

ESSAYS ON MEASURING ASSET PRICING ANOMALIES

By

MICHAEL GORMAN

A dissertation submitted to the

Rutgers Business School, Newark and New Brunswick

Rutgers, The State University of New Jersey

In partial fulfillment of the requirements

For the degree of

Doctor of Philosophy

Graduate Program in Finance

Written under the direction of

YANGRU WU

And approved by

Newark, New Jersey

May 2016

©2016

MICHAEL GORMAN

ALL RIGHTS RESERVED

ESSAYS ON ASSET PRICING ANOMALIES

By Michael Gorman

Dissertation Director: Dr. Yangru Wu

Abstract

Traditional methods of measuring asset pricing anomalies have historically relied on full sample tests of static parameters. With the increase of computational power and data available we are able to allow for time varying factor loadings for a portfolio based on asset rotation and also time varying factors by asset.

In the first paper we find that commonly used estimates of time varying asset pricing anomalies contain significant bias. We are able to show that the historical returns used to select momentum portfolios result in biased data in the short window asset level regressions which the literature uses to estimate portfolio parameters. This is caused through a non-random selection criterion which systematically chooses high epsilon assets. These nonrandom epsilons, when regressed upon bias estimates of Alpha, and through the correlation structure of the parameters they also bias the estimates of Beta.

We present a new methodology that is not subject to this bias, and allows for an accurate measurement of the size of anomalies. In executing this we find that inefficient portfolio rotation in the original portfolio level estimates is also indicative of bias. As such we suggest that the new methodology we propose is more accurate and less susceptible to bias than those currently in use in the literature. The new model suggests that to this point the risk adjusted returns of the momentum portfolio have been underestimated in the literature.

In the second paper we demonstrate that the momentum anomaly is driven by a small number of assets in the market using our new model and the methodology of False Discovery Rates. We show that these assets, behave differently in long and short portfolios, and also perform differently during the first month reversal period. Finally we demonstrate that an appropriately risk adjusted momentum alpha shows that extreme months are not sufficient to explain away momentum, and that poor returns in extreme months are overstated by traditional methods of measuring momentum. To this extent we claim that market downturns in the last 15 years have been insufficient to effectively eliminate the momentum anomaly as has been suggested.

Preface

The following work was originally written as two separate papers. For the purpose of this document Chapter 1 covers the material of the first paper, and Chapter 2 the material of the second paper.

Chapter 1 primarily deals with methodological improvements to the methods used to evaluate the momentum anomaly. Particular attention is paid to the bias induced by some methods currently being used in the literature, and improvements to those methods.

Chapter 2 covers the use of this method to generate results of interest, of particular note demonstrating that momentum survives with significant positive risk adjusted returns over the most recent period (since 2000) despite poor raw returns and poor traditional risk adjusted returns.

Table of Contents

Abstract.....	ii
Preface	iv
Table of Contents.....	v
List of Tables	vi
List of Figures	vii
Chapter 1.....	viii
I. Data and existing methodology	xv
II. Existing estimates and issues in the estimation of parameters.....	xxi
III. Estimates skipping the sorting period.....	xxiii
IV. Proposed methodology.....	xxv
V. Simulations.....	xxxi
VI. Conclusions	xxxvi
References	xxxix
Tables	xli
Figures.....	l
Chapter 2.....	lii
I. Data and Methodology	lvii
II. Results.....	lxvii
Using False Discovery Rate methodology to analyze momentum.....	lxix
Is Momentum changing over time?.....	lxxiv
Has Momentum disappeared?	lxxx
III. Conclusions	lxxxiv
References	lxxxvii
Tables	xc
Figures.....	xcv

List of Tables

Chapter 1

Returns of momentum portfolios during and after sorting	xli
Asset level regressions including the sorting period	xlvi
Asset level regressions excluding the sorting period	xlvi
Changing Beta estimates of momentum portfolios over time	xlvii
Monthly and daily data comparison for asset level tests	xlviii
Momentum alpha comparing 12 month and 1 month regressions	xlix
Regressions on holding period return using 1 month	l
Churn rates for momentum portfolios	li
Comparing new to existing methodology for momentum	lii

Chapter 2

Risk adjusted monthly momentum alpha by decade	xc
Alpha and FDR estimate by investment month	xc
Alpha under and over the \$5 cutoff	xcii
Raw momentum returns and alpha for extreme months	xciii
Comparison of adjusted monthly alpha with traditional estimates	xciv

List of Figures

Chapter 1

Momentum returns before, during and after sorting periods I

Simulating estimated Alpha in momentum strategies li

Chapter 2

Alpha of the Long Momentum Portfolio xcv

Alpha of the Short Momentum Portfolio xcvi

Chapter 1

Bias in Momentum Estimates

Abstract

We demonstrate that constructing momentum portfolios using traditional methodology in both portfolio level and asset level tests produces biased regression parameters. We suggest an alternative methodology that is not subject to the bias. Using this new methodology we compute more accurate estimates of risk adjusted returns which show standard momentum strategies to have risk adjusted returns of 40 basis points higher than previously estimated.

It has been well demonstrated that the CAPM does not fully account for the cross section of stock returns. Momentum portfolios are one example that has been used to illustrate the difficulties experienced by the CAPM. Portfolios of assets with high historical returns have been shown to consistently outperform portfolios of stocks with low historical returns over an intermediate time horizon of 3 to 12 months.

The issue of momentum in stock prices was covered in depth by Jegadeesh and Titman (1993) looking at portfolios of assets formed by ranking the preceding 3 to 12 months of data. Momentum was shown to persist across 32 combinations of portfolios providing excess returns of around 1% per month.

Since the initial paper there have been several follow up studies that have attempted to validate the original results through out of sample testing (Jegadeesh and Titman, 2001) and those that have expanded the scope of momentum to consider broader classes of assets (Fama and French, 2012; Moskowitz, Ooi and Pederson, 2012 and Asness, Moskowitz and Pedersen, 2013). The general consensus is that through many trials in different markets, time periods and asset classes, momentum is robust over a number of strategies.

While factors have been suggested that partially explain momentum there have not yet been compelling arguments made that are able to completely eliminate the anomaly. Pastor and Stambaugh (2003) and Sadka (2006) provide evidence to suggest that momentum is at least correlated with liquidity in US stocks, with Asness, Moskowitz and Pedersen (2013) expanding this to further markets and asset classes. Chordia and Lakshmanan (2002) use macroeconomic variables as an explanatory factors and reduce

the impact of momentum substantially, although this is countered by work from Cooper, Gutierrez and Hameed (2005) who do not find macroeconomic factors to be helpful. Moskowitz and Grinblatt (1999) use industry factors to attempt to explain momentum, which is shown to at least reduce but not eliminate momentum profits. Extensive efforts have also been made to explain momentum through behavioral models (e.g. Daniel, Hirshleifer, and Subrahmanyam, 1998; Barberis, Shleifer, and Vishny, 1998 and Hong and Stein, 1999), although these avenues are undermined by papers such as Assness, Moskowitz, and Pedersen (2013) who find evidence indicating the presence of common global risk factors, for which momentum (amongst other anomalies) may compensate.

Several studies have also argued that allowing a time varying estimate of portfolio Beta has helped to explain some of the more persistent anomalies such as the size and book to market effects. This has been extensively covered by Jagannathan and Wang (1996), Lettau and Ludvigson (2001), Santos and Veronesi (2006) and Lustig and Van Nieuwerburgh (2005). This was also considered for momentum by Wang and Wu (2011), the methodology of which we shall return to later in the paper. This is a relatively recent shift in the momentum literature and is distinct from the methodology proposed in Jegadeesh and Titman (1993) which relied on static parameter estimates. While these estimates were able to drastically reduce the size of size and book to market anomalies the momentum anomaly proved more persistent in the face of this methodology. The persistence of this anomaly remains the most enduring thorn in the side of efficient market theory.

In this paper we question whether the current commonly used methodology accurately measures the momentum anomaly. Our analysis first looks at asset level tests of momentum to establish where bias is introduced to the process. Secondly the paper looks at the mechanisms by which this bias is transferred into the portfolio level estimates currently favored in the literature. Last the paper presents a methodology that is not impacted by the bias identified and which allows the unbiased measurement of momentum alpha. Our improved measures of the momentum anomaly show that the excess returns generated by momentum are significantly higher than previously thought, by approximately 40 basis points per month.

The initial proposition of this paper is to establish bias in existing estimates of momentum parameters. The issue that makes momentum unique to this study amongst anomalies is the strategy by which it is implemented. As a technical trading strategy, momentum uses return data to select assets. By selecting based on extreme realized returns it is not unrealistic to expect that such returns are formed by expected excess return, market risk exposure, and idiosyncratic risk realizations that are not typical of the market. These observations being non-normal, selected from distributional tails, impacts the assumptions made by the standard models.

For an asset level test it is common to use the most recent return data available to estimate the firm's parameters. In the case of the momentum portfolio this often means that the period of returns used to select assets is included in the regression. Lewellen and

Nagel (2006)¹ used as few as 12 observations (1 year of monthly data) to conduct their conditional CAPM estimates, however more data is more common (Wang and Wu, 2011, use 3 years). Observations in the portfolio formation period can be extreme in magnitude, as we are selecting extreme deciles of data for use in momentum.

Given that there are large shocks hitting the assets through the portfolio formation period, and that our selection methodology selects for these shocks in a systematic manner, we now have data with a shock that does not have an expected value of zero. It is intuitively sensible to think that, for example, assets that have observed unusually high firm specific shocks are more likely to end up being in the top performing assets selected for a momentum portfolio than their peers with negative firm specific shocks.

Forcing the model to set the sum of error terms to zero has the simple impact of transferring the mean of the error term caused by the selection bias through to the alpha parameter. This leads to an overestimate (underestimate) of the alpha parameter in the model in the instance of the portfolio being formed from assets that have been in the highest (lowest) recent return decile.

The simplest method to validate this is to simply run regressions on the data at the asset level and observe the changes in alpha estimates when regressions are run using the most recent data available, which by definition includes the sorting period. Using this methodology we should expect to see increasing estimates of alpha (due to bias) for

¹ Extensive work on the conditional CAPM has been done by Wang (2003), Ang and Chen (2006) and Petkova and Zhang (2005). Further, tests of the conditional CAPM such as Harvey (1989), Shanken (1990), Jagannathan and Wang (1996) and Lettau and Ludvigson (2001) rely on state variable choice which poses challenges. Lewellen and Nagel (2006) side step this issue by simply assuming that asset parameters are stable within short windows, allowing them to avoid specifying conditioning information sets.

longer sorting periods, and generally higher alpha estimates than are reported in the momentum literature. This is confirmed by tests in this paper.

To get around this problem, one approach utilized in the literature (Lewellen and Nagel, 2006 and Wang and Wu, 2011) is to run these regressions but discard the alpha estimates, and utilize instead the beta estimates. Using beta estimates along with holding period returns a researcher can then identify implied holding period excess returns assuming an independent alpha and beta to allow unbiased estimation.

The assumption, is also problematic. It has been documented since at least Jegadeesh and Titman (1993) that assets included in long momentum portfolios have abnormally high alpha and beta compared to the market. The positive correlation between these two factors is problematic, as bias in the construction of alpha is also likely to be reflected in bias of a similar directionality, but different magnitude, in the beta estimates. This would in turn impact the implied estimates of alpha for the holding period, as the beta estimates used in the calculation would be biased. Given these problems it would appear that this methodology in current form is not fit for purpose. There is also an apparent relationship between alpha and beta in the short portfolio, with higher beta assets being systematically selected, which suggests not only that the relationship exists in both portfolios (exposing both to bias), but that the relationship is nonlinear. The functional form of this relationship is not the focus of this paper, but it suffices to say that the relationship does call into question estimates based on a lagged beta estimate.

Given the issues with asset level estimates of momentum alphas the simplest response would be to obtain alpha estimates by running alpha estimates at the portfolio level.

A potential problem with portfolio level estimates is that they are subject to the same biases established for asset level estimates. The asset level estimates are challenged by the use of most recent data (including sorting periods) and the bias this induces into the parameter estimates. On the surface it would seem that the portfolio level tests are immune from such bias because they only regress on data from the holding period. In theory, given an efficient market, this should not over or under sample from the holding periods of assets that were included in the long or short portfolios. Unfortunately the assumption of efficient markets is easily violated when measuring anomalies.

While it is true that only the holding period is regressed upon, there is a non-efficient churn observed with momentum. The issue arises when observing the assets sold and purchased over a given year. Rates of portfolio churn are significantly higher than would be expected in an efficient market. This higher churn means an asset currently in the portfolio is less likely than a randomly selected asset to be included in the replacement portfolio formed 6 months later.

This underrepresentation causes a problem. We now know that the assets included in the portfolio at any given time will have their sorting period underrepresented in the regression. Given that we know that the epsilon estimates of the long portfolio sorting period are expected to be highly positive, and that they are underrepresented, this implies that there is a downwards bias on alpha estimates extracted from portfolio level

momentum regressions for the long portfolio. A similar logic can be used to show bias in the short portfolio. In theory if the biases were of similar magnitude it could be argued that they are offsetting, however the churn of the short portfolio is not generally the same distance from the efficient expected churn as the long portfolio, and the ratio of differences is not constant over portfolio formation periods, leading to unstable bias, which is almost certainly not fully offsetting.

Given the issues outlined, it is natural to look for a more accurate measure of momentum alpha. Using an asset level holding period only estimate of alpha we are able to generate models that are protected to the greatest degree possible against accusations of bias, and demonstrate alpha estimates that are significantly higher than conventional estimates.

I. Data and existing methodology

The return data used in this study is collected from the Center for Research in Security Prices (hereafter CRSP). Data is collected from 1960 to 2012, and consists of both daily and monthly returns. The standard momentum methodology requires two years of data to be included, so return data for the portfolio begins in 1962. Market factors used for CAPM regressions are obtained from Ken French's website. These include monthly and daily estimates of the return on the market as well as the risk free rate of return.

Assets are considered for our study if they are traded on the NYSE, AMEX or NASDAQ exchanges. All primary domestic stocks are considered (CRSP share code 10 or 11) which excludes assets such as Real Estate Investment Trusts (REITs) and American Depository Receipts (ADRs).

The methodology used to calculate momentum returns is broadly consistent with that of Jegadeesh and Titman (1993), Asness (1994), Fama and French (1996), and Grinblatt and Moskowitz (2004). To be included in our portfolio an asset has to have two continuous years of return data present, and at the time of investment have a price of greater than \$5. An asset that survives the filters has returns calculated for a defined holding period of the previous 3 to 12 months returns. These returns are then ranked every month.

Each month, based on the ranking, a zero cost portfolio is formed using a long portfolio consisting of the top decile of ranked assets, and a short portfolio consisting on the bottom decile of ranked assets. The assets in their respective portfolios are equally weighted. It is worth noticing that some studies use the top third of assets, instead of the top decile (Jegadeesh and Titman, 2001).

These portfolios are then held for 3 to 12 months. For a 12 months holding period, each month one portfolio expires, and a new one is formed, so that at any given time the investment strategy will have 12 long portfolios and 12 short portfolios, all set up at zero total cost. The portfolio return is then calculated given the returns of the individual assets which are all equally weighted in this portfolio.

Common practice in the literature has evolved to add a skipped month between sorting and holding periods. The first month post sorting period for the zero cost portfolio has historically had poor returns, which is often attributed to liquidity or microstructure issues. Jegadeesh (1990), Lo and MacKinley (1990), Jegadeesh and Titman (1993), Boudoukh, Richardson and Whitelaw (1994), Asness (1994), Grinblatt and Moskowitz (2004) all skip that first month for at least a portion of their results. This procedure is common, and this paper implements all tests both with and without skipping that month. The results are not significantly different, with slightly higher Alpha estimates generally observed for the methodology that implements the skipped month. While only one set of results is included in the following paper, the results using the alternative method are available from the authors upon request.

Given the returns of the portfolio the standard method of measuring excess return is then to run a portfolio level regression using monthly data. This methodology was popularized in the momentum literature in Jegadeesh and Titman (1993). This specification involves running a single regression using the entire sample of portfolio returns available. In our case, using over 50 years of available data, we have access to over 600 observations which is more than sufficient degrees of freedom with which to generate statistically significant results and relatively stable estimates.

Standard methods of estimating excess return involve simple CAPM regressions, or a Fama French three factor model as presented in Fama and French (1995). Historically these two model specifications have not led to a large observed difference in excess

return estimates for the momentum anomaly. The consistency of these two methods is further confirmed in this paper. The respective models are listed below.

CAPM Model:

$$R_{i,t} - R_{f,t} = \alpha_i + \beta_i (R_{m,t} - R_{f,t}) + \varepsilon_{i,t} \quad (1)$$

Fama and French three factor model:

$$R_{i,t} - R_{f,t} = \alpha_i + \beta_{i,m} (R_{m,t} - R_{f,t}) + \beta_{i,smb} SMB_t + \beta_{i,hml} HML_t + \varepsilon_{i,t} \quad (2)$$

The only difference between the models is the inclusion of terms to allow for book to market and size anomalies, with the models otherwise being consistent.

An asset level test is an alternative method to calculate excess return. In this case, rather than simply perform a single regression at the portfolio level, each asset in the market has a regression performed at each month of interest using historical data to estimate the parameters of interest for that asset. Once asset level parameters have been estimated the parameters of assets from the portfolios are then combined to form an estimate of the portfolio parameters.

In a theoretical efficient market these two models should arrive at the same parameter estimates. Unfortunately when measuring anomalies it is rarely appropriate to assume that markets are efficient. We will briefly outline some of the common issues already identified in the literature when using these methodologies.

Pitfalls of using portfolio level estimates include a failure to account for dynamic portfolio parameters and constant estimates of historical excess returns.

When estimating at the portfolio level, a portfolio will be assigned a single parameter explaining sensitivity to market volatility (Beta) for the entire sample, and based on this

an excess return is calculated. This can cause problems when the portfolio changes dynamics dramatically during the sample. If the portfolio generates more or less excess return from this market timing it will be erased by the portfolio level estimate, which obscures such information in favor of a simple output.

Another drawback of the portfolio level tests is that they require a full sample to run. In this case all of the data is used to estimate the excess return, which means that the estimate of excess return for the initial years of the sample will continue to change as more years are added to the data. It appears inconsistent to have a continually changing estimate of excess return for the initial decade of data (1960s) based on the data that is being generated in the years most recently added (2013 in this case). An ideal methodology should be able to estimate excess returns using information available at that time, and provide estimates that do not change given future market movements assuming that the data used was not later corrected. It does not make intuitive sense that the excess return of this portfolio in 1967 is estimated differently in 1990 and 2012 using the same methodology. The performance of the portfolio between 1990 and 2012 could be argued to not shed any reliable light on the portfolio dynamics of 30 years prior without fairly restrictive assumptions being placed on the model. Such assumptions might be constant portfolio risk exposures and excess returns, which given the changes in trading environments over the 50 year interval appears unreasonable.

Asset level estimates do not suffer from these same limitations, they are far more flexible, and allow for changing portfolio dynamics. The primary challenge with asset level estimates is the computational time required to run regressions for every asset at every

time period in the model. In a simple restricted market with 1,000 assets over a 50 year window that is 600,000 regressions. At the time of the initial publications in momentum (1993), this may have been infeasible. However with the computing power available today these estimates can be obtained within reasonable timeframes, so this should no longer be considered a limitation.

Utilizing portfolio level regression, each regression has access to a large sample of data, hence there is no noticeable advantage to using more granular data than monthly. Asset level estimates are attempting to measure asset level parameters while allowing them to change over time. As such, assuming such parameters change, there is incentive to use the shortest time period possible with which to estimate asset parameters. The shorter the time period used in the regression, the closer data can get to the investment month, and the more accurate estimates can be of time varying parameters. Regressions over shorter time horizons benefit from more granular data, which allow less investment time to pass in order to acquire sufficient data to run the regressions. For this reason we will consider regressions using daily return data in our asset level analysis.

Lo and MacKinlay (1990) showed that daily data is more susceptible to lagged reactions to market data, so we will need lagged estimates of our asset level parameters. The theory proposed by Dimson (1979) has been previously implemented by Lewellen and Nagel (2006) in the momentum literature, and will be utilized here.

Lagged market returns are included stretching back five time periods, with older market returns being given lesser weight, to calculate the relevant asset parameters.

$$R_{i,t} - R_{f,t} = \alpha_i + \beta_{i,1} (R_{m,t} - R_{f,t}) + \beta_{i,2} (R_{m,t-1} - R_{f,t-1}) + \beta_{i,3} [(R_{m,t-2} - R_{f,t-2} + R_{m,t-3} - R_{f,t-3} + R_{m,t-4} - R_{f,t-4})/3] + \varepsilon_{i,t} \quad (3)$$

These three estimates of Beta are combined to give a combined daily Beta.

$$\beta_i = \beta_{i,1} + \beta_{i,2} + \beta_{i,3}$$

(4)

The inclusion of further lagged market returns do not affect the results of this study, and are not included to reduce the data required to estimate model parameters.

II. Existing estimates and issues in the estimation of parameters

The challenge with asset level estimation methodology is in estimating the asset level factors. In a random sample of assets it would simply be possible to regress on an assets historical return to obtain estimates of an asset's risk factors. If it is believed that these factors change over time the factor estimates can be done over a short window of history, with 2-5 years being common. The challenge in our case is that the selection of assets is not random, and we are selecting assets based on the endogenous variable in our regressions, the return of the asset.

In order to demonstrate the bias in the asset level regressions to be significant we need to understand the size of such a bias, and demonstrate that it is a first order of magnitude problem. A simple look at the returns of the momentum portfolios is illustrative in this case and sheds light on the impact of such a bias. Table 1 shows the

returns of assets included in the top and bottom deciles both before, during and after the sorting and holding periods. As would be expected the shorter sorting periods obtain the most extreme expected holding period returns, with average returns of over 10% per month being observed for both the 3 and 6 month sorting window of the long portfolio. With the mean Beta and market return both being positive over the sample it is not surprising that the short portfolio saw smaller magnitude averages, but still both shorter sorting windows (3 and 6 month) observe mean returns below -6% per month.

Figure 1 shows the distribution of returns over the sorting period. Returns are generally fairly evenly distributed, with the first and last month of the sorting period offering slightly more extreme returns than sorting months in the middle of the sorting window. Spikes in the last sorting month could be driven by established phenomenon such as pre-announcement drift, with spikes at the beginning of a sorting period similarly being driven by post-announcement drift. While interesting as an aside this does not cause concern with respect to the momentum methodology, so is not investigated further in this paper.

Both Table 1 and Figure 1 illustrate the significant extreme returns observed in the momentum sorting periods. As there appears to be opportunity for bias to be imparted in regressions using the sorting period, and that bias appears to be linked to the magnitude of sorting period returns, these results provide motivation for further inquiry.

Given that the source of bias in the current model is the inclusion of the sorting period in the data, a natural first step would be to exclude that from our analysis.

III. Estimates skipping the sorting period

The rationale for the currently used sorting window for analyzing momentum returns is simply that of most recent data available. If we believe that parameters change over time (Alpha and/or Beta) then the closer to the investment period we can be the closer our estimates will be to the true investment period parameters. As the sorting period falls immediately before the investment period in momentum strategies, this placed the sorting period directly in the regression sample.

The simplest means of excluding the sorting period is to use data from immediately prior to sorting to run our regressions. As a first look we will consider this, using data from 12 months prior to the sorting month going backwards. This analysis uses 24 months of data for each regression which is sufficient to get relatively stable estimates of the regression parameters.

Table 3 shows the results of these regressions. Implied Alphas are calculated by using Beta parameters from the regressions and raw returns from the holding period. Using this methodology we can see that the estimated momentum excess returns now fall within reasonable bounds, and we have no reason to suspect bias as the sorting period is excluded from the regression sample.

The challenge now is that the distance between holding period and regression data is increasing. The distance between the beginning of the regression window and the end of the holding period is between 3 and 4 years. It is not unreasonable to expect an asset to maintain constant parameters over a short window, but in this case the constraint is more

severe. There are two potential forces at work here changing the parameters, one is the constant gradual shift of parameters over time. The other is the potential for a dramatic shift in parameters between holding and sorting period. First we will address the former.

Table 4 demonstrates the changing nature of the regression parameters over time. The estimates are calculated over the 6 individual years immediately preceding the holding period. This is to cover all years that might be used in a 5 year moving window estimate giving the researcher the option of skipping the sorting period. The regressions are done over a short window but the sheer volume of regressions within a period, averaged across 50 years of data lends stability to the results. The analysis looks specifically at Beta estimates as this is the variable used to calculate implied Alpha over the holding period, so is of most interest.

Looking at the Beta estimates we can see that for all 16 long portfolios the Beta estimates move from an average of 1.21 with a 5 year lag, up to 1.38 immediately prior to the investment period. The short portfolio also shows such a change, moving from 1.34 to 1.48 as an average of all 16 portfolios. The results show that even over a short window regression there can be noticeable changes in Beta which would have a significant impact on estimates of Alpha. Given an average market risk premium of 0.45% per month from 1960 to 2012 (data from Ken French's website), a 0.17 shift in Beta is an 8 basis point shift in the estimates of Alpha generated from the strategy. In general this result would encourage us to use as recent data as possible to measure Beta, in order that it is as similar to the Beta found in the holding period as possible.

Aside from the generally smooth increase in Beta, the long portfolio shows a dramatic jump for the period including the sorting months, to the period immediately before (the analysis without lag). While other periods differ by less than 2%, the final step sees estimates of Beta change by 10% over and above estimates taken 12 months prior. This combined with earlier results that demonstrate bias induced into results utilizing the sorting period does not provide reassurance that the most accurate results can be gleaned from using the most recent data (including the sorting period).

With both angles considered the challenges from using the data can now be seen to apparently limit our ability to accurately model momentum using historical monthly data. If we use the most recent data available we are subject to bias caused by the inclusion of the sorting period in our regression sample period, and if we lag results further behind, we risk underestimating Beta due to the apparent upward drift of Beta as we approach the holding period.

One of the limiting factors in this analysis is the necessity for large numbers of observations which provide stability to the regressions performed. Using monthly data it has been common in the literature to require at least 3-5 years of data in order to generate stable estimates, and this exposes the results to ignoring short horizon changes in parameters. A natural way to reduce our dependence on these longer samples would be to move to shorter time horizons and more granular data.

IV. Proposed methodology

In order to reduce the dependence on long sample periods a natural progression is the change from monthly to daily return data. This change allows the use of samples that cover shorter time horizons, which use data from ever closer to the holding period. In theory this would reduce issues with time varying parameters. If it is possible to get a reliable estimate from a short sample an estimate from within the holding period would be ideal in order that this issue of time varying parameters could be eliminated completely.

A major challenge inherent in using daily returns is the greater difficulty in measuring Beta. Lo and Mackinlay (1990) identified delayed reactions to market movements when working at shorter time intervals, such as when using daily data. Their work found that the smaller assets tended to be impacted disproportionately highly with delayed reactions to market movements. As a means of accounting for such delays we use a lagged estimate of Beta in our daily regressions, as suggested by Dimson (1979). The Beta estimate is therefore a composite of the lagged Beta of the last five market daily returns, with more distant returns given diminishing weight.

$$R_{i,t} - R_{f,t} = \alpha_i + \beta_{i,0}R_{m,t} + \beta_{i,1}R_{m,t-1} + \beta_{i,2}[(R_{m,t-2} + R_{m,t-3} + R_{m,t-4})/3]$$
(5)

The final three lagged returns are combined to reduce the number of parameters estimated by the model. The overall Beta estimated by this model is taken to be the sum of all three Beta terms. The Alpha is simply taken directly, with no need to account for the lags. As this Alpha is a daily estimate we follow Lewellen and Nagel (2006) and multiply by 21 to obtain a monthly estimate.

As the daily data requires different treatment from the monthly data, it is important to establish that any results obtained are being driven by changes in the data, not changes in the methodology. To allay fears that this may be the case we will compare monthly and daily regressions. Table 5 contains 12 month regressions done using data taken backwards from 12 months before the sorting month (to avoid sorting period bias) for daily and monthly data. As can be seen from the results the estimates are broadly consistent, with the estimates individually differing from each other by an average of 4 basis points for both the long and short portfolio. This result gives indication that we can use the daily methodology with confidence knowing that results are not being driven by the change in methodology.

Given the daily asset level regressions are consistent with the results from monthly asset level regressions for similar time periods a natural progression is to ask, how much data is required to reasonably run the daily regressions. Ideally we wish to use as short an interval as possible to minimize the effect of time varying parameters on our results. The risk of using shorter horizons is that it could potentially lead to unstable estimates. Fortunately this style of analysis compensates for such instability at the asset level regression level by aggregating results sufficiently that stability is obtained with the aggregated estimates.

Table 6 shows the results of comparing the aggregated results of 1 month daily regressions to those of 12 month daily regressions. Both regressions use data going backward from 12 months prior to the sorting month. The table shows remarkable similarity of stability of estimates, and overall average estimates differing by

approximately 1 basis point. This is evidence to suggest that due to the methodology employed the results retain their stability even if limited to one month samples. This stability can be attributed to the volume of asset level regressions in each portfolio, the number of portfolios held at each holding period, and the number of holding periods, all contributing power in allowing the estimates to stabilize.

Given the result that daily regressions can be performed on only 1 month of data, it seems natural to attempt to perform the asset level regressions in, rather than prior to, the holding period. This eliminates the problem caused by the sorting period biasing results, and also completely eliminates the issues associated with time varying parameters, as the regressions occur as the asset is providing returns. Each asset within the holding period will obtain an estimate of Alpha directly, and there is no need to use the Beta estimates provided.

Table 7 provides the results of asset level regressions using daily data performed over the holding period months. For an asset which remains in the holding period for 6 months 6 separate regressions will be run, one each month to estimate a monthly Alpha. Those Alphas will then be aggregated by portfolio and aggregated over the sample period.

The results show Alphas generated from the momentum portfolio to be higher than anticipated in accordance with prior estimates. Over the 16 strategies tested the average momentum returns are estimated at 1.48% per month. This change implies that the estimates of momentum Alpha using historical data were failing to adequately measure the asset parameters as they changed through the period, and as such underestimated the true Alpha.

The manner in which this has been tested to this point demonstrates issues with and a solution to asset level estimates of momentum Alpha. However, the argument could be made that portfolio level estimates are not susceptible to such issues. While portfolio level estimates do have their own intrinsic issues which must be accepted, such as the restriction of a single constant exposure to market risk over the entire data range, they do not explicitly include sorting data in their regressions. Without this data being included there would not appear to be a bias issue.

In an efficient market assets included in momentum portfolios would be equally likely to be included or not in the momentum portfolio given prior inclusion (assuming non-overlapping sorting periods for simplicity). If we depart from efficient markets and assume momentum then we might expect assets that were in the long (short) portfolio to do better (worse) than expected under efficient markets, so be more likely to end up in the long (short) portfolio again in the next period (where the prior holding period would become the sorting period). We would therefore expect to see portfolio churn, the proportion of portfolio assets turned over in the course of a year, to be lower than expected under efficient markets.

Table 8 shows the results for churn levels observed when running a momentum portfolio versus the churn results that we expect in theory under efficient markets. Counter to intuition, assuming the existence of momentum, the portfolios show consistently higher churn than theory suggests, with all portfolios seeing 4%-20% higher than expected churn.

This high churn presents problems for portfolio based estimates. We know from asset based estimates that sorting period returns have higher (lower) than expected epsilon estimates in the long (short) portfolios. Given that the market epsilon has mean zero by construction we therefore know that the remainder of the market has epsilons biased in the opposite direction.

Given that we know due to churn that these portfolios contain less than expected of an assets "sorting period" returns, or holding period returns from this asset in a prior period, we know that these results are being under sampled. Under sampling positive (negative) epsilon returns in the long (short) portfolio will lead to the portfolio estimates of Alpha being downwardly (upwardly) biased. This implies that the overall estimate of momentum returns calculated at the portfolio level is biased, and further is downwardly biased on both the long and short portfolio.

Given the evidence presented above that suggests that the one month regressions using daily data it is of interest to compare to more traditional methodologies. Table 9 compares a standard asset level test (using data prior to the holding month of interest to calculate parameters) to the asset level tests using only the daily data from the holding month of interest. The results, consistent with the biases identified earlier, show that the newer method provides higher estimates of risk adjusted return. As an example, the 6 month sorting 6 month holding period momentum strategy shows a risk adjusted return increase of 40 basis points per month, from 1.24% to 1.64%. This indicates that the bias in the original estimates is of significant magnitude.

V. Simulations

Table 1 and Figure 1 clearly show the extremely high returns observed in the sorting period. As this is the claimed source of bias it is of interest to investigate how much bias this magnitude of return could produce in the momentum regression output. It is intuitively clear that a larger return in the sorting period should be the result of larger sorting period epsilon, and hence larger bias. However the question of whether this bias is economically significant is still unanswered.

For the purposes of estimating this bias we will utilize simulations as they allow the precise measurement of bias and an explicit statement of assumptions to simplify the interpretation of results. By simulating market data based on parameters driven by our raw data, the true and estimated model parameters can be easily measured, which makes the identification of bias relatively straightforward.

For these simulations we use the basic CAPM model to generate asset returns. This in turn requires us to generate market returns, Alphas and Betas for each asset. Market returns are generated as simple IID normal random variables drawn for as many consecutive months as necessary. For each iteration of the simulation these observations are refreshed.

The second step requires the generation of the epsilons, or random error terms, that occur for each asset at each time period.

To simplify the interpretation of the results all assets were given an Alpha parameter of 1%, and a Beta parameter of 1. Allowing variability of Alpha and Beta is possible, and

the results are not significantly different from those with static parameters, but they are more difficult to interpret simply (varying Alpha and Beta adds complexity without adding information). Also results in the momentum literature suggest that Alpha and Beta are not independent, and as such should not be generated as independent random variables. Exploring the functional form of the relationship between Alpha and Beta is considered beyond the scope of this paper, hence the use of static parameter.

Given the Alpha, Beta, epsilon and market excess return it is then possible to generate 30 months of excess return data for 500 assets using the standard CAPM structure below.

$$R_{i,t} - R_{f,t} = \alpha + \beta (R_{m,t} - R_{f,t}) + \varepsilon_{i,t} \quad (6)$$

Using this we now have 500 assets with simulated returns. Assets are then ranked and deciled on their last 6 months of excess returns. The portfolios of most interest for the momentum strategy are the highest and lowest decile, in this case each containing 50 assets.

Separate from the deciling each asset has a regression performed on the returns available. For the purposes of this simulation the regressions are run on 24 observations, equivalent to two years of data. These 24 months of data are the most recent 24 months so include the period used for sorting prior to deciling. The regressions take the market excess return and asset raw return as given, and estimate Alpha and Beta parameters. By construction the regression will set the sum of epsilons equal to zero. A test case was also

measured, with asset level regressions being run on the returns prior to the sorting period to provide a benchmark to judge results covering the sorting period.

This process (generating the returns on the market and assets, deciling, regressing, and combining regression results by decile) was performed 100 times. After 100 iterations the estimates become sufficiently stable that further iterations were deemed unnecessary.

Figure 2 shows the results of the simulations. Regressions over the pre-sorting period data indicate that the regressions methodology works well on this sample, with all deciles having Alpha estimates within 0.02% of the true background Alpha (1%). There is no discernible trend in the results by decile.

However, the regressions using the most recent data available (covering the sorting period) show a remarkable trend, with the lowest decile assets showing an estimated Alpha of -1.89%, a downward bias of 289 basis points. On the opposite end of the deciling the highest decile shows estimated Alphas of 3.86%, this time an upward bias of 286 basis points. The top and bottom decile are of most interest to us given that they are the targeted investment portfolios of the momentum strategy, but it is worth noting that the relationship is consistent across all 10 deciles, with observations further from the median exhibiting stronger bias.

Clearly this bias is an issue to be concerned with when considering momentum portfolio excess returns. The results imply that the bias that is exhibited when regressing over the momentum sorting period overstates the excess returns of the high decile assets, and understates the excess returns of the low decile assets. Both the bias on the long and

short portfolio will therefore contribute to an upward bias of the overall momentum excess returns.

In order to check the output of the simulations the simplest method is simply to compare the results of the simulation to the regression results observed using real data. Table 2 illustrates the results obtained using a 24 month regression over the pre-holding period data which includes the sorting period. The results are extremely consistent over different holding periods due to the fact that this estimate is static and obtained pre-holding period. Any changes in the Alphas observed can be attributed to assets delisting from later data having been present in shorter holding periods.

The results show returns on the long and short portfolios that provide momentum returns drastically higher than observed in the literature. While a return of approximately 1% per month have been documented fairly consistently starting from Jegadeesh and Titman (1993), these estimates show monthly returns for the 6 month holding period with 24 month regression of 3.8%. This result alone is reason to suspect bias in the results, and further inspection provides evidence to this effect.

The results from Table 2 show more extreme estimated Alphas for longer sorting periods. This is consistent with the idea that the sorting period is driving bias in the regressions. The longer the sorting period, the more extreme observations the methodology is likely to collect, and the higher the proportion of the regression window they will have become. This implies that bias would be greatest over the 12 month sorting periods, and least severe over the 3 month sorting periods. This is confirmed in Table 2.

Table 2 details regressions for 12, 24 and 36 months of observations. The observations are noticeably more extreme for the shortest window regressions, and least extreme for the longest window regression. This is consistent with the notion that the sorting period is causing bias in the estimates, as the shorter the regression becomes, the higher the proportion of data that falls within the sorting window, hence the exposure to the bias.

Traditionally these regressions have been the basis for estimates of momentum returns at the asset level. The magnitude of this bias is very high, so an obvious question is why this bias has not surfaced in the literature. Previously in the literature it has been common to use these regressions to estimate Beta (and discard the Alpha estimates) then use the estimates of Beta to compute implied Alpha for the holding periods (sometimes referred to as Alpha plus epsilon). The issue with these results is that the Beta estimates are being drawn from a clearly biased regression. If it could be established that Beta and Alpha are in no way correlated, then it could be argued that the estimates of Beta from this regression remain unbiased. Unfortunately it has been long established in the momentum literature, since at least Jegadeesh and Titman (1993), that there is a relationship between the variables. Figure 1 of the afore mentioned paper shows that in their analysis the long portfolio and the short portfolio both exhibited Betas significantly above average, and we are aware of the extreme Alphas that the portfolios generate, so it would appear unwise to assume no relation between the variables.

It might be thought of as interesting to use simulations to estimate the bias that could be caused by the relationship between Alpha and Beta. Unfortunately to do so would

require the assumption of a structural relationship between the variables, for which no single theoretical model is widely accepted. For the purposes of this paper it is sufficient to understand that the model has significant bias, and due to the correlation observed between the parameters being estimated, this bias is likely to impact both parameters, and any estimates of momentum returns that these models can in future provide for us. As such, it is prudent to look for a better model.

VI. Conclusions

The results presented have demonstrated that the use of most recent data to perform asset level regressions for momentum portfolios has the potential to induce serious bias in the data. This bias is induced by the method of identifying the momentum portfolio, namely sorting on returns. As a first check regressions prior to the sorting period were considered as viable alternatives, but the identification of systematically time varying parameters made such regressions also susceptible to bias.

As a solution we consider working with daily data and performing regressions over the holding period itself. Daily regressions were shown to be consistent with monthly regressions over periods where both could be run, and also were shown to provide accurate and stable estimates of the portfolio parameters. Running the momentum portfolio with daily data asset level holding period monthly regressions allows us to observe what appears to be a more accurate estimate of the Alpha generated by momentum portfolios.

Further to this work, a problem which remains outstanding is how to allow for the effect of delisting in the sample. Clearly when identifying high Beta assets with extreme returns (as the momentum strategy does) the investor is likely to see a significant number of delisting events.

Work by Shumway (1997) and Shumway and Warther (1999) have demonstrated that CRSP data does have some limitations when it comes to working with delisted assets. One side of the issue is that of measuring the final prices as an asset approaches delisting. For many such assets trading wanes, and bid ask spread rises, making the calculation of a simple return somewhat problematic, and making the returns provided by CRSP potentially unreliable.

Beyond such data, which can at least hopefully be corrected for, we have the issue of prices following delisting. If a delisting was unknown in advance then the investor does not have time to exit their position prior to the removal of the asset from the market. If the investor still holds the asset after delisting then the calculation of returns becomes yet more difficult. Many delisted assets go on to trade on the pink sheets, which is generally associated with low liquidity trading. Shumway (1997) was able to identify a significant portion of assets in this manner, but some were not accounted for. Of the assets found on the pink sheets it was observed in the paper that the bid ask spreads were frequently higher than 50% of the midpoint of the bid ask spread. Equally, finding a bid price in the pink sheets is not equivalent to observing a bid price with the NYSE. Limited liquidity on the pink sheets implies that even assuming you are able to find the asset

trading, and obtain a price, there is no guarantee that this will allow the liquidation of a large portfolio.

These problems with delisting returns are problematic when considering momentum returns, and could provide fruitful areas for further research.

References

- Ang, Andrew, and Joseph Chen, 2002, Asymmetric correlations of equity portfolios, *Journal of Financial Economics* 63, 443-494.
- Asness, Clifford S., Tobias J. Moskowitz, and Lasse H. Pedersen, 2013, Value and momentum everywhere, *Journal of Finance* 68, 929-985.
- Barberis, Nicholas, Andrei Shleifer, and Robert Vishny, 1998, A model of investor sentiment, *Journal of Financial Economics* 49, 307-343.
- Chordia, Tarun, and Lakshmanan Shivakumar, 2002, Momentum, business cycle, and time-varying expected returns, *The Journal of Finance* 57, 985-1019.
- Cooper, Michael J., Roberto C. Gutierrez, and Allaudeen Hameed, 2004, Market states and momentum, *Journal of Finance* 59, 1345-1365.
- Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam, 1998, Investor psychology and security market under- and overreactions, *Journal of Finance* 53, 1839-1885.
- Dimson, Elroy, 1979, Risk measurement when shares are subject to infrequent trading, *Journal of Financial Economics* 7, 197-226.
- Fama, Eugene F., and Kenneth R. French, 1996, Multifactor explanations of asset pricing anomalies, *The Journal of Finance* 51, 55-84.
- Fama, Eugene F., and Kenneth R. French, 1995, Size and book-to-market factors in earnings and returns, *Journal of Finance* 50, 131.
- Fama, Eugene F., and Kenneth R. French, 2012, Size, value, and momentum in international stock returns, *Journal of Financial Economics* 105, 457-472.
- Grundy, Bruce D., and J. S. Martin, 2001, Understanding the nature of the risks and the source of the rewards to momentum investing, *Review of Financial Studies* 14, 29-50.
- Harvey, Campbell R., 1989, Time-varying conditional covariances in tests of asset pricing models, *Journal of Financial Economics* 24, 289-317.
- Hong, Harrison, and Jeremy C. Stein, 1999, A unified theory of underreaction, momentum trading, and overreaction in asset markets, *Journal of Finance* 54, 2143-2184.
- Jagannathan, Ravi, and Zhenyu Wang, 1996, The conditional CAPM and the cross-section of expected returns, *Journal of Finance* 51, 3-53.
- Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: implications for stock market efficiency, *Journal of Finance* XLVIII, 65.
- Jegadeesh, Narasimhan, and Sheridan Titman, 2001, Profitability of momentum strategies: an evaluation of alternative explanations, *Journal of Finance* 56, 699-720.
- Jegadeesh, Narasimhan, 1990, Evidence of predictable behavior of security returns, *Journal of Finance* 45, 881-898.
- Lesmond, David A., Michael J. Schill, and Chunsheng Zhou, 2004, The illusory nature of momentum profits, *Journal of Financial Economics* 71, 349-32.
- Lettau, Martin, and Sydney Ludvigson, 2001, Resurrecting the (C)CAPM: A cross-sectional test when risk premia are time-varying, *Journal of Political Economy* 109, 1238-1287.
- Lewellen, Jonathan, and Stefan Nagel, 2006, The conditional CAPM does not explain asset-pricing anomalies, *Journal of Financial Economics* 82, 289-314.

- Lo, AW, and AC MacKinlay, 1990, When are contrarian profits due to stock market overreaction? *Review of Financial Studies* 3, 175-205.
- Ľuboš Pástor, and Robert F. Stambaugh, 2003, Liquidity risk and expected stock returns, *Journal of Political Economy* 111, 642-685.
- Lustig, Hanno N., and Stijn G. Van Nieuwerburgh, 2005, Housing collateral, consumption insurance, and risk premia: an empirical perspective, *Journal of Finance* 60, 1167-1219.
- Moskowitz, Tobias J., and Mark Grinblatt, 1999, Do industries explain momentum? *Journal of Finance* 54, 1249-1290.
- Moskowitz, Tobias J., Yao H. Ooi, and Lasse H. Pedersen, 2012, Time series momentum, *Journal of Financial Economics* 104, 228-250.
- Petkova, Ralitsa, and Lu Zhang, 2005, Is value riskier than growth? *Journal of Financial Economics* 78, 187-202.
- Sadka, Ronnie, 2006, Momentum and post-earnings-announcement drift anomalies: The role of liquidity risk, *Journal of Financial Economics* 80, 309-349.
- Santos, Tano, and Pietro Veronesi, 2006, Labor income and predictable stock returns, *Review of Financial Studies* 19, 1-44.
- Shanken, Jay, 1990, Intertemporal asset pricing: an empirical investigation, *Journal of Econometrics* 45, 99-120.
- Shumway, Tyler, and Vincent A. Warther, 1999, The delisting bias in CRSP's NASDAQ data and its implications for the size effect, *Journal of Finance* LIV, 2361.
- Shumway, Tyler, 1997, The delisting bias in CRSP data, *Journal of Finance* LII, 327.
- Wang, Jun, and Yangru Wu, 2011, Risk adjustment and momentum sources, *Journal of Banking and Finance* 35, 1427.
- Wang, Kevin Q., 2003, Asset pricing with conditioning information: a new test, *Journal of Finance* 58, 161-196.

Table 1**Returns of momentum portfolios during and after sorting**

The table below shows the returns of before during and after being selected as part of the momentum portfolio. Assets have their cumulative sorting period returns ranked, and the top decile becomes the long portfolio and the bottom decile becomes the short portfolio. The assets included in this table are all those on the NYSE, NASDAQ and AMEX exchanges between 1960 and 2012.

Sorting Period Months	Portfolio	Returns	
		Sorting Period	6 months post sorting period
3	Long	14.84%	1.40%
	Short	-9.93%	0.57%
6	Long	10.71%	1.62%
	Short	-6.72%	0.44%
9	Long	8.97%	1.66%
	Short	-5.28%	0.34%
12	Long	7.95%	1.56%
	Short	-4.40%	0.35%

Table 2**Asset level regressions including the sorting period**

The data below shows regression results from the long and short momentum portfolio using data from 1960-2012. All assets with 24 months of data are deciled, with the top and bottom decile being utilized for investment. Each asset in these portfolios of interest then has Alpha and Beta estimated over the most recent data pre-holding period, which includes the sorting period. The regression is completed on 24 months of data, therefore always includes data before the sorting period. Results across holding periods show extreme consistency due to minimal asset die off, and a static estimate of Alpha from a pre-holding period regression.

Regression	Sorting	Holding							
Months		Long				Short			
		3	6	9	12	3	6	9	12
12	3	4.34%	4.34%	4.33%	4.33%	-1.70%	-1.69%	-1.69%	-1.68%
	6	5.43%	5.43%	5.43%	5.42%	-3.07%	-3.06%	-3.05%	-3.05%
	9	6.12%	6.12%	6.11%	6.11%	-4.17%	-4.16%	-4.15%	-4.14%
	12	6.65%	6.65%	6.64%	6.63%	-5.15%	-5.14%	-5.13%	-5.12%
24	3	2.51%	2.51%	2.51%	2.50%	-0.13%	-0.12%	-0.12%	-0.12%
	6	3.07%	3.07%	3.07%	3.07%	-0.73%	-0.72%	-0.72%	-0.72%
	9	3.49%	3.49%	3.49%	3.49%	-1.25%	-1.25%	-1.24%	-1.24%
	12	3.82%	3.83%	3.82%	3.82%	-1.74%	-1.74%	-1.73%	-1.72%
36	3	1.85%	1.85%	1.85%	1.85%	0.33%	0.33%	0.33%	0.33%
	6	2.20%	2.20%	2.19%	2.19%	0.00%	0.00%	0.00%	0.00%
	9	2.47%	2.47%	2.46%	2.46%	-0.29%	-0.28%	-0.28%	-0.28%
	12	2.69%	2.69%	2.69%	2.69%	-0.57%	-0.57%	-0.56%	-0.56%

Table 3**Asset level regressions excluding the sorting period**

The data below shows regression results from the long and short momentum portfolio using data from 1960-2012. All assets with 24 months of data are deciled, with the top and bottom decile being utilized for investment. Each asset in these portfolios of interest then has Alpha and Beta estimated over the 24 months of most recent data prior to the sorting month – 12 months. None of these regressions include the sorting period in their sample.

Months	Long				Short				Long - Short			
	3	6	9	12	3	6	9	12	3	6	9	12
	3	0.42%	0.45%	0.46%	0.45%	-0.14%	-0.36%	-0.35%	-0.35%	0.56%	0.81%	0.81%
6	0.65%	0.65%	0.61%	0.53%	-0.44%	-0.54%	-0.55%	-0.45%	1.09%	1.19%	1.16%	0.98%
9	0.76%	0.70%	0.59%	0.49%	-0.52%	-0.64%	-0.56%	-0.41%	1.27%	1.33%	1.15%	0.90%
12	0.72%	0.61%	0.50%	0.43%	-0.64%	-0.63%	-0.48%	-0.33%	1.36%	1.24%	0.98%	0.76%

Table 4**Changing beta estimates of momentum portfolios over time**

The table below displays Beta estimates of regressions done on assets within the long and short portfolios of the momentum portfolio. Given the sorting month t , regressions are done on data from month $(t - a * 12)$ to $(t - a * 12 - 11)$ inclusive. Each regression is run on 12 months of data, and the intervals are non-overlapping. The estimates of Beta can be seen to be increasing as the regressions approach the holding period.

Lagged Years	Sorting	Holding				Short			
		Long				Short			
		3	6	9	12	3	6	9	12
0	3	1.32	1.32	1.32	1.32	1.46	1.46	1.46	1.46
	6	1.34	1.35	1.35	1.35	1.50	1.50	1.50	1.50
	9	1.39	1.40	1.40	1.40	1.49	1.49	1.49	1.49
	12	1.45	1.45	1.46	1.46	1.47	1.47	1.47	1.47
1	3	1.27	1.27	1.27	1.27	1.41	1.41	1.41	1.41
	6	1.25	1.26	1.26	1.26	1.46	1.46	1.46	1.46
	9	1.25	1.25	1.26	1.26	1.50	1.50	1.50	1.50
	12	1.25	1.25	1.25	1.26	1.51	1.51	1.51	1.51
2	3	1.25	1.25	1.26	1.26	1.35	1.35	1.35	1.35
	6	1.25	1.25	1.25	1.25	1.37	1.37	1.37	1.37
	9	1.24	1.24	1.24	1.24	1.41	1.41	1.41	1.41
	12	1.23	1.23	1.24	1.24	1.44	1.44	1.44	1.44
3	3	1.24	1.24	1.24	1.24	1.33	1.33	1.33	1.33
	6	1.24	1.24	1.24	1.24	1.36	1.36	1.36	1.35
	9	1.24	1.24	1.24	1.24	1.36	1.36	1.36	1.36
	12	1.24	1.24	1.24	1.24	1.38	1.38	1.38	1.38
4	3	1.22	1.22	1.22	1.22	1.30	1.30	1.30	1.30
	6	1.21	1.21	1.22	1.22	1.34	1.34	1.34	1.34
	9	1.21	1.21	1.21	1.21	1.35	1.35	1.35	1.35
	12	1.20	1.20	1.21	1.21	1.37	1.37	1.37	1.37
5	3	1.21	1.21	1.21	1.21	1.30	1.30	1.30	1.30
	6	1.21	1.22	1.22	1.22	1.30	1.30	1.30	1.30
	9	1.21	1.21	1.22	1.22	1.31	1.31	1.31	1.31
	12	1.21	1.21	1.21	1.21	1.32	1.32	1.32	1.32

Table 5**Monthly and daily data comparison for asset level tests**

The data below compares two methods of computing momentum Alphas, using monthly and daily return data. Both estimates use the 12 months of data starting 12 months prior to the sorting month (to avoid the sorting period for all strategies). The regressions are used to provide Beta estimates which in turn are used to compute implied Alphas for the holding period which are reported here.

Data	Sorting	Holding							
		Long				Short			
		3	6	9	12	3	6	9	12
Daily	3	0.43%	0.47%	0.49%	0.48%	-0.21%	-0.41%	-0.39%	-0.40%
	6	0.69%	0.70%	0.65%	0.56%	-0.50%	-0.58%	-0.59%	-0.49%
	9	0.80%	0.75%	0.63%	0.52%	-0.59%	-0.69%	-0.61%	-0.45%
	12	0.77%	0.67%	0.55%	0.45%	-0.72%	-0.68%	-0.53%	-0.37%
Monthly	3	0.42%	0.44%	0.44%	0.43%	-0.22%	-0.45%	-0.44%	-0.44%
	6	0.66%	0.65%	0.60%	0.51%	-0.53%	-0.64%	-0.65%	-0.54%
	9	0.76%	0.70%	0.57%	0.46%	-0.62%	-0.75%	-0.67%	-0.50%
	12	0.71%	0.59%	0.48%	0.38%	-0.71%	-0.71%	-0.56%	-0.40%

Table 6**Momentum alpha comparing 12 month and 1 month regressions**

The table below shows momentum Alpha calculated using daily data. The estimates are calculated using 12 and 1 month of data respectively going backwards from the sorting month – 12. Each regression is computed at the asset level and the Beta obtained is used to calculate the implied Alpha over the holding period.

Data	Sorting	Holdin g							
		Long				Short			
		3	6	9	12	3	6	9	12
1 Month	3	0.45%	0.50%	0.47%	0.45%	-0.20%	-0.39%	-0.40%	-0.43%
	6	0.70%	0.74%	0.65%	0.55%	-0.42%	-0.52%	-0.57%	-0.49%
	9	0.82%	0.80%	0.64%	0.51%	-0.52%	-0.66%	-0.59%	-0.47%
	12	0.81%	0.72%	0.55%	0.43%	-0.67%	-0.66%	-0.54%	-0.41%
12 Month	3	0.43%	0.47%	0.49%	0.48%	-0.21%	-0.41%	-0.39%	-0.40%
	6	0.69%	0.70%	0.65%	0.56%	-0.50%	-0.58%	-0.59%	-0.49%
	9	0.80%	0.75%	0.63%	0.52%	-0.59%	-0.69%	-0.61%	-0.45%
	12	0.77%	0.67%	0.55%	0.45%	-0.72%	-0.68%	-0.53%	-0.37%

Table 7**Regressions on holding period return using 1 month**

The table below shows risk adjusted momentum returns from regressions on daily data. The regressions use 1 month of observations (approximately 20 observations each) to estimate the Alpha and Beta for an individual asset in the holding period. As the regressions are being done over the holding period itself Alphas can be taken directly from the regressions unlike in traditional asset level tests.

Sorting	Holding							
	Long				Short			
	3	6	9	12	3	6	9	12
3	1.10%	1.06%	1.02%	0.99%	-0.10%	-0.21%	-0.16%	-0.16%
6	1.35%	1.27%	1.19%	1.06%	-0.33%	-0.37%	-0.37%	-0.25%
9	1.43%	1.31%	1.15%	1.00%	-0.45%	-0.51%	-0.39%	-0.22%
12	1.41%	1.22%	1.06%	0.93%	-0.60%	-0.50%	-0.30%	-0.13%

Table 8**Churn rates for momentum portfolios**

Churn rates below are calculated as the total volume of assets that the portfolio is required to buy and sell as a percentage of total portfolio dollar value. In this way a portfolio of \$100m which sold \$5m assets per month and replaces them with other assets of equivalent value would be $5 * 12 / 100 = 60\%$. The theoretical churn rates are calculated assuming efficient markets. As an example, a 3 month holding period necessitates that 1/3 of the portfolio is sold every month, but by random chance we would expect a 10% overlap between portfolios we buy and sell. This means a 30% churn is observed monthly, which is equivalent to 360% per year.

Sorting	Holding	Portfolio	Observed Churn	Theoretical Churn	Difference
3	3	Long	374%	360%	14%
6	6	Long	186%	180%	6%
9	9	Long	124%	120%	4%
12	12	Long	94%	90%	4%
3	3	Short	380%	360%	20%
6	6	Short	191%	180%	11%
9	9	Short	128%	120%	8%
12	12	Short	97%	90%	7%

Table 9**Comparing new to existing methodology for momentum**

This table provides a comparison of Alpha estimates provided by differing estimation methodologies for momentum portfolios. Both estimates use the same strategy, using data from 1960 to 2012 excluding assets priced under \$5, investing in the top decile and shorting the bottom decile based on sorting period performance. Both methods use daily data in their estimations. The holding period 1 month regressions runs regressions every month using the holding month for the asset level regression. The pre holding period uses the historical methodology of taking the 12 months of data pre holding period to estimate asset level parameters.

Data	Sorting	Holding			
		3	6	9	12
Holding Period 1 month	3	1.20%	1.27%	1.18%	1.15%
	6	1.68%	1.64%	1.56%	1.31%
	9	1.88%	1.83%	1.54%	1.22%
	12	2.02%	1.72%	1.36%	1.06%
Pre-holding Period 12 months	3	0.64%	0.86%	0.85%	0.85%
	6	1.15%	1.24%	1.22%	1.03%
	9	1.34%	1.40%	1.21%	0.94%
	12	1.47%	1.32%	1.06%	0.82%

Figure 1

Momentum returns before, during and after sorting periods

The charts below show the returns of assets selected in the momentum portfolios for 20 months before and after the final month of the sorting period. All assets in the NYSE, AMEX and NASDAQ exchanges are included in this analysis.

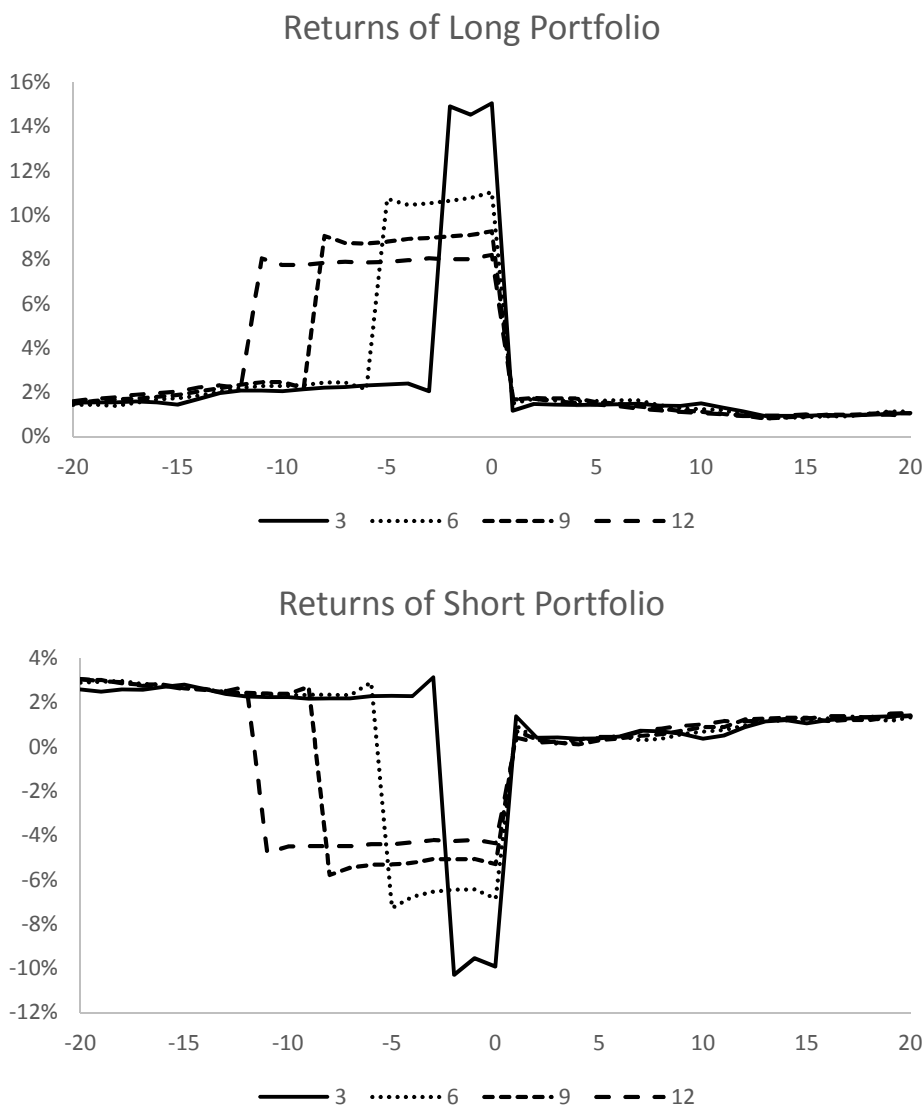
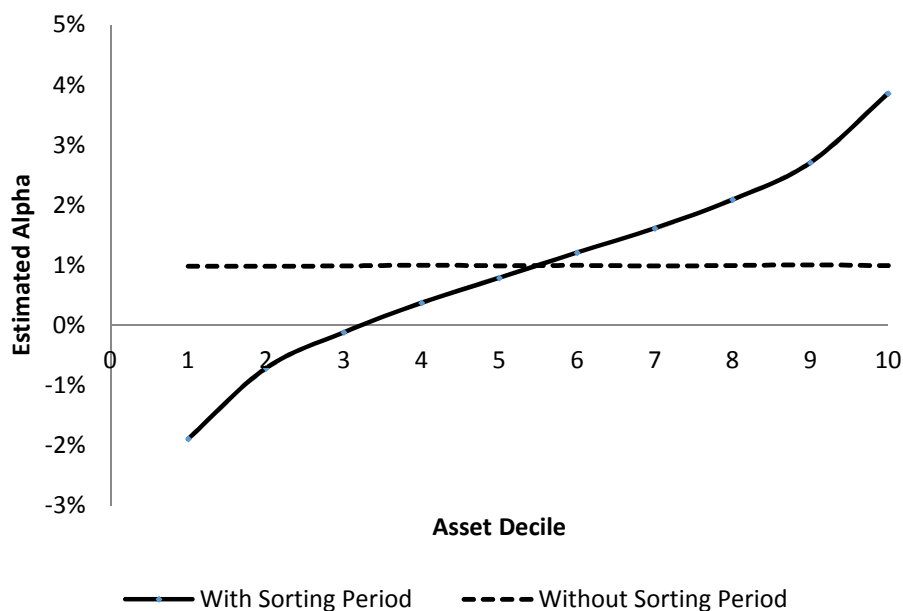


Figure 2

Simulating estimated alpha in momentum strategies

The chart below shows the output of regressions run on simulated asset level data. Simulations generate returns consistent with CAPM formulation for 500 assets. These 500 assets are then sorted into momentum deciles using 6 months of data. For each asset a regression is run using the most recent 24 months of data with and without the sorting period. This simulation is run 100 times at which point the results became sufficiently stable as seen in the chart below.



Chapter 2

Does Momentum Still Exist in US Equity Markets?

Abstract

We demonstrate that the momentum anomaly is driven by a small number of assets using asset level momentum estimates. We show that these assets behave differently in long and short portfolios, and also perform differently during the first month reversal period. We also show that some traditional facets of momentum portfolio construction may have in fact been subject to model selection bias. Despite these results an appropriately risk adjusted momentum return (alpha) remains positive and economically significant (unlike traditional estimates), even over the most recent data which has provided poor conditions for the momentum anomaly.

Momentum is the most studied, published and debated technical trading based market anomaly of the last 25 years. The notion that past winners will retain their superior performance in future and past loser firms will continue to underperform is so intuitively appealing and statistically well documented that a quarter of century of research has failed to disprove its existence. It remains one of the primary thorns in the side of even the most basic form of market efficiency.

There are several perspectives that can reasonably be taken, and have been argued, with respect to momentum. If the momentum anomaly is taken at face value it would appear that it offers substantial evidence against the existence of even weak form market efficiency. It is by far the most well publicized and accepted technical trading strategy in the academic asset pricing literature, and is considered, alongside the fundamental anomalies of size and value, as one of the more prominent unexplained risk factors of the last 25 years.

While initially posited as a challenge to just the US equities market by Jegadeesh and Titman (1993) there has since been a substantial amount of research that has provided evidence of the momentum anomaly outside of this market. A significant portion of the literature (Chan, Hamao and Lakonishok, 1991; Fama and French, 1998; Rouwenhorst, 1998; Griffin, Ji and Martin, 2003; Asness, Moskowitz and Pedersen, 2009; Chui, Titman, and Wei; 2010) has established the prevalence of momentum (with size and value) as an anomaly internationally, which lends weight to the idea that the anomaly exists and is persistent in the market.

An alternative view of momentum is that it is a spurious anomaly formed by the mis-measurement of an unspecified (in the original models) risk factor. The idea persists in the literature that a better, more complete understanding of the modelling of asset prices, and an appropriately specified set of risk factors should be sufficient to account for the momentum anomaly. There is literature on integrated international asset pricing, attempting to explain the asset pricing anomalies using international risk factors. Karolyi and Stulz (2003) provides a comprehensive overview, with notable examples being Griffin (2002), Hou, Karolyi and Kho (2011) and Fama French (2012). However the literature does not yet agree on an international model that explains significantly more than the country or region specific models.

Alternative risk factors have also been proposed by the literature. Pástor and Stambaugh (2003) and Sadka (2006) look at liquidity as a factor that attempts to explain away momentum. Asness, Moskowitz and Pedersen (2013) take this idea beyond the US equities market. Moskowitz and Grinblatt (1999) use industry factors, while Chordia and Lakshmanan (2002) and Cooper, Gutierrez and Hameed (2004) look at macroeconomic factors. Ultimately, although some of these papers claim to reduce the magnitude of the momentum anomaly by some degree, none of these efforts have come close to fully explaining away the momentum anomaly.

There is also a section of the literature that looks at measuring risk associated with momentum, but rather than look for new risk factors, it looks for methods of better measuring existing factors; namely by allowing time varying factor exposure. Extensive

coverage has been given by Jagannathan and Wang (1996), Lettau and Ludvigson (2001), Santos and Veronesi (2006) and Lustig and Van Nieuwerburgh (2005) to these time varying estimates, although primarily focusing on size and value estimates. More work directly related to momentum has been completed by Wang and Wu (2011) and Lewellen and Nagel (2009).

The third perspective is that momentum is an artefact of data mining. The assumption is that the finance literature tests many orders of magnitude more anomalies than are ever published, so by the nature of publishing only those that provide dramatic results, we have the potential for a significant model selection bias. The primary method used to combat such allegations is an out of sample test of the anomaly. There is always the challenge of arguing the validity of such tests if there has been a significant change in the data over the period. Improved methods of measurement leave the interested researcher balanced between accusations of failing to use the appropriate models, and adapting the models so far as to suggest overfitting the data.

This paper will consider both a risk based solution to the momentum anomaly, and also the argument that the anomaly is the result of overfitting through data mining. Unlike prior papers, we are not attempting to provide additional measures of different types of risk to explain the anomaly. Instead, we are focusing on a better method of evaluating market risk, and allowing a portfolio risk exposure to vary over time. With respect to accusations of overfitting the model through data mining, we will evaluate some of the changes to the model that have occurred over the history of the literature.

And in so doing, consider their impact on the strength of the anomaly over different time periods, including out of sample tests, to see if there is evidence to support such conclusions.

Overall, our paper finds that there is significant evidence to suggest a decline in the profitability of momentum over the last 50 years of return data. The number of assets within the momentum portfolio that generate alpha appears to be dwindling fairly consistently over time. We also show that there is severely limited data to suggest that momentum is as persistent as proposed in the literature. The standard momentum strategies suggest that momentum is valid for up to 12 months after the initial investment. However we find that the vast majority of momentum alpha is generated in the first half of the holding period, and this is obscured by using a combined portfolio.

We also find that there is evidence to suggest that some elements of the standard measurement of the momentum anomaly are open to accusations of overfitting the data, in particular we look at the decision to remove all assets under \$5. This was an addendum to the traditional momentum methodology, and appears to have inflated alphas which provides evidence to suggest overfitting and model selection bias.

However, despite both these results we find that in the most recent out of sample test, over the worst decade for momentum since the 1930s, and using the best methodology we have to date, momentum continues to deliver positive alpha.

Our work provides an extension to current papers looking at risk based explanations for the momentum anomaly and other anomalies. We show that using a novel

methodology we can obtain accurate and rapidly changing factor rotations, which account for the changing risk profile of the momentum portfolio. The output of these models suggest that momentum is alive and well in the most recent data, and also that traditional models fail to capture the full texture of the risk exposure of the investor pursuing momentum alpha.

Our paper is organized as follows. Section I summarizes the data and methodology used to evaluate the momentum anomaly, from the portfolio formation to the multiple methods of controlling for risk exposure. Section II deals with our primary results, detailing challenges with measuring momentum and looking at data that often undermines the momentum anomaly. Section IV concludes and provides thoughts on future research.

I. Data and Methodology

The data used in this study is obtained from the Center for Research in Security Prices (hereafter CRSP) as is typical in the asset pricing literature. Daily and monthly price data is obtained beginning in 1960 through to 2013. Earlier CRSP data is not used due to concerns with backfilling.

Supporting data is obtained through Ken French's website. This includes estimates of market returns, the risk free rate, and anomaly returns. For the purposes of this study the anomalies size, value and momentum are utilized. Monthly and daily data is used for the purposes of our results.

For the purposes of this study we are only concerned with American listed stocks, and have a preference for liquid and highly traded assets. For this reason we filter the data such that all assets had to be listed on the New York Stock Exchange (NYSE), American Stock Exchange (AMEX) or NASDAQ. CRSP share codes are required to be 10 or 11 which removes Real Estate Investment Trusts (REITs), American Depository Receipts (ADRs) and some other small asset classes from the data.

For purposes of liquidity we will generally exclude assets with a price of less than \$5 from our investment universe. This is intended as a protection against looking at assets with low liquidity, and follows the standard set by Jegadeesh and Titman (2001). In the rare instances where low price assets are included, we will identify the change explicitly.

For the purposes of this study we will focus primarily on the momentum anomaly. The methodology we use will be similar to that seen in Jegadeesh and Titman (1993), Asness (1994), Fama and French (1996), and Moskowitz and Grinblatt (1999). The premise of the momentum strategy is that assets that have done well over the last 3-12 months, and returned in the top segment of the market, will continue to do well over the next 3-12 months. Likewise assets that have done poorly will continue to do poorly. A portfolio long in assets that have done well, and short in assets that have done poorly, has been found to provide positive risk adjusted return.

The most common strategy is the 6 month sort, most popularized by Jegadeesh and Titman (1993). The premise is that at time $t=0$ we sort all assets on the past 6 months of raw return data (t from -1 to -6). We then invest in the top decile of assets, and short

the bottom decile of assets. The initial momentum strategy proposed by Jegadeesh and Titman (1993) uses these deciles, but some later papers such as Jegadeesh and Titman (2001) have changed to using the top 33% of the data. There has been little evidence of a material change to momentum returns historically by changing methods, so for the sake of this study we maintain the original method to reduce the influence of overfitting the most recent data. While we are sure the intentions of prior authors was good, we cannot account for the model selection bias that may be incurred by tweaking a methodology with each publication, hence we have a preference to stick to the original where possible.

The most popular investment horizon is 6 months from the initial papers, and so the investment will be held from time $t=0$ to time $t=5$. There is also a variant of the strategy again proposed by Jegadeesh and Titman (1993) that suggests skipping one month between sorting months and investing months to enable the investment to obtain better returns. While our base strategy for the purposes of this paper will not include the skipped month, we will revisit this variant later.

Using this strategy the investor can obtain a new portfolio every month. If the investor decides to hold each portfolio for 6 months, then each month the investor will be forced to liquidate $1/6$ of her portfolio, and invest that $1/6$ into the new portfolio calculated based on the last 6 months returns.

Given this portfolio it seems trivial that we can calculate portfolio returns. In the event of equal weighted portfolios, the investor simply averages the returns of all assets in the 6 currently held portfolios, then averages those 6 returns to create a monthly

return. This raw return can be simply calculated for the entire sample period and the investors long term return can then be computed.

With this portfolio level return, the literature has then traditionally calculated the risk adjusted return of the portfolio by regressing that return against the market return, and any other risk factors that the authors deem important. Popular models include (but are not limited to) the simple CAPM, and the Fama French three factor model popularized by Fama and French (1995, 1996). Models are shown below, where $R_{i,t}$ represents the portfolio “i” return at month “t”.

CAPM Model:

$$R_{i,t} - R_{f,t} = \alpha_i + \beta_i (R_{m,t} - R_{f,t}) + \varepsilon_{i,t} \quad (1)$$

Fama and French three factor model:

$$R_{i,t} - R_{f,t} = \alpha_i + \beta_{i,m} (R_{m,t} - R_{f,t}) + \beta_{i,smb} SMB_t + \beta_{i,hml} HML_t + \varepsilon_{i,t} \quad (2)$$

The results typically show that the overall return has a beta close to 0, and as such the portfolio risk adjusted return is not significantly different from the raw return. If anything the beta on the short portfolio typically slightly exceeds that on the long portfolio, which leads to an increase of a few basis points when comparing risk adjusted returns to raw returns.

The traditional method uses a single estimate of portfolio beta over the sample. With a reliable data window now stretching over 50 years this assumption of single constant beta is being stretched further and further as more data is added to our sample by CRSP. Additional pressure has been brought to traditional methods by research

suggesting that the momentum portfolios have poor performance months under periods of intense pressure on the overall market. When the market goes through a period of stress, momentum strategies can lose large sums. This indicates that the momentum strategy may not have constant risk exposure, which only further undermines the idea that the researcher or investor should be calculating a single point estimate of risk exposure over the entire sample.

To combat this, methods have been suggested that utilize the increasing computational power available to the researcher. Among the methods suggested have been moving window estimates of risk adjusted returns. This allows the use of monthly data to calculate portfolio level alpha over a 3-5 year window. Unfortunately with a 6 month sorting and investing window, it is possible for a momentum portfolio to completely reinvent itself over a 6 month window, making even these moving window estimates somewhat unreliable in periods of rapid and large magnitude market shocks.

A solution to this problem proposed by Wang and Wu (2011) and also Lewellen and Nagel (2006) is to perform asset level regressions using monthly data. By cycling asset level estimates into and out of the investment portfolio we can allow changes to portfolio representation. Unfortunately this also assumes constant asset risk exposure over the modeling window (3-5 years) and in the case of momentum, involves regressing over the sorting period. Regressing over the sorting period violates several assumptions, most notably that the sample is not chosen with bias in the parameters.

To resolve these problems the authors suggest in a prior paper (Gorman and Wu, 2015) that a better method is to use daily data to run asset level regressions. The proposed method runs one regression for each asset, each month. Each regression will have 20-23 observations, and will allow the estimation of risk adjusted return for an individual asset in an individual month. Using this we can then aggregate regression output to the portfolio level to get risk adjusted returns, that adjust immediately to changes in asset risk exposure, or portfolio composition. While each individual regression might be quite noisy, the aggregate of a large number of regressions (in the portfolio returns) has been shown to give consistent results with small margins of error. The power of this test is not coming from the size of an individual regression, but the number of regressions performed.

To regress data at the daily level we need to adjust our models to account for factors that are particular to this model. Most notable of these is the propensity of the asset level daily data to demonstrate lagged response to market return. Lo and MacKinlay (1990) showed that daily data is more susceptible to lagged reactions to market data, so we will need lagged estimates of our asset level parameters. The theory proposed by Dimson (1979) has been previously implemented by Lewellen and Nagel (2006) in the momentum literature, and will be utilized here.

Lagged market returns are included stretching back five time periods, with older market returns being given lesser weight, to calculate the relevant asset parameters.

$$R_{i,t} - R_{f,t} = \alpha_i + \beta_{i,1} (R_{m,t} - R_{f,t}) + \beta_{i,2} (R_{m,t-1} - R_{f,t-1}) + \beta_{i,3} [(R_{m,t-2} - R_{f,t-2} + R_{m,t-3} - R_{f,t-3} + R_{m,t-4} - R_{f,t-4})/3] + \varepsilon_{i,t} \quad (3)$$

These three estimates of Beta are combined to give a combined daily Beta.

$$\beta_i = \beta_{i,1} + \beta_{i,2} + \beta_{i,3} \quad (4)$$

The inclusion of further lagged market returns do not materially impact the results of this study, and are not included, to reduce the data required to estimate model parameters.

The existence of hundreds, if not thousands, of observations each month that have been estimated provides a new challenge to the researcher. While a traditional approach might be to aggregate these results in order to ascertain the most secure estimate available of portfolio level parameters, it is also possible to examine these results individually. Clearly there exists a multiple hypothesis testing challenge when it comes to working with such a large number of tests.

One option to challenge this would be the use of Bonferroni² corrections to our p-values. Essentially this simply requires that we adjust our p-values such that we can make the claim that the odds of even one observation exceeding the cutoff is less than the predetermined percentage (commonly 5%). However this is a somewhat crude and

² From the applied statistical literature there is coverage of Bonferroni's method in Schweder and Spjotvoll (1982) and Hochberg and Benjamini (1990). The usage in Finance is not widespread, but can be found in (among others) Shanken (1990), Ferson and Harvey (1999), Boudoukh et al. (2007) and Patton and Timmermann (2010). For a comparison with other methodology see Harvey and Liu (2015)

cautious approach, and does not tell us of the likelihood of other observations violating the null hypothesis.

For a more nuanced analysis of this problem we will use the methodology of false discovery rates³. Essentially, rather than attempting to classify observations as able to reject or not the null hypothesis, we look at the distribution of a test statistic produced and estimate the number of observations that would have been generated within a given sample under the null. The percentage of the sample expected to be produced by the null, divided by the total number of observations in the sample space, is then the false discovery rate. A high (low) false discovery rate in this instance indicates that an observation in this portion of the distribution is more (less) likely to fall under the null hypothesis.

This is particularly important when dealing with large samples of data. Momentum anomaly tests used most commonly in the literature consist of 50 years of data, each month being invested in 3 to 12 portfolios, and each portfolio consisting of hundreds of assets, it is clear that we have a large number of observations. As an example, let us consider a sample of 1 million observations. In this case, if we correctly identify the distribution of the null hypothesis, we would expect that of 1 million observations, 50,000 test statistics would exceed the threshold to be considered significant at the 5% level of a one tailed test under the null hypothesis. In this example, if on examination of the data

³ Benjamini and Hochberg (1995), Benjamini and Yekutieli (2001), Sarkar (2002) and Storey (2003) all contribute to the development of the false discovery rate literature. Applications in finance, see Barras, Scaillet and Wermers (2010), Bajgrowicz and Scaillet (2012) and Kosowski, Timmermann, White and Wermers (2006).

we find that 60,000 observations exceed the 5% significance cutoff, we could conclude conservatively that 10,000 observations violate the null hypothesis. We can also say that our estimate of a false positive rate from observations exceeding the cutoff to be significant is $50,000/60,000$. This rate, $5/6$, is the false discovery rate. It is the rate that predicts the probability of an observation over the significance cutoff being generated by the null, and thus being a false positive.

When looking at the distribution of risk adjusted returns, we can take a parametric approach and assume a distribution that the null follows, or use an empirical distribution to estimate the null hypothesis, as a non-parametric method. The parametric methods are useful, and often necessary in cases with small samples, samples of known distribution, and samples subject to two tailed testing, where non-parametric methods are more difficult to apply. As we have very large empirical distribution to work with it is reasonable to assume that the non-parametric approach has sufficient power to be reasonably accurate. As non-parametric methods require less assumptions to be made we will use these methods for our analysis.

In order to use a non-parametric estimate of the null hypothesis we are required to have such a distribution from which to estimate. When looking at a large number of risk adjusted return observations generated under the null hypothesis we could look for segments of the data that do not have documented risk adjusted returns. As an example, with a top decile of the market we might expect to find positive momentum alpha, and the lowest deciles may provide negative alpha, however we have no reason to believe that the central deciles provide alpha, so these deciles could provide our baseline.

Unfortunately, assets in central percentiles are more likely to be less volatile, and as such not representative of the null hypothesis in outer deciles. Using them as a baseline would be expected to underestimate the variance within the null hypothesis and as a result underestimate the False Discovery Rate, and over estimate the number of significant observations in the sample.

As an alternative we could look within the extreme deciles. By construction, the null hypothesis is mean zero when considering risk adjusted returns. Within the long portfolio, because we are looking at a one tailed test, there is no evidence that any observation with risk adjusted alpha less than zero could reject the null hypothesis. The lower half of the distribution can be considered to be entirely generated under the null, and used to generate the full null distribution. We do not have a parametric framework on which to hang the data, so we are forced to make one assumption, that the data is approximately symmetrically distributed around the mean of zero. Using this assumption, we can then mirror the empirical distribution below zero, on the unknown space above zero.

This mirroring can be used on the short portfolio also, although in this case we assume that the portion of the distribution above zero is generated by the null, and therefore we are estimating the lower portion of the distribution.

Once the distribution of the null hypothesis has been established the actual distribution can be compared to that of the null. Where the actual distribution exceeds the null hypothesis we can then claim evidence of observations that reject the null hypothesis. The false discovery rate is the proportion of data expected under the null, over the total observed data for that portion of the distribution.

Within this methodology there are more choices to be made in the method of implementation. For large statistically robust samples we can carefully analyze within the distribution, which percentiles exhibit non-null data, and where the false discovery rates are the highest. Alternatively, for smaller samples where the overall FDR for that side of the distribution matters, we can estimate simply by comparing using large intervals containing all observations on one side of the mean. In practice these methods give slightly different results, but the total distribution large interval estimate will provide a lower bound for the sum of the more granular estimates. In this way it can be considered the more conservative estimate of the two.

II. Results

Since at least the mid 1990s momentum has been a thorn in the side of arguments for market efficiency. Most frustratingly for proponents of efficiency the anomaly does not relate to any fundamental characteristic of the firm, and is entirely price driven, undermining even weak form market efficiency. Perhaps most challenging is that the momentum anomaly was obtained trading on the NYSE (later AMEX and NASDAQ), trading relatively large liquid assets, and large numbers of them. The momentum does not suggest that the inefficiency is small and contained. It suggests that momentum impacts all assets, and is extremely profitable when trading 20% of the market at any given time (top and bottom decile).

This seems to be a reasonable conclusion to draw. We invest in a portfolio of hundreds of assets, and make risk adjusted return, and as such the entire portfolio can be classified as flouting market efficiency. However, initial forays into momentum have not really been able to answer the question, are momentum returns driven by the entire portfolio exhibiting the anomaly, or a small number of assets exhibiting the anomaly very significantly. If there is evidence that a small number of assets drive the anomaly, then it is valuable to attempt to separate these assets and identify what distinguishes them from the remainder of the data. In so doing it is hoped that we can better understand the driving forces behind the momentum anomaly.

Alternatively we could think about how to come up with a better strategy. If a small number of assets drive the momentum anomaly, investors should be targeting those assets in our investment portfolios, at least to the point where we might believe that the anomaly has been traded out of existence. The remainder of the momentum strategy (the entire top decile) is not of interest if we can find a smaller subset that supplies all of (or the majority of) the risk adjusted return.

In our attempts to identify the specific assets from within the momentum portfolio that exhibit momentum, the first challenge in the literature is to identify the risk adjusted return of each asset each month over the entire sample of interest. The vast majority of work on the momentum anomaly works with returns at the portfolio level, and regresses upon the data at that level. This does not account for the difference in asset level risk

exposure, and makes it impossible to discern the risk adjusted performance of individual assets.

Improvements to this method were initially attempted using sliding window estimates of asset parameters. 3-5 years of data and regressions at the asset level could at least get to risk adjusted asset returns. Unfortunately this method also ran into difficulty, in that regressing over the momentum sorting period induced bias. Also, the assumption of constant market risk exposure over the sample is still problematic.

The solution to this is the asset level estimates of risk adjusted return using one month of data. Each asset has a new regression run every month using that month of asset return data, in order to estimate its risk exposure. The data is somewhat noisy on an individual asset-month basis, but when aggregated we can begin to understand how the data behaves and how it changes over time.

The following results are split into three sections. Section 1 covers using false discovery rates with momentum. Section 2 will cover the changes occurring in momentum results over time, and section 3 will look at whether there is evidence that momentum no longer exists in the most recent data.

Using False Discovery Rate methodology to analyze momentum

There is a clear challenge when testing thousands of test statistics simultaneously. To evaluate these thousands of monthly risk adjusted returns we turn to the statistical

method of False Discovery Rates. Through the use of this statistic the goal is not to identify whether an individual observation does or does not manage to reject the null hypothesis. The goal is to estimate the probability of an observation with a given alpha coming from the null hypothesis distribution. Given this, we can then calculate an estimate of the number of observations that are not generated by the null hypothesis.

The main decision a researcher has to face when implementing the FDR is how to estimate the null hypothesis. Standard parametric methods of full sample parameter estimation have been criticized for being wildly inaccurate when there are a significant number of observations generated from a distribution other than the null. To counter this it is common to choose only the data that you can be most certain fall under the null hypothesis in order to estimate your null distribution.

In a two tailed test, it would therefore be fair to use only a central cut of the data in order to estimate the mean and spread of the distribution. In our example, keeping long and short portfolios separate, we clearly have prior expectations as to the direction of the likely deviations from the null hypothesis. We would expect that the long portfolio has positive alpha deviations from the null, and that the short portfolio has negative alpha deviations from the null. We also do not need to estimate the mean of the null hypothesis, the null is that we have efficient markets, and the mean of that null hypothesis alpha is zero.

Let us first look at the long portfolio. We want to estimate the distribution of alpha under the null hypothesis, using as little data as possible that may fall under the

alternative hypothesis. And to do this given that we know that the mean of the null must be zero. The authors suggest that a simple way to do this is to fit the distribution on just the observations that are below zero alpha.

The next challenge is how to parametrize such a distribution. A common approach is to assume that the data is normally distributed, but this does impose additional assumptions. Generally this parametric approach is necessary when dealing with a small number of observations so fitting a standard distribution eliminates sample variability driving the results. In our case we are fortunate to have over 2.5 million observations on which to draw. With this number of observations the benefits of using a parametric distribution are less clear, and so the nonparametric methods become more appropriate tools for the task at hand.

Using all observations below alpha equal to zero, we then set a distribution. We know that the data represents the shape of the null hypothesis below zero, so mirroring that distribution above zero we have a reasonable estimate of the null distribution. Any instances of the observed frequencies on the positive side exceeding those seen on the low side would represent a frequency not fully explained by the null hypothesis. As with any method, we are still forced to make some assumptions with the use of this nonparametric test, but those assumptions are less arduous than in the case of a parametric test. For example, we are forced to assume a null distribution that is symmetrical. This is a less stringent restriction than fitting a parametric distribution to the data. A normal distribution for example, which also implicitly assumes the symmetry of

the null distribution, along with other assumptions that, while often reasonable, we would rather not make if we have alternative methods available.

The analysis done represents all assets in the long portfolio of the momentum investment strategy from 1963 to 2013, 50 years of data. For each month an asset is in the investment portfolio we run a regression on that month's daily data and get a single estimate of alpha. That one asset held for 1 month, and the associated alpha, represent one data point. The histogram contains over 1 million such observations. This sample size is what allows us to use a non-parametric estimate with relative precision.

Figure 1 shows the histogram of the long portfolio risk adjusted returns. As we can see, the center of the distribution is symmetrical, which lends support to our non-parametric null hypothesis. The black bars represent the null hypothesis, and the white bars represent the cases where the observations exceed the expected observations under the null hypothesis.

From Figure 1 we can see that the long portfolio shows assets which do not appear to follow the null hypothesis as we move into the tail of the distribution. The further into the tail we move the greater the proportion of assets that appear to be generated by a distribution other than the null hypothesis. In fact, for the last bar in the distribution, the observations with monthly alpha greater than 40%, only 65% can be explained by the null hypothesis, and 35% are determined to have been generated by a different distribution. It is these observations that exhibit positive alpha over and above the null hypothesis, that appear to be generating our momentum alphas.

Overall the chart shows 71,170 observations that we do not believe came from the null hypothesis, out of a total of over 1.8 million observations. It is fairly shocking that even in this portfolio that was designed with the intent of generating alpha, we only see evidence that 3.8% of observations appear generated by a distribution with non-zero mean.

While there are perhaps less assets than the researcher may have predicted that appear to have this excess return, it is clear that they are still capable of driving large monthly returns simply by their magnitude. Half of the observations outside the null show risk adjusted return of greater than 24% per month. They appear so far in the tail of the distribution, even in small number they can have a dramatic effect on portfolio alpha.

It might be reasonable to expect that the short portfolio exhibit similar returns to the long portfolio, simply on the other side of the distribution. However figure 2 shows that the short side of the distribution has a quite different distribution of non-null observations. In this case the observations that occur in greater frequency than predicted by the null hypothesis are those closest to zero. Compared to the long portfolio, where half of the observations not from the null exhibited risk adjusted returns over 24%, here we see all observations over the null frequencies between -24% and 0%.

In the case of the short portfolio we see more of the data exceed the expected null distribution, 6.3% of all observations, versus 3.8% for the long portfolio. Clearly, in this case their influence on the risk adjusted returns for the short portfolio are limited by their reduced distance from the mean. Even though there are more of these observations

than there were with the long portfolio, the long portfolio is still able to deliver more momentum return as a result.

As the short portfolio deviations from the mean are concentrated in the area of highest data concentration under the null, we do have a harder time identifying areas rich in non-null data. The most fertile areas of the short portfolio distribution are -6% to -12% risk adjusted returns, which show 18% of data not generated by the null.

Both tests of the long and short portfolio appear to conclude that momentum is being driven by a small number of assets each month, which casts some doubt on whether momentum is as prevalent in the market as previously thought. While alone this is interesting, these results are based on 50 years of continuous data, which begs the question, what if the market has changed over that period?

Is Momentum changing over time?

One of the most pertinent questions in asset pricing today is, to what extent have the identified anomalies been reduced over time. Or in another way of asking, to what extent are our markets getting more efficient over time. There are several arguments that could be made for the expected reduction of the anomaly over time.

Most obvious would be the increase in the accessibility of the market over time to capital. The times of requirements of large sums of money or connections to access the markets are long gone, as are physical location based limitations which once hindered the ability of people to apply their capital to the market. Today all an investor requires is an

internet connection, less than \$10 to cover the transaction fees, and their money may be invested in the market. The more capital acquires access to the market, the more liquid the market becomes, the less concentrated the market becomes, and by extension the more efficient we would expect it to be. If this is the case we would hypothesize that the common asset pricing anomalies should be reduced as the markets become more and more efficient.

Another argument for the increasing efficiency of the market is the publication of asset pricing literature. The publication of anomalies draws attention to their existence, or at least widens awareness of their existence. Assuming that there are at least some rational well capitalized investors that learned something from this information, it is reasonable to say they may decide to trade on this information. The more the anomalies are traded, the less value should remain. As such it is reasonable to suggest that anomalies should decrease in magnitude post publication (with presumably the visibility of the publication significantly impacting the extent to which more trading is focused on the anomaly).

To look at these arguments, and consider how the anomaly has changed over time, we will break up the risk adjusted returns by decade. The first three decades contained in the data (1960s, 1970s, 1980s) are all considered to be prior to the publication of momentum in the literature, by Jegedeesh and Titman (1993). Changes prior to 1990 may be due to the overall improvements in the efficiency of the market,

where as changes from the 1990s onwards may be considered, at least partially, influenced by the publication of research drawing attention to the anomaly.

Table 1 shows the results of breaking the momentum anomaly down by decade. The returns in this table are not true investment returns for the momentum portfolios, but equally weighted risk adjusted returns for all assets in the momentum portfolios that decade. The most recent decade in the data is only comprised of 3 years of data, and as such, volatility in the results should perhaps be expected.

The results show a noticeably consistent momentum anomaly prior to 1990, driven primarily by the long portfolio. In the 1990s, where the argument can be made that there was an informational shock to the market with respect to momentum, returns from the long portfolio dropped considerably. This poor performance of the long portfolio continued worsening into the 2000s. Overall the short portfolio appears to show less consistent trends in portfolio returns.

Beyond looking at aggregate returns, we can look at the number of assets driving those returns. While individual results in the long and short portfolio vary, the total percentage of assets exhibiting momentum appears to show a consistent decline. From 6.9% of assets in the investment portfolio showing non efficient alpha, to consistently below 4% since 1990. 25 years of fairly constant non null alpha may imply that the remaining assets are simply too illiquid to allow trading for the anomaly, or that it is a result of a delisting bias, or that the barrier to entry in the market is being lowered at a similar rate to that which investors can trade away prior inefficiency.

Beyond splitting out momentum by decade, it is also interesting to consider how momentum returns change over the investment horizon. The typical portfolios in a momentum strategy are from 3 to 12 months in duration. Of these portfolios the alpha have been typically documented as strong through 6 months, but then weaker at the longer horizons. The primary benefit of a longer holding period appears to be the reduction in transaction costs as opposed to real benefits to alpha.

Month one returns have also been documented to show typically poor performance, and it has been suggested that a better strategy would be to begin investing in month 2 in order to improve returns to the investment portfolio. There is however no deeper understanding to these poor month 1 returns. So it is of interest to look at them under the new methodology and see if we can draw insight from the availability of a true risk adjusted month 1 return.

Table 2 shows the return alpha broken out by investment month for the first 12 months post sorting period. The most striking facet of the data is the clear and gradual decline of the long portfolio alpha. The first 8 months show strong returns with a gradual decline, and with the declining alpha picking up speed as we approach month 12. There does not appear to be a noticeable change of behavior in month 1 in the long portfolio.

The short portfolio does show some discontinuity in month 1, but retains positive overall alpha (long – short). The following months do provide superior alpha, and it appears that month 2 through 8 offer the best overall returns on the portfolio. From the results of this analysis it could be inferred that the last 3 months of momentum do not

truly exhibit momentum, and we should perhaps think of momentum as a 9 month phenomenon, as opposed to a 12 month phenomenon, as is often referenced in the literature.

Also noticeable is the rapid decline in the number of assets that appear to be significant within this sample as we move further from the sorting months. Both long and short portfolios peak in the first 3 months of the investment period, and then see the percentage of true alpha assets decline from 7.76% to 1.15% in the short portfolio, and from 5.96% to 0.16% in the long portfolio. This dramatic decline paired with the declining alpha does cast doubt on the idea that momentum should be considered a 12 month anomaly when the last 3 months of the data show very weak evidence for momentum.

Jegadeesh and Titman (2001) present their second paper on momentum as an out of sample test of their first paper (Jegadeesh and Titman (1993)). The rationale behind this line of inquiry, was criticism that dogged the initial results claiming that the anomaly was merely the result of data mining. Given sufficient anomalies to check, it is trivial to show that the finance research community could find anomalies that provide significant alpha at any given significance level. The 2001 paper attempted to show that the result was consistent even in an out of sample test.

Unfortunately, the 2001 paper made a number of changes to the methodology that make direct comparison with the 1993 results difficult. One of the changes implemented was a removal of all assets with a price under \$5 at the time of investment from the momentum portfolio. The rationale for this data cut was to exclude small value

assets with high transaction costs from the portfolios, which insulates the results somewhat from accusations that these anomalies are not realizable, as they require trading illiquid assets.

In the spirit of the out of sample test, this methodology makes sense if the assets under \$5 also exhibit momentum. By removing them we would then be hampering our ability to find momentum, and making it all the more impressive if an out of sample test were able to find such momentum. The concern here is that, if the cut removed assets that did not exhibit momentum, or that exhibited reversals instead of momentum, then the cut is making it easier to find momentum in the remainder of our data. In this instance we have what is often referred to as data mining, but is perhaps more accurately referred to as model selection bias, or model overfitting.

Table 3 shows the results for momentum assets above and below the \$5 cutoff. The results above the \$5 cutoff are the same results shown before, with evidence of momentum. The assets under \$5 however do not show evidence of momentum. In fact they show a strong reversal, which is exactly the opposite of momentum. The challenge is not coming from the long portfolio, where returns remain positive and of significant magnitude, but from the short portfolio. In the short portfolio we see positive risk adjusted returns, far above those of the long portfolio. The momentum portfolio of assets below \$5 shows negative returns of approximately 2% per month.

We do not suggest that this reversal anomaly is tradable or offering positive alpha that has not yet been corrected for. There are many explanations that could make this

anomaly very untradeable. Transaction costs are one of the simplest explanations, and the effects of asset delisting, especially prevalent among small assets going through dramatic shocks, can also account for a significant segment of the returns on this hypothetical portfolio. Shumway (1997) and Shumway and Warther (1999) have done rigorous work on evaluating the delisting returns of assets in this sample. However even their work still leaves us unable to accurately and simply estimate accurate returns on delisting assets, often due to excessive time to liquidation prices appearing, and the unstable nature of those prices due to illiquidity.

Despite not necessarily being tradable, the existence of this reverse momentum in the sample that was removed from the data by Jegadeesh and Titman (2001) is a challenge. This result does ask questions of the methodology and raises the issue of model selection bias. Of primary importance here are questions as to whether this cutoff is appropriate for the stated reason of increasing the liquidity of the sample. This is not the focus of this paper, but is a question we leave open for future research.

Has Momentum disappeared?

It is well documented that the momentum anomaly does particularly badly in periods of extreme market stress. Single months have been observed giving double digit negative returns, which is clearly difficult for any strategy to overcome. With the increase in the frequency of periods of market uncertainty it is fair to question whether the momentum anomaly still exists in the market.

The argument that the momentum anomaly does not exist is essentially arguing that the anomaly is simply betting against a catastrophic market event. In periods where these events are less frequent or less severe the portfolio will make positive alpha. However, in times of extreme market stress, the portfolios can be expected to lose all that was gained in the good markets. Essentially the argument is that momentum is not a mispricing. That it is instead just betting on an event that is so rare, that in the short term it appears to be a mispricing.

Market volatility observed since 2000 has been more extreme than observed historically, with the collapse of the dot com bubble in 2001, and the real estate collapse of 2007 and subsequent period of market uncertainty. The argument has been made that these crashes, and subsequent losses from the momentum portfolio, show that the markets are pricing momentum correctly, and that momentum has not made alpha over the period.

To begin to look at whether this argument holds we will first look at the worst performing months for the momentum portfolio. Taking a look at raw return, there does appear to be a change in the performance of the strategy in recent periods. In the years covered by the momentum literature most commonly (1960-present) we have 5 complete decades of data. In that, the most recent decade contains the only three examples of monthly returns below -15%. Beyond that it contains 75% of observations below -10%, and 40% of observations below -5%. Months of poor performance do seem to appear more frequently in the most recent data. This is the type of distribution that

lends weight to the argument that momentum is not particularly profitable in the most recent periods, and may even have disappeared completely.

To evaluate this we first take a look at the most extreme observations in the data and compare the raw momentum returns to the monthly risk adjusted returns. Table 4 shows the results of this analysis. As we see from the table the worst performing months appear to cluster around 2000-2001 and the dot com bubble, and 2009, with the market still reeling from the housing market collapse. These 10 observations are the most extreme negative returns post 2000. The table shows each of them with raw returns of below -8.5%, which is a dramatic monthly loss. However, mitigating this loss we find that the risk adjusted return of the portfolios in these months was substantially lower than indicated by the raw returns.

A prime example of this is April 2009. In this month the momentum portfolio loses 35% of its raw value. This is a substantial loss by any standards, but the risk adjusted returns clearly show that the alpha of the month is only -10%. The loss in raw return was caused primarily by the risk profile of the portfolio during this month. While it was a bad month for momentum, it is not nearly as bad as it first appears.

Given these large distinctions between the raw returns and the risk adjusted returns for some of the extreme months in the portfolio it seems that it may be worth revisiting the question of whether the momentum portfolio offers positive alpha returns since the year 2000.

Table 5 shows raw returns on the momentum portfolio from the 1960s, 1970s, 1980s, 1990s and since 2000. It is clear, with 0.15% return since 2000, that momentum has struggled of late, relative to the recent performance in the 1990s (1.13%), or the 1980s (0.75%). Given these raw numbers it would be reasonable to form the opinion that momentum has been unsuccessful in recent years. The traditional risk adjusted momentum estimates also bear this message out, with each sample (1980s, 1990s, 2000-present) differing from the raw returns by small margins (5, 7 and 14 basis points). This is a reflection of the fact that portfolio level long window momentum estimates generally return estimates of momentum beta close to zero. Given this, the raw returns end up being very similar to the risk adjusted returns from the model.

We can see from the results of table 5 that for periods of relative calm in the markets, such as the 1990s, the monthly adjusted momentum estimates are very close to the traditional long window estimates. In the case of the 1990s, this resulted in estimated alphas within 7 basis points of each other. However, in more volatile markets, such as the period 2000 to present, the monthly adjusted alphas for the momentum portfolio are significantly higher than the traditional estimate, by 81 basis points. This reflects the inability of the traditional estimates to adjust for swings in holding period risk exposure. We see a similar story in the 1980s, with adjustable risk exposure providing estimates of alpha 77 basis points above the traditional estimates.

While this result does not change the fact that momentum portfolios did not realize significant raw returns over the latest decade or more, it does help to show that

there is evidence that momentum continues to generate significant risk adjusted returns. The evidence suggests that, over the last 12 years, momentum has often been caught holding very high risk loadings at inopportune moments, leading to high declines in raw return. We do not take a stance on whether the momentum anomaly systematically times the market, such that it has high loading on negative beta assets when the market increases, or whether this is a statistical anomaly that has occurred over the past 12 years of data, with no future predictability. Only an out of sample test could answer such a question, and for that we will have to wait for the data to become available.

III. Conclusions

Through the use of a novel model which allows a truly monthly estimate of an asset pricing strategy's risk exposure, we have been able to examine the momentum anomaly in more detail than prior literature, and begun to examine more closely how individual assets behave within the momentum strategy.

Upon inspection there is significant evidence that would suggest a waning of the power of the momentum anomaly. Our estimates suggest that less than 10% of assets within a momentum portfolio truly exhibits the anomaly. Even then, in recent years the number of assets that appear to exhibit this anomaly have decreased by approximately half since the 1960s. The data does not suggest economically significant returns to the momentum portfolio beyond 7 months, and a \$5 cutoff that has been applied for reasons of reducing transaction cost, actually inflated alphas and undermined prior work

attempting out of sample testing of the momentum anomaly. Beyond that the average return of the momentum anomaly since 2000 is seen to be significantly smaller than in prior years, and the realized return is in fact negative since 2000 for some momentum portfolios. All of this gives the authors reason to doubt the persistence of the momentum anomaly.

Yet the anomaly stands strong. We provide evidence that shows strong risk adjusted returns of the momentum portfolio when the time varying risk exposure of the portfolio is appropriately measured. The momentum anomaly continues to generate alpha, even in the most difficult market conditions for momentum since the 1930s. We are also able to show that the poorest days for raw returns in the momentum portfolio history are, in large part driven by risk exposure, and do not undermine the ability of the portfolio to generate alpha as much as was previously thought.

Beyond the scope of this study, there is clearly work to be done identifying which assets are driving the momentum anomaly. We have identified that a small number of assets appear to be driving the anomaly on a month to month basis. If there were a way to further focus the momentum anomaly on these assets there is the potential to learn far more about the nature of the anomaly, and obtain cleaner measurement of it.

Also there remain questions about the risk exposure of the anomaly. While it is reassuring to those who might be invested, that positive alpha is indeed generated by the momentum strategy to this day, it is still a strategy with poor raw returns. The large varying betas which have cost the strategy so much in raw return are still somewhat

unknown. If they can be predicted reasonably far in advance, then they can be hedged, and the investor can extract the portfolio alpha. If not then the investor is simply strapped to a time varying beta they cannot control or hedge against. It is of course still possible that this is not a trend, but simply a facet of the most recent data that will not prove to be persistent. Only time, and an out of sample test, will be able to answer that question with any real degree of certainty.

Beyond the area of asset pricing and momentum the methodology remains potentially useful for looking at other areas of asset pricing. There remains no reason this methodology could not be applied to size and value anomalies, to attempt to better understand their time varying risk profiles, and their changes over time. Consistent evidence across the main asset pricing anomalies would certainly be another way to provide support to the conclusions of research, before true out of sample testing is available.

References

- Asness, Clifford S., Tobias J. Moskowitz, and Lasse H. Pedersen, 2013, Value and momentum everywhere, *Journal of Finance* 68, 929-985.
- Bajgrowicz, Pierre and Oliver Scaillet, 2012, Technical trading revisited: False discoveries, persistence tests, and transaction costs, *Journal of Financial Economics* 106, 473-491.
- Barras, Laurent, Oliver Scaillet and Russ Wermers, 2010, False discoveries in mutual fund performance: Measuring luck in estimated alphas, *Journal of Finance* 65, 179-216.
- Benjamini, Yoav and Yosef Hochberg, 1995, Controlling the false discovery rate: A practical and powerful approach to multiple testing, *Journal of the Royal Statistical Society. Series B* 57, 289-300.
- Benjamini, Yoav and Daniel Yekutieli, 2001, The control of the false discovery rate in multiple testing under dependency, *Annals of Statistics* 29, 1165-1188.
- Boudoukh, Jacob, Roni Michaely, Matthew Richardson and Michael R. Roberts, 2007, On the importance of measuring payout yield: implications for empirical asset pricing, *Journal of Finance* 62, 877-915.
- Chan, Louis, Yasushi Hamao and Joseph Lakonishok, 1991, Fundamentals and stock returns in Japan, *Journal of Finance* 46, 1739-1764.
- Chordia, Tarun, and Lakshmanan Shivakumar, 2002, Momentum, business cycle, and time-varying expected returns, *Journal of Finance* 57, 985-1019.
- Cooper, Michael J., Roberto C. Gutierrez, and Allaudeen Hameed, 2004, Market states and momentum, *Journal of Finance* 59, 1345-1365.
- Chui, Andy, Sheridan Titman and John Wei, 2010. Individualism and momentum around the world. *Journal of Finance* 65, 361-392.
- Dimson, Elroy, 1979, Risk measurement when shares are subject to infrequent trading, *Journal of Financial Economics* 7, 197-226.
- Fama, Eugene F., and Kenneth R. French, 1996, Multifactor explanations of asset pricing anomalies, *Journal of Finance* 51, 55-84.
- Fama, Eugene F., and Kenneth R. French, 1995, Size and book-to-market factors in earnings and returns, *Journal of Finance* 50, 131.
- Fama, Eugene F., French, Kenneth R., 1998, Value versus growth: the international evidence. *Journal of Finance* 53, 153-193.
- Fama, Eugene F., French, Kenneth R., 2012, Size value and momentum in international stock returns. *Journal of Financial Economics* 105, 457-472.
- Ferson, Wayne E., and Campbell R. Harvey, 1999, Conditioning variables and the cross section of stock returns, *Journal of Finance* 54, 1325-1360.
- Gorman, Michael J., and Wu, Yagru, 2015, Bias in momentum estimates, *Working Paper*.
- Griffin, John M., 2002, Are the Fama and French factors global or country specific? *Review of Financial Studies* 15, 783-803.
- Griffin, John M., Xiuqing Ji and J. Spencer Martin, J.S., 2003, Momentum investing and business cycle risk: evidence from pole to pole. *Journal of Finance* 58, 2515-2547.

- Harvey, Campbell R., Yan Liu and Heqing Zhu, 2015, ...and the cross section of expected returns, *Review of Financial Studies*, *Forthcoming*.
- Harvey, Campbell R. and Yan Liu, 2015, Multiple testing in economics, *Working Paper*
- Hochberg, Yosef and Benjamini, Y., 1990, More powerful procedures for multiple significance testing, *Statistics in Medicine* 9, 811-818.
- Hou, Kewei, G. Andrew Karolyi and Bong Chan Kho, 2011, What factors drive global stock returns? *Review of Financial Studies* 24, 2527-2574.
- Jagannathan, Ravi, and Zhenyu Wang, 1996, The conditional CAPM and the cross-section of expected returns, *Journal of Finance* 51, 3-53.
- Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* XLVIII, 65.
- Jegadeesh, Narasimhan, and Sheridan Titman, 2001, Profitability of momentum strategies: An evaluation of alternative explanations, *Journal of Finance* 56, 699-720.
- Karolyi, G. Andrew, Stulz, Rene M., 2003, Are assets priced globally or locally? In: Constantinides, G.M., Harris, M., Stulz, R. (Eds.), *Handbook of the Economics of Finance (Chapter 16)*, Elsevier Science, North Holland, pp. 2551-2595.
- Kosowski, Robert, Allan Timmermann, Russ Wermers and Hal White, 2006, Can mutual fund "stars" really pick stocks? New evidence from a bootstrap analysis, *Journal of Finance* 61, 2551-2595.
- Lo, Andrew W. and A. Craig MacKinlay, 1990, When are contrarian profits due to stock market overreaction? *Review of Financial Studies* 3, 175-205.
- Lettau, Martin, and Sydney Ludvigson, 2001, Resurrecting the (C)CAPM: A cross-sectional test when risk premia are time-varying, *Journal of Political Economy* 109, 1238-1287.
- Lewellen, Jonathan, and Stefan Nagel, 2006, The conditional CAPM does not explain asset-pricing anomalies, *Journal of Financial Economics* 82, 289-314.
- Lustig, Hanno N., and Stijn G. Van Nieuwerburgh, 2005, Housing collateral, consumption insurance, and risk premia: An empirical perspective, *Journal of Finance* 60, 1167-1219.
- Moskowitz, Tobias J., and Mark Grinblatt, 1999, Do industries explain momentum? *Journal of Finance* 54, 1249-1290.
- Ľuboš Pástor, and Robert F. Stambaugh, 2003, Liquidity risk and expected stock returns, *Journal of Political Economy* 111, 642-685.
- Patton, Andrew J. and Allan Timmermann, 2010, Monotonicity in asset returns: New tests with applications to the term structure, the CAPM and portfolio sorts, *Journal of Financial Econometrics* 98, 605-625.
- Rouwenhurst, K. Geert, 1998, International momentum strategies. *Journal of Finance* 53, 267-284
- Sarkar, Sanat K., 2002, Some results on false discovery rate in stepwise multiple testing procedure, *Annals of Statistics* 30, 239-257.
- Sadka, Ronnie, 2006, Momentum and post-earnings-announcement drift anomalies: The role of liquidity risk, *Journal of Financial Economics* 80, 309-349.
- Santos, Tano, and Pietro Veronesi, 2006, Labor income and predictable stock returns, *Review of Financial Studies* 19, 1-44.

- Schweder, T. and E. Spjotvoll, 1982, Plots of p-values to evaluate many tests simultaneously, *Biometrika* 69, 439-502.
- Shanken, Jay, 1990, Intertemporal asset pricing: An empirical investigation, *Journal of Econometrics* 45, 99-120.
- Shumway, Tyler, and Vincent A. Warther, 1999, The delisting bias in CRSP's NASDAQ data and its implications for the size effect, *Journal of Finance* LIV, 2361.
- Shumway, Tyler, 1997, The delisting bias in CRSP data, *Journal of Finance* LII, 327.
- Storey, John D., 2003, The positive false discovery rate: A bayesian interpretation and the q-value, *Annals of Statistics* 31, 2013-2035.
- Wang, Jun, and Yangru Wu, 2011, Risk adjustment and momentum sources, *Journal of Banking and Finance* 35, 1427.

Table 1**Risk adjusted monthly momentum alpha by decade**

Data obtained from CRSP, Daily and Monthly return data from 1960-2013. This table shows an implementation of the 6 month sorting period, 12 month investment period, momentum portfolio. The long portfolio represent the top decile of assets based on the past 6 months of raw return, and the short portfolio represents assets that were in the bottom decile based on the past 6 months raw return. Assets under \$5 are removed. Each asset each month has its daily data regressed against the market to estimate a risk adjusted return for that month. This table shows those results, aggregated by decade.

The estimates of true alpha assets are gained from estimating the False Discovery Rate for each portfolio (long and short) each decade for the side of the distribution of interest. For example, on the long portfolio we estimate the FDR of all assets with risk adjusted return above zero. The true alpha percentage is just 1-FDR.

Decade	Average Alpha		Non-Efficient Alpha		
	Long	Short	Long	Short	Total
1960	2.21%	0.42%	5.91%	7.93%	6.89%
1970	1.59%	0.42%	4.43%	6.10%	5.11%
1980	1.31%	-0.08%	3.09%	6.30%	4.11%
1990	0.86%	-0.33%	2.60%	5.34%	3.39%
2000	0.98%	0.11%	4.92%	0.84%	3.62%
2010	0.11%	-1.32%	1.19%	9.06%	3.69%
All	1.16%	0.02%	3.73%	5.26%	4.26%

Table 2**Alpha and FDR estimate by investment month**

Data obtained from CRSP, Daily and Monthly return data from 1960-2013. This table shows an implementation of the 6 month sorting period, 12 month investment period, momentum portfolio. The long portfolio represent the top decile of assets based on the past 6 months of raw return, and the short portfolio represents assets that were in the bottom decile based on the past 6 months raw return. Assets under \$5 are removed. Each asset each month has its daily data regressed against the market to estimate a risk adjusted return for that month. This table shows those results, aggregated by investment month given that performance was measured over month -5 to 0.

Investment Month	Average Alpha		Assets with True Alphas	
	Long	Short	Long	Short
1	1.45%	0.02%	4.47%	3.07%
2	1.54%	-0.44%	5.96%	6.76%
3	1.38%	-0.49%	5.32%	7.67%
4	1.28%	-0.33%	4.97%	7.37%
5	1.26%	-0.15%	4.64%	6.80%
6	1.29%	-0.08%	4.75%	6.19%
7	1.35%	-0.12%	5.01%	6.21%
8	1.12%	0.04%	3.81%	5.90%
9	0.92%	0.21%	2.46%	4.85%
10	0.84%	0.38%	1.56%	3.91%
11	0.77%	0.54%	1.15%	2.93%
12	0.56%	0.77%	0.16%	1.15%

Table 3**Alpha under and over the \$5 cutoff**

Data obtained from CRSP, Daily and Monthly return data from 1960-2013. This table shows an implementation of the 6 month sorting period, 12 month investment period momentum portfolio. The long portfolio represent the top decile of assets based on the past 6 months of raw return, and the short portfolio represents assets that were in the bottom decile based on the past 6 months raw return.

In this table we compare the returns on assets under the \$5 price cutoff suggested by Jegadeesh and Titman and compare to assets over that cutoff.

Investment Month	Average Alpha			
	Assets under \$5		Assets over \$5	
	Long	Short	Long	Short
1	2.03%	6.67%	1.45%	0.02%
2	3.05%	5.10%	1.54%	-0.44%
3	3.08%	4.75%	1.38%	-0.49%
4	3.09%	4.61%	1.28%	-0.33%
5	3.28%	4.45%	1.26%	-0.15%
6	3.26%	4.34%	1.29%	-0.08%
7	3.65%	4.02%	1.35%	-0.12%
8	3.56%	4.14%	1.12%	0.04%
9	3.37%	4.34%	0.92%	0.21%
10	3.09%	4.34%	0.84%	0.38%
11	2.93%	4.30%	0.77%	0.54%
12	2.63%	4.56%	0.56%	0.77%

Table 4**Raw momentum returns and alpha for extreme months**

The Table lists the 10 most extreme months of raw returns for the momentum portfolio since the year 2000 (coincides with the end of Jegadeesh and Titman's out of sample momentum test, Jegadeesh and Titman 2001). Raw momentum returns were obtained from Ken French's website.

Monthly alpha estimates are obtained by regressing daily CRSP data over 1 month samples for each asset in the momentum portfolio. These results are then aggregated within their portfolios, within their investment months.

Year	Month	Momentum Raw Return	Monthly Alpha	Difference
2009	4	-34.6%	-10.1%	24.5%
2001	1	-25.0%	-18.6%	6.4%
2002	11	-16.3%	-12.7%	3.6%
2009	5	-12.3%	-1.7%	10.6%
2009	3	-11.5%	-7.5%	4.0%
2003	4	-9.4%	-1.2%	8.3%
2000	5	-9.1%	-2.4%	6.7%
2009	8	-8.9%	-3.9%	5.0%
2001	11	-8.6%	0.5%	9.1%
2000	4	-8.5%	-3.4%	5.1%

Table 5**Comparison of adjusted monthly alpha with traditional estimates**

The average raw returns on the momentum portfolio are obtained from Ken French's website, as are estimates of the market risk factors necessary to estimate alpha at both the monthly and daily level. Daily price data is obtained from CRSP from 1960-2012. Average momentum portfolio alphas are calculated by simply running a single regression using portfolio level returns over the entire sample. Monthly adjusted momentum alphas are calculated using regressions each month for each asset in the momentum portfolio, using a single month of daily return data. These alphas are then aggregated to the portfolio level.

Time period	Average Raw Return	Average Momentum Portfolio Alpha	Monthly Adjusted Momentum Alpha
2000-2012	0.15%	0.20%	1.01%
1990-1999	1.13%	1.06%	1.19%
1980-1989	0.75%	0.61%	1.38%
1970-1979	0.83%	0.85%	1.17%
1960-1969	0.92%	0.92%	1.78%

Figure 1**Alpha of the long momentum portfolio**

The chart below is built from CRSP data spanning from 1960 to 2013. For each month assets are ranked on their prior 6 month performance, and then the top decile are bought. The positions are held for the following 12 months. Any asset that has returns in these portfolios has a regression run using one month of daily data for every month that they are present in the portfolio. Each of these regressions gives a single risk adjusted return for the month, and these are plotted below. Only assets with a price above \$5 are considered valid investment assets.

The false discovery rate methodology is then applied to determine the fraction of each histogram category that we can claim is not generated by the null hypothesis. The fraction of observations that cannot be justified by the null are identified in white.

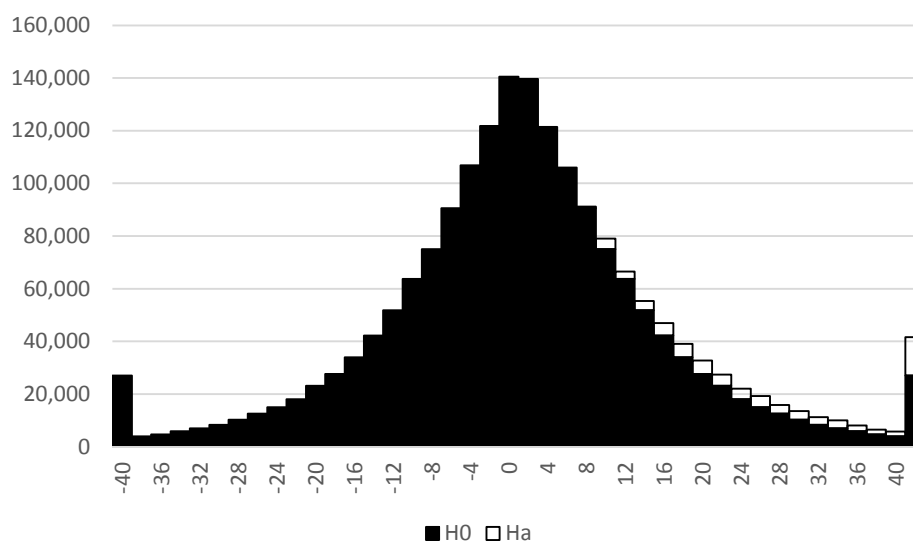


Figure 2**Alpha of the short momentum portfolio**

The chart below is built from CRSP data spanning from 1960 to 2013. For each month assets are ranked on their prior 6 month performance, and then the bottom decile assets are bought. The positions are held for the following 12 months. Any asset that has returns in these portfolios has a regression run using one month of daily data for every month that they are present in the portfolio. Each of these regressions gives a single risk adjusted return for the month, and these are plotted below. Only assets with a price above \$5 are considered valid investment assets.

The false discovery rate methodology is then applied to determine the fraction of each histogram category that we can claim is not generated by the null hypothesis. The fraction of observations that cannot be justified by the null are identified in white.

