premium, see Welch and Goyal (2008). Polk et al. (2006) use the cross-sectional price of risk to forecast the equity premium. Ferreira and Santa-Clara (2011) use a sum-of-parts method that combines the three components of market returns (dividend-price ratio, earnings growth, and price-earnings ratio growth). They find that this performs better than other variables as well as the three components individually. Guo (2006) uses the consumption-to-wealth ratio (cay) of Lettau and Ludvigson (2001) to forecast market returns. Guo finds that this measure has substantial out-of-sample predictive power. This is in contrast to results of Bossaerts and Hillion (1999) and Goyal and Welch (2003), which find little out-of-sample predictive power of traditional forecast variables such as dividend yield, term premium, default premium, etc. In addition to Kelly and Pruitt (2013), Li et al. (2013) is another recent study that forecasts market returns. Li et al. create their own implied cost of capital measure, which they find has significant out-of-sample forecast power.

3. Partial least squares methodology

Kelly and Pruitt (2013) note that their 3PRF reduces to PLS when using the forecast variable in the first pass instead of an alternative proxy, de-meaning and variance-standardizing the predictors prior to the first pass, and removing the constants from the first two passes. This paper finds that using PLS instead of 3PRF and portfolio returns instead of book-to-market ratios significantly improves forecast performance based on OOS-R squared values.[1] PLS is most commonly used in the field of chemometrics. In addition to the NIPALS algorithm that is frequently used in this field, de Jong (1993) creates a SIMPLS algorithm. De Jong notes that this

[1] One other advantage of PLS compared to the 3PRF is that PLS (SIMPLS specifically) allows for more than one factor to be created (which is not done in this paper).

algorithm avoids the construction of deflated data matrices while also calculating the factors directly from the variables. This allows for faster computation compared to NIPALS. Nonetheless, the SIMPLS algorithm is shown to produce basically identical results compared to NIPALS when there is only one dependent variable (as is the case in this paper). For these reasons, this paper follows the SIMPLS algorithm.[2]

The de Jong SIMPLS algorithm is essentially as follows.[3] The matrix Y of the variables to be forecasted is demeaned, and the dominant eigenvector of $X^T YY^T X$ is selected. The eigenvector is post-multiplied by $X^T YY^T X$ to give the X factor weights, which weight the X matrix to give a linear combination of the X variables in order to produce the X factor, or X-score. The X-score is normalized to produce X loadings (X * X-score) and Y loadings (centered Y *X-score). The Y matrix times its loading gives the Y-score. The loadings are orthogonalized to all previous loadings and stored. Current loadings are then removed from the $X^T YY^T X$ matrix and the process is repeated for any additional factors. The PLS regression coefficients can also be conveniently found by multiplying the X weights by the Y loadings.

Partial least squares can be thought of as a variant of principal-component analysis. Rather than maximizing the covariance matrix $X^T X$, PLS maximizes $Y^T XX^T Y$ or $X^T YY^T X$, which maximizes the covariance between Y and X (Y being the forecast variable(s) and X being the predictor variables in this case). Theoretically, this should give PLS greater forecasting power compared to PCA, as the latter creates a linear combination of the variables that explains only variations in X itself. Kelly and Pruitt (2012) find that this is indeed the case. Another key difference between PLS and PCA is that PLS removes each successive factor from the data and

[2] For a summary of PLS, see Hoskuldsson (1988).
[3] For a full explanation and the detailed algorithm, please see de Jong (1993) appendix.

then starts the procedure over to produce the next factor, which will be orthogonal to all previous factors.

An additional advantage of PLS is that by design it is a type of shrinkage estimator. Shrinkage estimators give a stable forecast over time, with diffuse weights. OLS can give extreme weights due to over-fitting. Thus, when dealing with a large set of predictors, PLS should produce superior forecasts compared to OLS. Also, PLS allows for better examination of portfolio weights compared to OLS and PCA. The portfolio weights from PLS should be less volatile over the forecast period compared to OLS, as will be shown in Section 5. Stable weights and weights that are not as extreme as OLS would allow for an investor to better implement the information from the PLS forecasts. Knowing which portfolios matter for the next month's market return can allow the investor to hedge properly. It also would allow for a simple check of a few portfolios to get a quick guess as to which direction the market should be headed in the next month.

4. Data and methodology

Log excess monthly returns are used for all portfolio returns.[4] This paper looks at the forecast power of the 6, 25, and 100 size and book-to-market sorted Fama-French portfolios. This allows for comparison to the results of Kelly and Pruitt (2013), which uses the book-to-market ratios of these portfolios in forecasting. This paper also uses the 5, 10, 12, 17, 30 and 49 Fama-French industry portfolios. The sample period starts January 1930 and ends December 2010. Recursive forecasting starts January 1980, to allow for at least a 50-year estimation

[4] Log returns are used to be consistent with Kelly and Pruitt (2013), which uses log market returns.

window. The sample period and forecasting period is consistent with Kelly and Pruitt. All data (including the market return and risk-free rate data) are obtained from Ken French's website.

The de Jong SIMPLS algorithm (explained in the previous section) is used to perform PLS for forecast purposes. Only the first PLS factor is used to forecast. As mentioned, the algorithm directly gives the regression coefficients:

(1) $B_t = R_t Q_t^T$, where B is the PLS regression coefficient matrix, R is the matrix of X weights, and Q is the matrix of Y loadings.

Note that here X is the matrix of portfolio returns and Y is a vector containing market returns. This is done up to time t (first the X matrix is lagged so that the relationship is $Y_t = X_{t-1}$). Then the forecast for Y_{t+1} is created as follows in order to implement the X information up to time t:

(2) $\hat{Y}_{t+1} = X_t^{'} B_t$

The process is recursively repeated by extending the estimation window to include t+1, while anchoring the starting point. This produces a series of $\hat{Y}$ and a series of realized values of Y. The OOS-R squared statistic based on the historical mean can be produced from here, as in Campbell and Thompson (2008). Also, the B vector can be stored each month to examine the stability of the weights and the forecast. The B vector can also be compared to OLS betas.

5. Forecast results and portfolio weights analysis

(i) Forecast results

Table 1 shows the in-sample and out-of-sample R squared values for each of the test assets. The p-values for the OOS-R^2 are constructed based on the MSPE f-statistic of Clark and

West (2007), as in Li et al. (2013). As can be seen, for all tests assets examined, the OOS-R^2 is positive at a 1% significance level. The methodology is robust to all test assets used. Also, little to no explanatory power is lost when moving out of sample, as can be seen by comparing the in-sample and out-of-sample R^2 statistics. In a few cases, most likely due to chance, the OOS-R^2 is higher than the IS-R^2 statistic.

The 5, 10 and 12 industry portfolios provide the highest OOS-R^2 values, while the 100 size-value sorted portfolios provide the highest IS-R^2 values. Interestingly, this shows that PLS works well even when the set of predictors is not very large. The industry portfolios overall seem to be slightly better at forecasting than the size-value portfolios, but this does not hold for all industry specifications. There does not appear to be a clear relationship between the number of predictors and the R squared values (both in sample and out of sample) for the industry portfolios. Increasing the number of size-value portfolios does appear to increase both in-sample and out-of-sample explanatory power, but the effect is not as drastic as Kelly and Pruitt (2013) finds (see below).

The values are also higher across the board compared to the results of Kelly and Pruitt (2013). Kelly and Pruitt obtain an OOS-R^2 of 0.65%, 0.77%, and 0.93% when using the book-to-market ratios of the 6, 25, and 100 size-value portfolios, respectively.[5] Kelly and Pruitt obtain IS-R^2 values of 0.60%, 1.12%, and 2.38% for these portfolios.[6] Save for the case with 100 portfolio book-to-market ratios, the in-sample results are improved as well. In unreported

[5] They do obtain much higher OOS-R^2 values when forecasting with annual data, which this version of the paper does not do.

[6] The higher IS-R^2 for the 100 size-value (SV) portfolios appears to be a product of the methodology difference. Using the methodology from this paper with the Kelly and Pruitt 100 BTM data still produces a comparatively lower IS-R^2. One additional reason is that Kelly and Pruitt sometimes drop missing observations, of which there are a lot for the 100 portfolios dating back to 1930. This paper simply replaces the missing observations with a 0% return.

regressions, the same test assets are used to forecast market returns using OLS. In all cases, OLS gives a (usually large) negative OOS-R^2.

[Table 1 goes here]

The OOS-R^2 values are statistically significant, but Cochrane (2005) shows a way to measure economic significance for a mean-variance investor. As in Kelly and Pruitt (2013), the Cochrane measure for an active investor's Sharpe ratio, $s^* = \sqrt{\frac{s_0^2 + R^2}{1 - R^2}}$, can be used along with the historical monthly Sharpe ratio of .108 from Campbell and Thompson (2008) to show the potential Sharpe ratio to an investor using the forecast information presented here. Kelly and Pruitt find that their 0.93% OOS-R^2 for the 100 BTM ratios gives a roughly 33% potential Sharpe ratio improvement over a buy-and-hold Sharpe ratio. From the results in Table 1 it can be found that this potential Sharpe ratio improvement ranges from 40% (for the case with the lowest OOS-R^2, the 30 industry portfolios) to almost 48% (using the highest OOS-R^2, from the 5 industry portfolios). The potential SR improvement from these forecasts for an investor using the Kelly and Pruitt data and methodology would range from 4% to nearly 10%.

Table 2 presents some additional results. The first panel in the table shows the "factor" portfolio mean returns and standard deviation for each of the test assets. Returns in the table are annualized excess returns. This shows the performance of a portfolio formed based on the PLS weights. This is similar to the factor portfolios given by PCA in an arbitrage pricing theory setting. The mean and standard deviation of each portfolio can be compared to a buy-and-hold the market strategy that would give a mean excess return of 5.5% and a standard deviation of 4.7%. It can be seen that all but two of the factor portfolios provide a higher return. Further, all but three provide a higher Sharpe ratio. Thus, this provides some evidence against the CAPM and perhaps in favor of an APT model. Note that transaction costs are ignored here, and to invest in the factor portfolio would require frequent rebalancing. However, all weights are positive so there would be no short selling required (unlike if using OLS weights).[7]

The second panel in Table 2 shows the portfolio performance of an investor with a simple trading strategy of investing in the market if the forecast of excess market returns is positive and investing in the risk-free asset if the forecast of market returns is negative. Again, returns in the table are annualized excess returns. Although the mean returns in most cases are lower than the average market return of 5.5% (buy-and-hold strategy), the risk is lower across the board compared to the market. Therefore, in all but one case the Sharpe ratio is higher than that of a buy-and-hold the market strategy. However, the factor portfolios from the first panel have higher Sharpe ratios in most cases. So it appears that there is a way for investors to better take advantage of the forecast information than with this simple strategy (a possibility is mentioned in the next section). Again, this simple analysis ignores trading costs, so a direct comparison to the returns of a buy-and-hold the market strategy should not be made. The adjusted returns in the second panel would likely be closer to those from the first panel after trading costs adjustments were made as well.

[Table 2 goes here]

(ii) Raw portfolio weight analysis

For reasons discussed earlier, PLS gives portfolio weights that should provide superior information to OLS weights. The other benefit is that the forecast performance is better. With

[7] The OLS weights produce factors with much lower Sharpe ratios (unreported). Using OLS weights would also require much more rebalancing due to their higher variance (see Table 3).

more explanatory power, the portfolio weights will be more useful. PLS weights are diffuse (as is the case with shrinkage estimators) and positive (no short selling required). Principal-component loadings would contain no information about market returns, unlike PLS weights. Table 3 also shows that PLS weights are much less volatile over the forecast sample compared to OLS weights. Since a new estimation and forecast is done each month, the set of betas/weights can be stored for each month to produce a panel of weights. The table below shows the standard deviation of the weights for each of the test assets, for both PLS and OLS. The standard deviation reported is based on an average of the 6, 25, etc., time series of individual weights. The standard deviation for the OLS weights range from 1.4% to 4.4%, while the PLS standard deviation is close to zero, the highest being 0.1%.

[Table 3 goes here]

With the PLS weights being much more stable, this would help an investor wanting to utilize the information. If the investor wanted to re-estimate the weights each month for either rebalancing or creating a new forecast, then there would be little cost to doing so (much less cost compared to OLS). Also, since the weights change little over time (the results in Table 3 are over the forecast sample of 30 years), the investor or practitioner could simply estimate the weights once and continue to use them (perhaps re-estimating every 1-2 years).

Since the portfolio weights from PLS should be more reliable than those from PCA or OLS, it allows for examination of which portfolios explain future market returns well. Table 4 shows the average weights on small size, large size, high book-to-market, and low book-to-market portfolios. This only looks at the smallest-size-firm portfolios and largest-size-firm portfolios (highest BTM portfolios and lowest BTM portfolios), with the portfolios formed on intermediate quantiles unreported here. This allows for a quick check of whether small or large size portfolios (and growth and value portfolios) matter more for future market returns. The first column in each of the three panels shows how much weight the small/large/growth/value portfolios were given in forecasting.[8] For example, the percentage given for small size portfolios is the ratio of weights put on small size portfolios compared to the weights of all portfolios. The second column in each panel shows where the portfolios rank on average. All portfolios in the group of test assets are sorted based on weight, with the portfolio with the highest weight being ranked first. Then the average rank for each characteristic is found, so that higher average weight means a smaller rank number. The table shows the results for the 6, 25, and 100 size-value sorted portfolios.

Examination of the results in Table 4 shows that small size portfolios carry a greater weight than large size portfolios. Although the percentage of total weight given to small versus large portfolios may not look significantly different, it is important to remember that weights are very diffuse in PLS. For comparison purposes, note that the difference between the highest weighted portfolio and the lowest weighted portfolio is roughly 5%, 2.5%, and 2% for the 6, 25, and 100 portfolios respectively. Thus, it can be seen that small size portfolios have a significantly higher weight than large size portfolios in explaining future market returns for each group of test assets. The pattern holds when looking at the intermediate size quantiles for the 25 and 100 portfolios as well, which is not shown in the table. The closer link between small firm returns and future market returns may be part of the explanation for why small-firm stocks have

[8] The weights were calculated based on full-sample estimation (i.e., the last forecast estimation run), but since weights are stable over time, the results obtained by averaging the weights from each monthly forecast are not significantly different.

higher returns. They are riskier because they signal a change in investment opportunities: low returns on small stocks followed by low returns in the wider market.

There is no clear pattern between growth and value portfolios in the table, however. For the set of six portfolios, value stocks have a higher weight on average. The gap closes for the set of 25 portfolios, though, and reverses for the set of 100 portfolios (low BTMs have a higher weight). When looking at the intermediate book-to-market portfolios (unreported) rather than just the highest and lowest, it does appear that higher book-to-market portfolios (value stocks) on average have a higher weight when forecasting the market return. Thus, again the higher returns of value stocks may be explained in part by their higher risk in that low returns on value stocks are followed by low market returns. Nonetheless, the book-to-market difference is not quite as clear and evident as the size difference.

[Table 4 goes here]

The weights on the industry portfolios can be examined as well to see which industries matter in explaining the next month's market return. One issue is that the number of industries specified is crucial to the results. However, a few general results can be gleaned from examining the weights for all the different industry portfolios.[9] The energy industry (oil, mines, coal in some specifications) consistently has a low weight. Other industries that have a low weight across all specifications include telecommunications, food, soda, and clothing. The finance/money industry consistently has a high weight, perhaps evidence against dropping these firms when collecting data. Note also that the financial sector has a lot to gain from improving market conditions and thus it makes sense that the stock returns in this sector are useful leading

[9] The industry weights are unreported, but available upon request.

indicators. In two of the seven specifications, the entertainment industry has the highest weights. Durables, paper, high-tech industries, books, automobile, and insurance also have high weights.

Future versions of this paper will explore the implication of portfolio weights further. One possibility is using the weights to hedge against market risk. The objective would be to minimize variance by holding some combination of the test asset portfolios according to the weights. Knowing which portfolios should be weighted higher than others can both allow an investor to update his or her prediction of the market and perhaps hedge against market risk.

(iii) Principal-component weight analysis

An additional test is performed to investigate whether the results presented so far are spurious. This will help avoid the brunt of the criticism of Novy-Marx (2014) mentioned in the introduction. Rather than running PLS on the raw portfolio returns, first PCA is performed. From there, the methodology for forecasting is the same as above. For example, with the 25 size-value portfolios, PCA is performed so that now there are 25 factors instead of the 25 raw portfolios. Note that these factors are simply linear combinations of the original returns, representing the same information. Therefore, using these factors to forecast with PLS will not change the explanatory power (R squared values do not change). However, now the weights should not be diffuse if there is indeed some predictive information contained in the test assets. The first few factors from PCA should contain most of the information from the portfolio returns. Other factors should have relatively little weight.

From looking at Figures 1 and 2, it can be seen that this is indeed the case. The figures show the absolute value of the PLS weights, where the test assets are the PCA factors. Figure 1 shows the case for the 25 PCA factors obtained from the 25 size-value portfolios. Figure 2 shows the case for the 30 PCA factors obtained from the 30 industry portfolios. The same pattern holds for all other test assets, so the figures are omitted. If the results were indeed spurious, then the weights on the factors would be random. However, this is not the case here. The first factor seems to capture all the relevant information, as it is the only factor with substantial weight. For the 25 factors in Figure 1, the first factor has a relative weight of around 57% compared to all other factors (no other factor has a relative weight above 7%). For the 30 factors in Figure 2, the first factor has a relative weight of 55%, and no other factor has a relative weight above 5%.[10] It seems that the PLS approach (whether with raw portfolio returns or PCA factors) captures a single risk factor which has predictive power. Most likely this is simply the market risk factor. This is confirmed by forecasting excess market returns with an AR(1) model. The AR(1) model gives an out-of-sample R squared value of 1.2%, which is roughly in the middle of the results obtained by PLS. Note that this is also higher than the results of Kelly and Pruitt (2013). Therefore, PLS and the 3PRF seem to offer little improvement over an AR(1) model; however, PLS does do a good job of capturing the market risk factor. PLS also allows for examination of which portfolios capture market risk and therefore do well in forecasting (see previous subsection).

[Figures 1 and 2 go here]

6. Conclusion

This paper offers three main contributions to the literature. The first is to improve on the out-of-sample performance when forecasting market returns of Kelly and Pruitt (2013) by both using different predictors and slightly changing the methodology from a 3PRF to a more

[10] Other test assets place a relative weight on the first factor of up to 86%. The number of test assets/factors is crucial to this weight; the more factors there are, the more the total weight is spread out.

traditional PLS approach. It is shown that the results are both statistically and economically significant. Further, the improvement over the Kelly and Pruitt results offers a mean-variance investor a potential 10% Sharpe ratio improvement.

The second contribution of the paper is to examine the portfolio weights that result from forecasting with PLS. It is shown that small size portfolios have more explanatory power than large size portfolios. Portfolios with higher book-to-market ratios also have a higher weight on average, but this is not as drastic as it is with the small size portfolios. Several industries are found to be consistently important in explaining future market returns. This paper also looks at a "factor"-type portfolio that could be formed based on the PLS weights. Most of the test assets examined produce a factor portfolio with a Sharpe ratio that is superior to a buy-and-hold the market strategy Sharpe ratio.

The final contribution of the paper is to address a concern raised by Novy-Marx (2014): is the predictive power due to beneficial information or is it simply due to chance? By using PCA factors as predictors, it is found that most of the PLS weight is placed on the first factor. If the results were random, then random (most likely diffuse) weights would be placed on all factors. Further, with only the first factor having significant weight, it is argued that only one risk factor matters: market risk. The predictive power of an AR(1) model offers evidence in favor of this argument. Nonetheless, using PLS to forecast can still offer slightly improved predictive power and allows for valuable examination of portfolio weights.

Both investors/practitioners and future research could benefit from the information in the portfolio weights. The PLS weights are found to be much more stable and less extreme (no negative weights) compared to OLS. Also, the forecast power is significantly greater than OLS.

Kelly and Pruitt (2012) find that the 3PRF (a PLS variant) performs significantly better than principal-components analysis. Since PLS better forecasts market returns compared to PCA and OLS and produces more reliable weights, the PLS results offer many benefits. One potential benefit would be using the weights to hedge against market risk, something a future version of this paper will attempt to do. Overall, though, there are many potential applications of PLS to finance, especially in any forecasting environment.

References

Bossaerts, P., and Hillion, P. 1999. "Implementing statistical criteria to select return forecasting models: what do we learn?" *Review of Financial Studies* 12 (2): 405-428.

Campbell, J. Y., and Thompson, S. B. 2008. "Predicting excess stock returns out of sample: Can anything beat the historical average?" *Review of Financial Studies* 21 (4): 1509-1531.

Clark, T.E., and West, K.D. 2007. "Approximately normal tests for equal predictive accuracy in nested models." *Journal of Econometrics* 138 (1): 291-311.

Cochrane, J.H. 2005. *Asset Pricing*. Princeton University Press.

De Jong, S. 1993. "SIMPLS: an alternative approach to partial least squares regression." *Chemometrics and Intelligent Laboratory Systems* 18 (3): 251-263.

Ferreira, M.A. and Santa-Clara, P. 2011. "Forecasting stock market returns: the sum of the parts is more than the whole." *Journal of Financial Economics* 100 (3): 514-537.

Goyal, A., & Welch, I. 2003. "Predicting the equity premium with dividend ratios." *Management Science* 49 (5): 639-654.

Guo, H. 2006. "On the out-of-sample predictability of stock market returns." *Journal of Business* 79 (2): 645-670.

Hoskuldsson, A. 1988. "PLS regression methods." *Journal of Chemometrics* 2 (3): 211-228.

Kelly, B.T., and Pruitt, S. 2012. "The three-pass regression filter: A new approach to forecasting with many predictors." Working paper, Chicago Booth.

Kelly, B.T., and Pruitt, S. 2013. "Market expectations in the cross-section of present values." *Journal of Finance* 68 (5): 1721-1756.

Lettau, M., and Ludvigson, S. 2001. "Consumption, aggregate wealth, and expected stock returns." *Journal of Finance* 56 (3): 815-849.

Li, Y., Ng, D.T., and Swaminathan, B. 2013. "Predicting market returns using aggregate implied cost of capital." *Journal of Financial Economics* 110 (2): 419-436.

Novy-Marx, R. 2014. "Predicting anomaly performance with politics, the weather, global warming, sunspots, and the stars." *Journal of Financial Economics* 112 (2): 137-146.

Polk, C., Thompson, S., and Vuolteenaho, T. 2006. "Cross-sectional forecasts of the equity premium." *Journal of Financial Economics* 81 (1): 101-141.

Welch, I. and Goyal, A. 2008. "A comprehensive look at the empirical performance of equity premium prediction." *Review of Financial Studies* 21 (4): 1455-1508.

Table 1: Market return forecast results

This table shows the R-squared values from forecasting. Out-of-sample (OOS) R-squared is based on the historical mean of the market return. OOS p-value is the p-value obtained by regressing the Clark and West (2007) adjusted-MSPE statistic on a constant, as in Li et al. (2013). IS R-squared is the percent variance of Y (market returns) explained by the first X-component. SV are Fama-French size-value sorted portfolio returns. Ind are the various Fama-French industry portfolio returns. Monthly data is used. Data is obtained from Ken French's website. Recursive forecasting is performed with the estimation sample starting at January 1930 (anchored starting point) to December 1979. Forecast sample starts at January 1980 and runs to December 2010. PLS regression estimates obtained from the SIMPLS algorithm of de Jong (1993) are used to forecast. Excess log returns are used, the CRSP value-weighted market index is used for market returns, and 3-month US Treasury bills are used for the risk-free rate.

	6 SV	25 SV	100 SV	5 ind	10 ind	12 ind	17 ind	30 ind	49 ind
OOS R-squared	1.10%	1.17%	1.24%	1.35%	1.25%	1.32%	1.11%	1.09%	1.12%
OOS p-value	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01
IS R-squared	1.18%	1.12%	1.34%	1.27%	1.29%	1.32%	1.18%	1.27%	1.31%

Table 2: mean-variance performance

This table shows the mean excess return and the standard deviation for two different portfolios for each of the test assets. The first, the factor portfolio, is a portfolio constructed based on the portfolio weights obtained from the PLS regression. This is done using the full-sample coefficients. The second portfolio is the portfolio performance that an investor could achieve based on the forecast results of PLS. The trading strategy is as follows: if the predicted excess market return is positive, invest in the market; otherwise, invest in the risk-free asset. SV are Fama-French size-value sorted portfolio returns. Ind are the various Fama-French industry portfolio returns. Monthly data is used. Recursive forecasting is performed with the estimation sample starting at January 1930 (anchored starting point) to December 1979. Forecast sample starts at January 1980 and runs to December 2010. PLS regression estimates obtained from the SIMPLS algorithm of de Jong (1993) are used to forecast. Excess log returns are used, the CRSP value-weighted market index is used for market returns, and 3-month US Treasury bills are used for the risk-free rate. Returns below are monthly, annualized excess returns.

	Factor Portfolio		Forecast trading strategy	
	Mean excess return	Standard deviation	Mean excess return	Standard deviation
6 SV	7.88%	5.35%	5.20%	4.13%
25 SV	7.00%	5.13%	5.17%	4.14%
100 SV	7.28%	5.12%	4.53%	4.09%
5 ind	6.00%	4.57%	5.30%	4.17%
10 ind	5.82%	4.64%	5.87%	4.11%
12 ind	5.63%	4.71%	5.93%	4.11%
17 ind	5.28%	4.98%	4.63%	4.20%
30 ind	5.53%	5.10%	4.78%	4.09%
49 ind	5.21%	5.15%	4.99%	4.10%

Table 3: Comparison of OLS and PLS coefficient volatility

This table compares, for each of the test assets, the standard deviation of OLS coefficients and PLS coefficients over the forecast sample period. Each time the recursive forecast regression is run, the coefficients for each portfolio are stored. Then for every portfolio's coefficients, the standard deviation is calculated. Reported in the table is the average standard deviation of the respective group of test assets. Recursive forecasting is performed with the estimation sample starting at January 1930 (anchored starting point) to December 1979. Forecast sample starts at January 1980 and runs to December 2010. PLS regression estimates obtained from the SIMPLS algorithm of de Jong (1993) are used to forecast. Excess log returns are used, the CRSP value-weighted market index is used for market returns, and 3-month US Treasury bills are used for the risk-free rate.

	6 SV	25 SV	100 SV	5 ind	10 ind	12 ind	17 ind	30 ind	49 ind
OLS std dev	4.42%	2.53%	1.71%	1.97%	1.95%	2.05%	1.87%	1.43%	1.41%
PLS std dev	0.06%	0.02%	0.01%	0.12%	0.06%	0.05%	0.04%	0.02%	0.01%

Table 4: Summary of size and value weights

This table gives a summary of weights on various size and value portfolios, obtained from PLS regressions that forecast the market return. % of weight is the sum of the coefficients for each of the respective size or value quantile's portfolios divided by the total sum of the coefficients. For example, small size gives the ratio of the weight put on small size portfolios to that of all weights (done for 6, 25, and 100 Fama-French size and book-to-market sorted portfolios, obtained from Ken French's website). Average rank shows where the average size or value sorted portfolio ranks in terms of weight compared to all the other portfolios. Rank is done by highest weight first, lowest weight last. Recursive forecasting is performed with the estimation sample starting at January 1930 (anchored starting point) to December 1979. Forecast sample starts at January 1980 and runs to December 2010. PLS regression estimates obtained from the SIMPLS algorithm of de Jong (1993) are used to forecast. Excess log returns are used, the CRSP value-weighted market index is used for market returns, and 3-month US Treasury bills are used for the risk-free rate.

	6 size-value sorted portfolios		25 size-value sorted portfolios		100 size-value sorted portfolios	
	% of weight	average rank	% of weight	average rank	% of weight	average rank
small size	51.09%	3.3	21.50%	10.2	10.55%	47.1
large size	48.91%	3.7	17.67%	17.4	7.03%	65.6
high BTM	37.28%	1.5	19.79%	13.6	7.41%	64.7
low BTM	30.24%	5.0	19.63%	13.4	8.52%	64.7

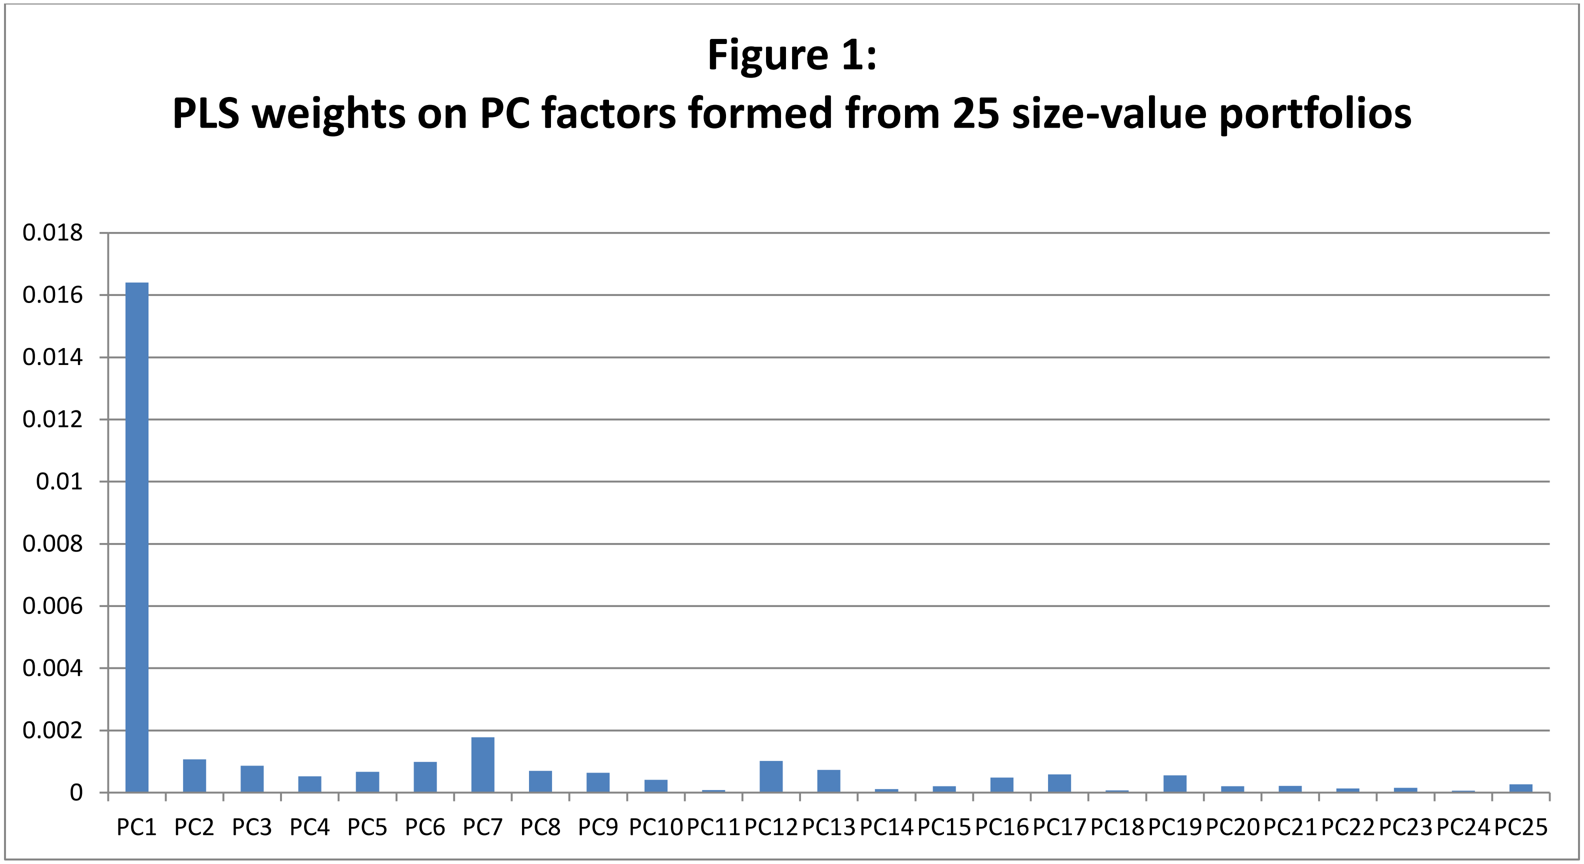

This figure shows the PLS weights when PCA factors are used as the predictor variables when forecasting. Here, 25 PCA factors are formed from the 25 size-value Fama-French portfolios. The weights are obtained from the PLS forecast regression results where excess market returns are the dependent variable. Recursive forecasting is performed with the estimation sample starting at January 1930 (anchored starting point) to December 1979. Forecast sample starts at January 1980 and runs to December 2010. PLS regression estimates obtained from the SIMPLS algorithm of de Jong (1993) are used to forecast. Excess log returns are used, the CRSP value-weighted market index is used for market returns, and 3-month US Treasury bills are used for the risk-free rate.

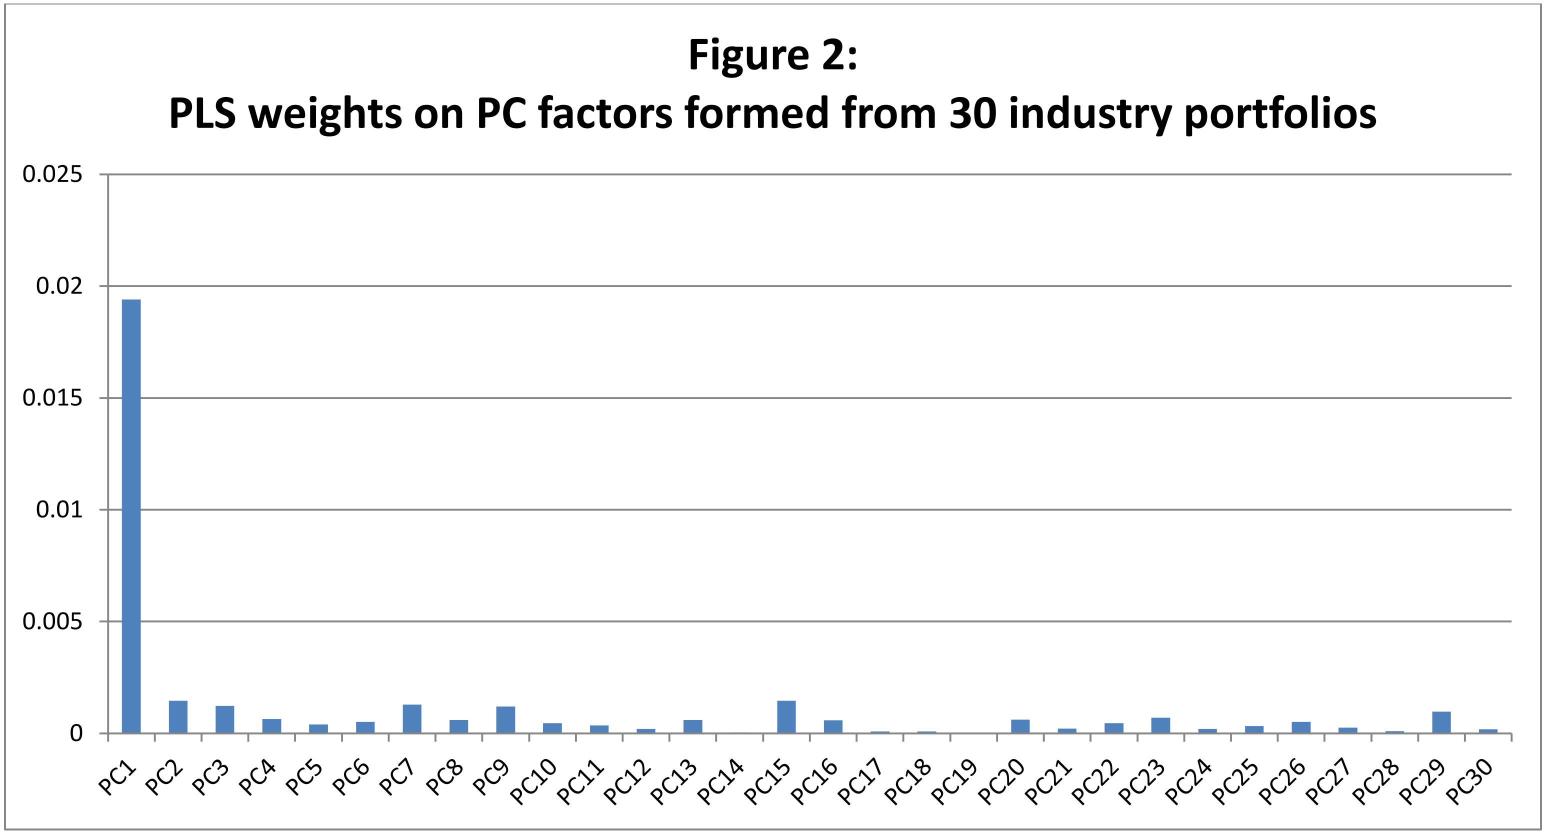

This figure shows the PLS weights when PCA factors are used as the predictor variables when forecasting. Here, 30 PCA factors are formed from the 30 Fama-French industry portfolios. The weights are obtained from the PLS forecast regression results where excess market returns are the dependent variable. Recursive forecasting is performed with the estimation sample starting at January 1930 (anchored starting point) to December 1979. Forecast sample starts at January 1980 and runs to December 2010. PLS regression estimates obtained from the SIMPLS algorithm of de Jong (1993) are used to forecast. Excess log returns are used, the CRSP value-weighted market index is used for market returns, and 3-month US Treasury bills are used for the risk-free rate.