Managing the Risk of Momentum*

Pedro Barroso† and Pedro Santa-Clara‡

This version: April 2012

Abstract

Compared to the market, value or size risk factors, momentum has offered investors the highest Sharpe ratio. However, momentum has also had the worst crashes, making the strategy unappealing to investors with reasonable risk aversion. We find that the risk of momentum is highly variable over time and quite predictable. The major source of predictability does not come from systematic risk but from specific risk. Managing this time-varying risk virtually eliminates crashes and nearly doubles the Sharpe ratio of the momentum strategy. Risk-managed momentum is a much greater puzzle than the original version.

*We thank Eduardo Schwartz for helpful comments.

†PhD student at Nova School of Business and Economics. E-mail: pmsb@novasbe.pt.

‡Millennium Chair in Finance. Nova School of Business and Economics and NBER and CEPR. Campus de Campolide, 1099-032 Lisboa, Portugal. Phone +351-21-3801600. E-mail: psc@novasbe.pt.

1 Introduction

Momentum is a pervasive anomaly in asset prices. Jegadeesh and Titman (1993) find that previous winners in the US stock market outperform previous losers by as much as 1.49 percent a month. The Sharpe ratio of this strategy exceeds the Sharpe ratio of the market itself, as well as the size and value anomalies. Momentum returns are even more of a puzzle since they are negatively related to the market and the value risk factors. From 1927 to 2011, momentum had a monthly excess return of 1.75 percent per month controlling for the Fama-French factors. This result has led researchers to use momentum as an additional risk factor.[1] Momentum is not just a US stock market anomaly. Momentum has been documented in European equities, emerging markets, country stock indices, industry portfolios, currency markets, commodities and across asset classes.[2] Grinblatt and Titman (1989,1993) found most mutual fund managers incorporate momentum of some sort in their investment decisions, so relative strength strategies are widespread among practitioners.

But the remarkable performance of momentum comes with occasional large crashes.[3] In 1932, the winners-minus-losers strategy delivered a -91.59 percent return in just two months. In 2009 momentum experienced another significant crash of -73.42 percent over three months. Even the large returns of momentum do not compensate an investor with reasonable risk version for these sudden crashes that take decades to recover from.

The two most expressive momentum crashes occurred as the market rebounded following large previous declines. One explanation for this pattern is the time-varying systematic risk of the momentum strategy. Grundy and Martin (2001) show that momentum has significant negative beta following bear markets.[4] They

[1] Carhart (1997)

[2] See Rouwenhorst (1998,1999) for international evidence, Asness, Liew, and Stevens (1997) for country stock indices, Moskowitz and Grinblatt (1999) for industry portfolios, Okunev and White (2003) and Menkhof, Sarno, Schmeling, and Schrimpf (2012) for currency markets, Erb and Harvey (2006) for commodities and Asness, Moskowitz, and Pederson (2011) for momentum across asset classes.

[3] Daniel and Moskowitz (2011)

[4] Following negative returns for the overall market, winners tend to be low-beta stocks and the reverse for losers. Therefore the winner-minus-losers strategy will have a negative beta.

argue that hedging this time-varying market exposure produces stable momentum returns but Daniel and Moskowitz (2011) show that using betas in real time does not avoid the crashes.

In this work we propose a different method to manage momentum risk. We estimate the risk of momentum by the realized variance of daily returns and find that it is highly predictable. An auto-regression of monthly realized variances yields an out-of-sample (OOS) R-square of 57.82 percent. This is 19.01 percentage points (p.p.) more than a similar auto-regression for the variance of the market portfolio which is already famously predictable.[5]

Making use of this predictability in risk management leads to substantial economic gains. We scale the long-short portfolio by its realized volatility in the previous 6 months, thereby obtaining a strategy with constant volatility. The Sharpe ratio improves from 0.53 for unmanaged momentum to 0.97 for its risk-managed version. But the most important benefit comes from a reduction in crash risk. The excess kurtosis drops from 18.24 to 2.68 and the left skew improves from -2.47 to -0.42. The minimum one-month return for momentum is -78.96 percent while for risk-managed momentum is -28.40 percent. The maximum drawdown of momentum was -96.69 percent versus -45.20 percent for its risk-managed version.

To assess the economic significance of our results, we evaluate the benefits of risk management for a risk-averse investor using a power utility function with Constant Relative Risk Aversion (CRRA) of four. The representative investor holding the market portfolio has an annual certainty equivalent of 0.14 percent. Adding momentum to the portfolio reduces this certainty equivalent to -5.46 percent. However, combining the market with risk-managed momentum already achieves a certainty equivalent of 13.54 percent. We find that the main benefit of risk management comes in the form of smaller crash risk. This alone provides a gain of 14.96 p.p. in annual certainty equivalent.

One pertinent question is why managing risk with realized variances works while using time-varying betas does not. To answer this question we decompose the risk of momentum into systematic risk (from time-varying exposure to the market) and specific risk. We find that the systematic component is only 23 percent of total risk on average, so most of the risk of momentum is specific. This

[5] Engle and Bollerslev (1986), Schwert (1989).

specific risk is more persistent and predictable than the systematic component. The OOS R-square of the specific component is 47 percent versus 21 percent for the systematic component. This is why hedging with time-varying betas fails: it focuses on the smaller part of risk and also the less predictable one.

The work that is more closely related to ours is Grundy and Martin (2001) and Daniel and Moskowitz (2011). But their work focuses on the time-varying systematic risk of momentum, while we focus on momentum's specific risk. Our results have the distinct advantage of offering investors using momentum strategies an effective way to manage risk without forward-looking bias. The resulting risk-managed strategy deepens the puzzle of momentum.

This paper is structured as follows. Section 2 discusses the long-run properties of momentum returns and its exposure to crashes. Section 3 shows that momentum risk varies substantially over time in a highly predictable manner. We analyze the implications of such predictability for risk management in Section 4. Section 5 decomposes the gains of risk management according to the moments of returns. In Section 6 we decompose the risk of momentum and study the persistence of each of its components. Finally, Section 7 presents our conclusions.

2 Momentum in the long run

We compare momentum to the Fama-French factors using a long sample of 85 years of monthly returns from July 1926 to December 2011.[6] This is the same sample period as in Daniel and Moskowitz (2011).

Figure 1 presents the cumulative returns of each factor. As each factor consists of a long-short strategy, we construct the series of returns assuming the investor puts \$1 in the risk-free asset at the beginning of the sample, buys \$1 worth of the long portfolio and sells the same amount of the short portfolio. Then in each subsequent month, the strategy fully reinvests the accumulated wealth in the risk-free asset, assuming a position of this same size in the long and short leg of the portfolio. The winners-minus-losers (WML) strategy offered an impressive performance for investors. One dollar fully reinvested in the momentum strategy

[6] See annex for a description of the data.

grew to \$68,741 by the end of the sample. This compares to the cumulative return of \$2,136 from simply holding the market portfolio.

There would be nothing puzzling about momentum's returns if they corresponded to a very high exposure to risk. However, running an OLS regression of the WML on the Fama-French factors gives (t-stats in parenthesis):

$$\begin{array}{lllll} r_{WML,t} = & 1.752 & -0.378\ r_{RMRF,t} & -0.249\ r_{SMB,t} & -0.677\ r_{HML,t} \\ & (7.93) & (-8.72) & (-3.58) & (-10.76) \end{array}$$

so momentum provided abnormal returns of 1.75 percent per month with *negative* exposure to the Fama and French (1992) risk factors. This amounts to a 21 percent per year abnormal return and the negative loadings on other risk factors imply momentum actually diversified risk in this extended sample.

Table 1 compares descriptive statistics for momentum in the long-run with the Fama-French factors. Buying winners and shorting losers has provided large returns of 14.46 percent per year, with a Sharpe ratio higher than the market. The impressive excess returns of momentum, its high Sharpe ratio and negative relation to other risk factors, particularly the value premium, make it look like a free lunch for investors.

But as Daniel and Moskowitz (2011) show there is a dark side to momentum. Its large gains come at the expense of a very high excess kurtosis of 18.24 combined with a pronounced left-skew of -2.47. These two features of the distribution of returns of the momentum strategy imply a very fat left tail – significant crash risk. The apparent free lunch of momentum returns can very rapidly turn into a free fall, wiping out decades of returns.

Figure 2 shows the performance of momentum in the two most turbulent decades for the strategy: the 1930's and the 2000's.

In July and August 1932, momentum had a cumulative return of -91.59 percent. From March to May 2009, momentum had another large crash of -73.42 percent. These short periods leave an enduring impact on cumulative returns. For example, someone investing one dollar in the WML strategy in July 1932 would only recover it in April 1963, 31 years after and with considerably less real value. This puts the risk to momentum investing in an adequate long-run perspective.

Both in 1932 and in 2009, the crashes happened as the market rebounded after experiencing large losses.[7] This leads to the question of whether investors could predict the crashes in real time and hedge them away.

Grundy and Martin (2001) show that momentum has a substantial time-varying loading on stock market risk. The strategy ranks stocks according to returns during a formation period, for example the previous 12 months. If the stock market performed well in the formation period, winners tend to be high-beta stocks and losers low-beta stocks. So the momentum strategy, by shorting losers to buy winners, has by construction a significant time-varying beta: positive after bull markets and negative after bear markets. They argue that hedging this time-varying risk produces stable returns, even in pre-WWII data, when momentum performed poorly. In particular, the hedging strategy would be long in the market portfolio whenever momentum has negative betas, hence mitigating the effects of rebounds following bear markets, which is when momentum experiences the worst returns. But the hedging strategy in Grundy and Martin (2001) uses forward looking betas, estimated with information investors did not have in real time. Daniel and Moskowitz (2011) show that using betas estimated solely on ex-ante information does not avoid the crashes and portfolios hedged in real time perform even worse than unhedged momentum.

3 The time-varying risk of momentum

One possible cause for the excess kurtosis of momentum is time-varying risk.[8] The very high excess kurtosis of 18.24 of the momentum strategy (more than twice the market portfolio) leads us to study the dynamics of its risk and compare it with the market (RMRF), value (HML) and size (SMB) risk factors.

For each month, we compute the realized variance RV_t from daily returns in the previous 21 sessions. Let $\{r_d\}_{d=1}^{D}$ be the daily returns and $\{d_t\}_{t=1}^{T}$ the time series of the dates of the last trading sessions of each month. Then the realized variance of factor i in month t is:[9]

[7] Daniel and Moskowitz (2011) argue this is due to the option-like payoffs of distressed firms in bear markets.

[8] See, for example, Engle (1982) and Bollerslev (1987).

[9] Correcting for serial correlation of daily returns does not change the results significantly.

$$RV_{i,t} = \sum_{j=0}^{20} r^2_{d_t - j} \tag{1}$$

Figure 4 shows the monthly realized volatility of momentum. This varies substantially over time, from a minimum of 3.04 percent (annualized) to a maximum of 127.87 percent.

Table 2 shows the results of AR (1) regressions with the realized variances of the WML, RMRF, SMB and HML:

$$RV_{i,t} = \alpha + \rho RV_{i,t-1} + \varepsilon_t \tag{2}$$

Panel A presents the results for RMRF and WML, for which we have daily data available from 1927:03 to 2011:12. Panel B adds the results for HML and SMB, for which daily data is available from 1963:07 onwards.

Momentum returns are the most volatile. From 1927:03 to 2011:12, the average realized volatility of momentum was 15.03, more than the 12.81 of the market portfolio. For the 1963:07 onwards sample, the average realized volatility of momentum was 16.40, the highest when compared to RMRF, SMB and HML.

In the full sample period, the standard deviation of monthly realized volatilities is higher for momentum (12.26) than the market (7.82). Panel B confirms this result for the other factors in the 1963:07 onwards sample. So the risk of momentum is the most variable.

The risk of momentum is also the most persistent. The $AR(1)$ coefficient of momentum in the 1963:07 sample is 0.77, which is 0.19 more than the market for the same sample period and higher than the estimate for SMB and HML.

To check the out-of-sample (OOS) predictability of risk, we use a training sample of 240 months to run an initial AR(1) and then use the estimated coefficients and last observation of realized variance to forecast the realized variance in the following month. Then each month we use an expanding window of observations to produce OOS forecasts and compare these to the accuracy of the historical mean $\overline{RV}_{i,t}$. As a measure of goodness of fit we estimate the OOS R-square as:

$$R^2_{i,OOS} = 1 - \frac{\sum_{t=S}^{T-1} (\widehat{\alpha}_t + \widehat{\rho}_t RV_{i,t} - RV_{i,t+1})^2}{\sum_{t=S}^{T-1} (\overline{RV}_{i,t} - RV_{i,t+1})^2} \tag{3}$$

where S is the initial training sample, $\widehat{\alpha}_t, \widehat{\rho}_t$ and $\overline{RV}_{i,t}$ are estimated with information available only up to time t.

The last column of table 2 shows the OOS R-squares of each auto-regression. The AR(1) of the realized variance of momentum has an OOS R-square of 57.82 percent (full sample), which is 50 percent more than the market. For the period from 1963:07 to 2011:12, the OOS predictability of momentum-risk is twice that of the market. Hence more than half of the risk of momentum is predictable, the highest level among risk factors. In the next section we explore the potential of this predictability for risk management.

4 Risk-managed momentum

We use estimates of momentum risk to scale the exposure to the strategy in order to have constant risk over time.

For each month we compute a variance forecast $\hat{\sigma}_t^2$ from daily returns in the previous 6 months.[10] Let $\{r_{WML,t}\}_{t=1}^{T}$ be the monthly returns of momentum and $\{r_{WML,d}\}_{d=1}^{D}, \{d_t\}_{t=1}^{T}$ be, as above, the daily returns and the time series of the dates of the last trading sessions of each month.

The variance forecast is:

$$\hat{\sigma}_t^2 = 21 \sum_{j=0}^{125} r^2_{d_{t-1-j}} / 126 \tag{4}$$

As WML is a zero-investment and self-financing strategy we can scale it without constraints. We use the forecasted variance to scale the returns:

[10] We also used one-month and three-month realized variances as well as exponentially-weighted moving average (EWMA) with half-lifes of 1, 3 and 6 months. All worked well with nearly identical results.

$$r_{WML^*,t} = \frac{\sigma_{\text{target}}}{\hat{\sigma}_t} r_{WML,t} \tag{5}$$

where $r_{WML,t}$ is the unscaled or plain momentum, $r_{WML^*,t}$ is the scaled or risk-managed momentum, and σ_{target} is a constant corresponding to the target level of volatility. Scaling corresponds to having a weight in the long and short legs that is different from one and varies over time, but it is still a self-financing strategy, so the choice of the constant is arbitrary. We pick a target corresponding to an annualized volatility of 12 percent.[11]

Figure 3 shows the cumulative returns of risk-managed momentum compared to plain momentum. The risk-managed momentum strategy achieves a higher cumulative return with less risk. So there are economic gains to risk-management of momentum. The scaled strategy benefits from the large momentum returns when it performs well and effectively shuts it off in turbulent times, thus mitigating momentum crashes. As a result, one dollar invested in risk-managed momentum grows to \$6,140,075 by the end of the sample, nearly 90 times more than the plain momentum strategy.[12] Also, the risk-managed strategy no longer has variable and persistent risk, so risk management indeed works.[13]

Table 3 provides a summary of the economic performance of WML^* in 1927-2011. The risk-managed strategy has a higher average return, with a gain of 2.04 (p.p. per year), with substantially less standard deviation (less 10.58 p.p. per year). As a result, the Sharpe ratio of risk-managed portfolios almost doubles from 0.53 to 0.97. The most important gains of risk-management show up in the improvement in the higher order moments. Managing the risk of momentum drops the excess kurtosis from a very high value of 18.24 to just 2.68 and reduces the left

[11] The annualized standard deviation from monthly returns will be higher than 12% as volatilities at daily frequency are not directly comparable to those at lower frequencies due to small positive autocorrelation of daily returns.

[12] This difference in cumulative returns is fundamentally due to risk management successfully avoiding the two momentum crashes. But in the post-war period from 1946 to 2007, not including the crashes, the Sharpe ratio of momentum was 0.86, versus 1.15 for risk-managed momentum. So risk management also contributes to performance in not so turbulent times.

[13] The AR(1) coefficient of monthly squared returns is only 0.14 for the scaled momentum versus 0.40 for the original momentum. Besides, the auto-correlation of momentum is significant up to 15 lags while only 1 lag for risk-managed momentum. So persistence in risk is much smaller for the risk-managed strategy.

skew from -2.47 to -0.42. This practically eliminates the crash risk of momentum. Figure 5 shows the density function of momentum and its risk-managed version. Momentum has a very long left tail which is much reduced in its risk-managed version.

The benefits of risk-management are specially important in turbulent times. Figure 6 shows the performance of risk-managed momentum in the decades with the most impressive crashes. The scaled momentum manages to preserve the investment in the 1930's. This compares very favorably to the pure momentum strategy which loses 90 percent in the same period. In the 2000's simple momentum lost 28 percent of wealth, because of the crash in 2009. Risk-managed momentum ends the decade up 88 percent as it not only avoids the crash but also captures part of the positive returns of 2007-2008.

Figure 7 shows the weights of the scaled momentum strategy over time – interpreted as the dollar amount in the long or short leg. These range between the values of 0.13 and 2.00, reaching the most significant lows in the early 1930's, in 2000-02, and in 2008-09. On average, the weight is just 0.90, slightly less than full exposure to momentum. As these weights depend only on ex ante information this strategy could actually be implemented in real time.

5 Economic Significance: An Investor Perspective

Momentum offers a trade-off between an appealing Sharpe ratio, obtained from the first two moments of its distribution, and less appealing higher order moments, such as high kurtosis and left skewness. An economic criterion is needed to assess whether this trade-off is interesting. Risk management offers improvements to momentum across the board, higher expected returns, lower standard deviation and crash risk. Still it is pertinent to evaluate the relative economic importance of each of these contributions.

We use an utility-based approach to discuss the appeal of momentum to a representative investor. We adopt the power utility function as it has the advantage of taking into consideration higher order moments instead of focusing merely on the mean and standard deviation of returns. This is particularly suitable as momentum has a distribution far from normal. The utility of returns is:

$$U(r) = \frac{(1+r)^{1-\gamma}}{1-\gamma} \tag{6}$$

where γ is the Constant Coefficient of Relative risk aversion (CRRA). Bliss and Panigirtzoglou (2004) estimate γ empirically from risk-aversion implicit in one-month options on the S&P and the FTSE and find a value very close to 4. This is a more plausible value for CRRA than previous estimates featured in the equity premium puzzle literature using utility over consumption. So we adopt this value for CRRA. We obtain the certainty equivalent from the utility of returns:

$$CE(r) = \{(1-\gamma)E[U(r)]\}^{\frac{1}{1-\gamma}} - 1 \tag{7}$$

This states the welfare a series of returns offers the investor in terms of an equivalent risk-free annual return, expressed in a convenient unit of percentage points (p.p.) per year.

For an economic measure of the importance of the mean return, variance and higher order moments, we use a Taylor series approximation to expected utility around the mean:

$$E[U(r)] = U(\bar{r}) + \frac{1}{2}U''(\bar{r})E(r-\bar{r})^2 + \phi_3(r) \tag{8}$$

where ϕ_3 is the Lagrangian rest corresponding to the utility from moments with order greater than 2. From this we obtain the certainty equivalent due to each moment:

$$\begin{aligned} CE(\mu_1) &= \{(1-\gamma)U(\bar{r})\}^{\frac{1}{1-\gamma}} - 1 && (9) \\ CE(\mu_2) &= \left\{(1-\gamma)[U(\bar{r}) + \frac{1}{2}U''(\bar{r})E(r-\bar{r})^2]\right\}^{\frac{1}{1-\gamma}} - CE(\mu_1) - 1 && (10) \\ CE(\mu_{i>2}) &= CE(r) - CE(\mu_1) - CE(\mu_2) && (11) \end{aligned}$$

We compute the certainty equivalent from annual overlapping returns, an adequate horizon from an investor perspective. Also, one-year horizons capture better the occasional large drawdowns of momentum documented in Section 2.

Table 4 shows the decomposition of the certainty equivalent for the representative investor holding the market portfolio. It also assesses whether it is optimal to deviate from the market portfolio to include (risk-managed) momentum.

The first row shows the results for holding only the market portfolio. The mean return had a positive contribution for the certainty equivalent of 11.72 percent per year, but the variance of returns reduces this by 7.39 p.p. Higher order moments diminuish the certainty equivalent by a further 4.18 p.p. As a result the certainty equivalent of the market portfolio was only 0.14 percent per year.

Adding momentum to the market portfolio increases returns. As a result the certainty equivalent of the mean return increases from 11.72 percent per year to 28.51 percent. The higher standard deviation partially offsets this gain by reducing the certainty equivalent by 6.51 p.p. Still, looking only at the first two moments of the combined portfolio leads to the conclusion that the investor is better off including momentum.

But the increase in higher order risk – the momentum crashes – reduces the certainty equivalent by 15.89 p.p. per year. As a result, including momentum actually dampens the economic performance of the market portfolio. The certainty equivalent of the market plus momentum is -5.46 percent per year versus 0.14 percent of the market only. The high Sharpe ratio of momentum does not compensate the investor for bearing the increased crash risk. So, in spite of the impressive cumulative returns of momentum in the long-run, crash risk is so high that a reasonable risk-averse investor would rather just hold the market portfolio.

This illustrates with an economic measure how far the distribution of momentum is from normality. Indeed, momentum has a distribution with many small gains and few, but extreme, large losses. Taking this into account the momentum puzzle of Jegadeesh and Titman (1993) is substantially diminished. Our discussion in terms of utility shows that rare observations with large losses are significant enough to change the interest of momentum for a risk-averse investor holding the market.

In contrast, risk-managed momentum produces large economic gains. These come from higher returns when compared to the market (a 19.92 p.p. gain) and less crash risk than the market with plain momentum (a 14.96 p.p. gain). As a result, the annual certainty equivalent of the market with risk-managed momentum is 13.54 percent, which compares very favorably to the 0.14 percent of the market alone and even more so with the -5.46 percent of the market combined with simple momentum.

Comparing the strategies with plain and scaled momentum, risk management produces a 3.13 p.p. gain in returns and a 15.86 p.p. gain from reduction of risk. So the bulk of the gains comes from less risk, especially high-order moments. Essentially, scaling momentum eliminates non-normal risk without sacrificing returns.

6 Anatomy of momentum risk

A well documented result in the momentum literature is that momentum has time-varying market risk (Grundy and Martin (2001)). This is an intuitive finding since after bear markets winners tend to be low-beta stocks and the inverse for losers. But Daniel and Moskowitz (2011) show that using betas to hedge risk in real time does not work. This contrasts with our finding that the risk of momentum is highly predictable and managing it offers strong gains. Why is scaling with forecasted variances so different from hedging with market betas? We show it is because time-varying betas are not the main source of predictability in momentum risk.

We use a CAPM regression to decompose the risk of momentum into systematic and specific risk:

$$RV_{wml,t} = \beta_t^2 RV_{rmrf,t} + \sigma_{e,t}^2 \tag{12}$$

The realized variances and betas are estimated with 6-months of daily returns. On average, the systematic component $\beta_t^2 RV_{rmrf,t}$ accounts for only 23 percent of the total risk of momentum. Almost 80 percent of the momentum risk is specific. Also, the different components do not have the same predictability. Table 5 shows the results of an AR (1) on each component of risk.

Either in-sample or out-of-sample (OOS), β_t^2 is the least predictable component of momentum risk. Its OOS R-square is only 5 percent. The realized variance of the market also has a small OOS R-square of 7 percent. When combined, both elements of systematic risk component already show more predictability (OOS R-square of 21 percent) but still less than the realized variance of momentum (OOS R-square of 44 percent). The most predictable component of momentum variance is the specific risk with an OOS R-square of 47 percent, more than double the predictability of the systematic risk.

Hedging with betas alone, as in Daniel and Moskowitz (2011) fails because most of the risk is left out.

7 Conclusion

Unconditional momentum has a distribution that is far from normal, with huge crash risk. We find that taking this crash risk into consideration, momentum is not appealing for a risk-averse investor.

Yet, the risk of momentum is highly predictable. Managing this risk eliminates exposure to crashes and increases the Sharpe ratio of the strategy substantially.

As a result, the momentum puzzle shows up stronger in its risk-managed version and the case for a peso explanation of momentum returns is severely weakened.

References

[1] Asness, C. S., J. M. Liew, and R. L. Stevens (1997). "Parallels between the cross-sectional predictability of stock and country returns", Journal of Portfolio Management 23, pp. 79-87.

[2] Asness, C. S., T. Moskowitz, and L. Pedersen (2011). "Value and momentum everywhere", University of Chicago, working paper.

[3] Bliss, R., and N. Panigirzoglou (2004). "Option-Implied Risk Aversion Estimates", Journal of Finance, Vol LIX, 1, pp. 407-446.

[4] Bollerslev, T. (1987), "A Conditional Heteroskedastic Time Series Model for Speculative Prices and Rates of Return", Review of Economics and Statistics, 69, pp. 542-547.

[5] Carhart, M. (1997). "On Persistence in Mutual Fund Performance", The Journal of Finance, vol. 52, no. 1, pp. 57-82.

[6] Daniel, K., and T. Moskowitz (2011). "Momentum Crashes", Working paper.

[7] Engle, R. (1982). "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of U.K. Inflation", Econometrica 50, pp. 987-1008.

[8] Engle, R., and T. Bollerslev (1986). "Modeling the Persistence of Conditional Variances", Econometric Reviews, 5, pp. 1-50.

[9] Erb, C. B., and C. R. Harvey (2006). "The strategic and tactical value of commodity futures", Financial Analysts Journal 62, pp. 69-97.

[10] Fama, E. and K. French (1992). "The Cross-Section of Expected Stock Returns", The Journal of Finance, vol. 47, no. 2, pp. 427-465.

[11] Fama, E. and K. French (1996). "Multifactor Explanations of Asset Pricing Anomalies", The Journal of Finance, vol. 51, no. 1, pp. 55-84.

[12] Grinblatt, M., and S. Titman (1989). "Mutual fund performance: An analysis of quarterly portfolio holdings", Journal of Business, 62, pp. 394-415

[13] Grinblatt, M., and S. Titman (1993). "Performance measurement without benchmarks: An examination of mutual fund returns", Journal of Business, 66, pp. 47-68.

[14] Grundy, B., 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, pp. 29-78.

[15] Jegadeesh, N., and S. Titman (1993). "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency", The Journal of Finance, vol. 48, n^o 1, pp. 65-91.

[16] Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2012). "Currency Momentum Strategies", Journal of Financial Economics, forthcoming

[17] Moskowitz, T., and M. Grinblatt (1999). "Do Industries Explain Momentum?", Journal of Finance, vol. 54, no. 4, pp. 1249-1290.

[18] Okunev, J., and D. White (2003). "Do momentum-based strategies still work in foreign currency markets?", Journal of Financial and Quantitative Analysis 38, pp. 425-447.

[19] Schwert, G. (1989). "Why Does Stock Market Volatility Change over Time?", Journal of Finance, 44, pp. 1115-1154.

[20] Rouwenhorst, K. (1998). "International momentum strategies", Journal of Finance, 53, pp. 267-284.

[21] Rouwenhorst, K. (1999). "Local return factors and turnover in emerging stock markets", Journal of Finance 54, pp. 1439-1464.

Annex: Data sources

We obtain daily and monthly returns for the market portfolio, the high-minus-low, the small-minus-big, the ten momentum-sorted portfolios and the risk-free (one-month Treasury-bill return) from Kenneth French's data library. The monthly data is from July 1926 to December 2011 and the daily data is from July 1963 to December 2011.

For the period from July 1926 to June 1963, we use daily excess returns on the market portfolio (the value-weighted return of all firms on NYSE, AMEX and Nasdaq) from the Center for Research in Security Prices (CRSP). We also have daily returns for ten value-weighted portfolios sorted on previous momentum from Daniel and Moskowitz (2011). This allows us to work with a long sample of daily returns for the winner-minus-losers (WML) strategy from August 1926 to December 2011. We use these daily returns to calculate the realized variances in the previous 21, 63 and 126 sessions at the end of each month.

For the momentum portfolios, all stocks in the NYSE, AMEX and Nasdaq universe are ranked according to returns from month t-12 to t-2, then classified into deciles according to NYSE cutoffs. So there is an equal number of NYSE firms in each bin. The WML strategy consists on shorting the lowest (loser) decile and a long position in the highest (winner) decile. Individual firms are value weighted in each decile. Following the convention in the literature, the formation period for month t excludes the returns in the preceding month. See Daniel and Moskowitz (2011) for a more detailed description of how they build momentum portfolios. The procedures (and results) are very similar to those of the Fama-French momentum portfolios for the 1963:07-2011:12 sample.

Tables

	Max	Min	Mean	STD	KURT	SKEW	SR
RMRF	38.27	-29.04	7.33	18.96	7.35	0.17	0.39
SMB	39.04	-16.62	2.99	11.52	21.99	2.17	0.26
HML	35.48	-13.45	4.50	12.38	15.63	1.84	0.36
WML	26.18	-78.96	14.46	27.53	18.24	-2.47	0.53

Table 1. The long-run performance of momentum compared to the Fama-French factors. 'KURT' stands for excess kurtosis and 'SR' for Sharpe ratio. The sample returns are from 1927:03 to 2011:12.

	α	t-stat	ρ	t-stat	R^2	OOS R^2	$\overline{\sigma}$	σ_σ
	Panel A: 1927:03 to 2011:12							
RMRF	0.0010	6.86	0.60	23.92	36.03	38.81	12.81	7.82
WML	0.0012	5.21	0.70	31.31	49.10	57.82	15.03	12.26
	Panel B: 1963:07 to 2011:12							
RMRF	0.0009	5.65	0.58	17.10	33.55	25.46	13.76	8.48
SMB	0.0004	8.01	0.33	8.32	10.68	-8.41	7.36	3.87
HML	0.0001	4.88	0.73	25.84	53.55	53.37	6.68	4.29
WML	0.0009	3.00	0.77	29.29	59.71	55.26	16.40	13.77

Table 2. AR (1) of 1-month realized variances. The realized variances are the sum of squared daily returns in each month. The AR (1) regresses each realized variance on its lag and a constant. The OOS R-squared uses the first 240 months to run an initial regression so producing an OOS forecast. Then uses an expanding window of observations till the end of the sample. In panel A the sample period is from 1927:03 to 2011:12. In panel B we repeat the regressions for RMRF and WML and add the same information for the HML and SMB. The last two columns show, respectively, the average realized volatility and its standard deviation.

	Max	Min	Mean	STD	KURT	SKEW	SR
WML	26.18	-78.96	14.46	27.53	18.24	-2.47	0.53
WML^{*}	21.95	-28.40	16.50	16.95	2.68	-0.42	0.97

Table 3. Momentum and the economic gains from scaling. The first row presents as a benchmark the economic performance of plain momentum from 1927:03 to 2011:12. The second row presents the performance of risk-managed momentum. The risk-managed momentum uses the realized variance in the previous 6 months.

	$CE(\mu_1)$	$CE(\mu_2)$	$CE(\mu_{i>2})$	$CE(r)$
RM	11.72	-7.39	-4.18	0.14
RM+WML	28.51	-13.90	-20.07	-5.46
RM+WML*	31.64	-13.00	-5.11	13.54

Table 4. The economic performance of momentum for a representative investor. The first row shows the performance of the market portfolio. The second row combines the market portfolio with momentum and the third one with scaled momentum. The first column shows the certainty equivalent of the mean return of each strategy. The second and third columns present the contribution to the certainty equivalent of standard deviation and higher moments, respectively. The last column shows the certainty equivalent obtained from annual non-overlapping returns. The returns are from 1927:03 to 2011:12. The decomposition uses a Taylor expansion of the utility function around the mean return of the portfolio with a CRRA of 4.

	α	t-stat	ρ	t-stat	R^2	R^2_{OOS}
σ^2_{wml}	0.0012	2.59	0.70	12.58	0.49	0.44
σ^2_{rmrf}	0.0012	4.29	0.50	7.37	0.25	0.07
β^2	0.3544	6.05	0.21	2.83	0.05	0.05
$\beta^2\sigma^2_{rmrf}$	0.0007	2.73	0.47	6.80	0.22	0.21
σ^2_{ε}	0.0007	2.69	0.72	13.51	0.52	0.47

Table 5. Risk decomposition of momentum risk. Each row shows the results of an AR (1) for 3-month, non-overlapping periods. The first row is for the realized variance of the WML and the second one the realized variance of the market. The third row is squared beta, estimated as a simple regression of 63 daily returns of the WML on RMRF. The fourth row is the systematic component of momentum risk and the last row the specific component. The OOS R-squares use an expanding window of observations after an initial in-sample period of 20 years.

Figures

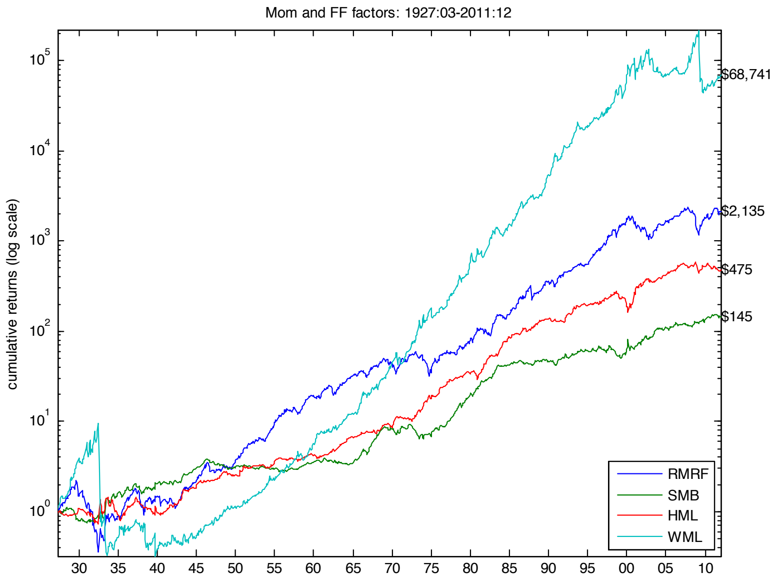

Figure 1. The long-run cumulative returns of momentum compared to the Fama-French factors. Each strategy consists on investing \$1 at the beginning of the sample in the risk-free rate and combine it with the respective long-short portfolio. The proceeds are fully reinvested till the end of the sample. On the right is the terminal value of each strategy.

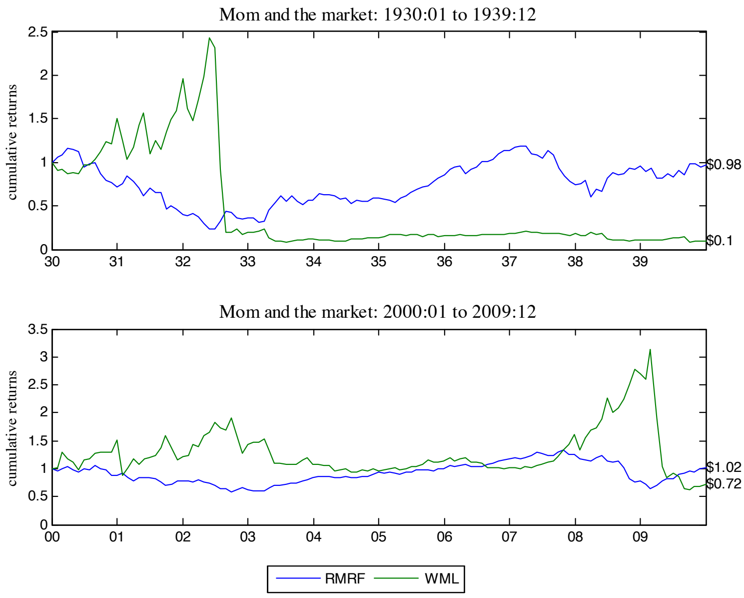

Figure 2. Momentum crashes. The figure plots the cumulative return and terminal value of the momentum and market portfolio strategies in its two most turbulent periods: the 1930's and the 2000's.

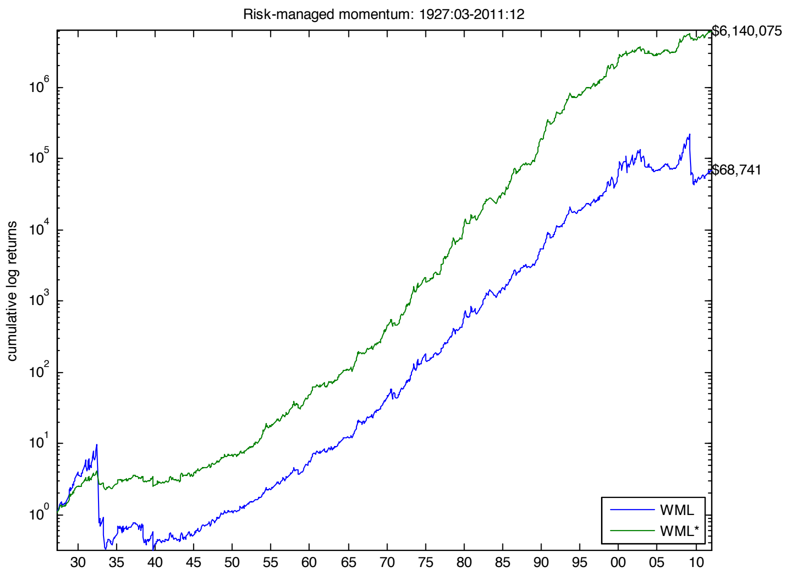

Figure 3. The long-run performance of risk-managed momentum. The risk-managed momentum (WML*) scales the exposure to momentum using the realized variance in the previous 6-months. In the beginning of the sample the strategy invests $1 in the risk-free asset and combines it with the long-short portfolio. The proceeds are fully reinvested till the end of the sample. On the right is the terminal value of the strategy.

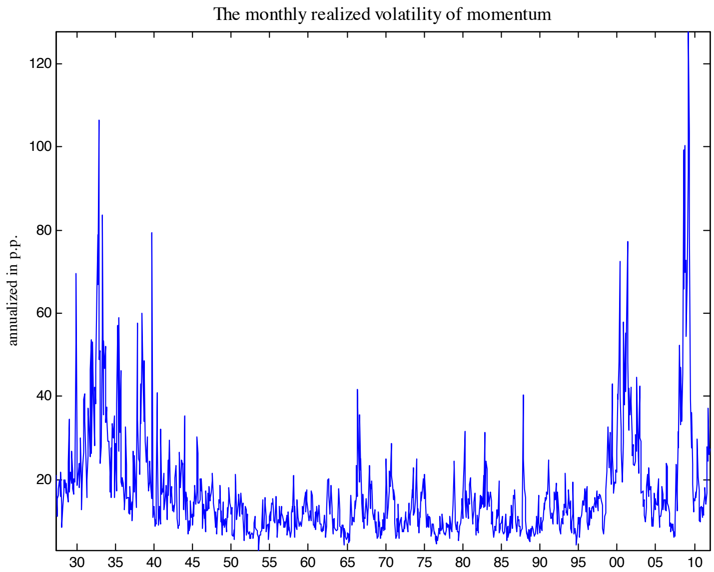

Figure 4. The realized volatility of momentum obtained from daily returns in each month from 1927:03 to 2011:12.

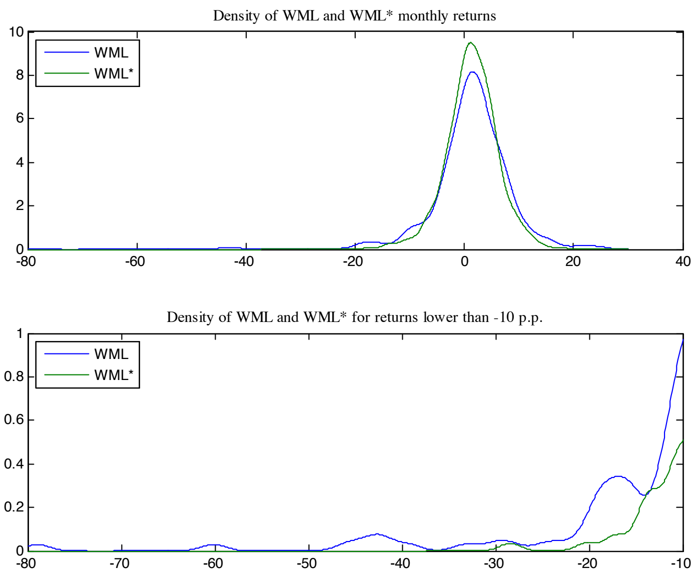

Figure 5. The density of plain momentum (WML) and risk-managed momentum (WML*). The risk-managed momentum uses the realized variance in the previous 6 months to scale the exposure to momentum. The returns are from 1927:03 to 2011:12.

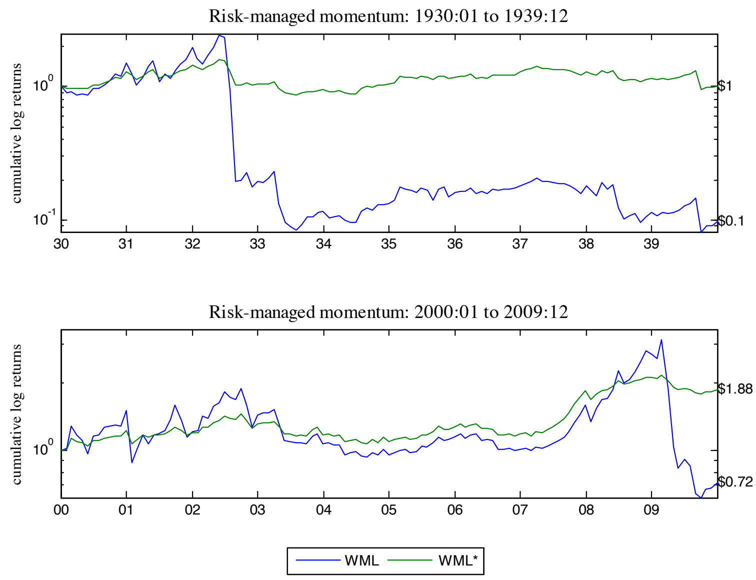

Figure 6. The benefits of risk-management in the 1930's and the 00's. The risk-managed momentum (WML*) uses the realized variance in the previous 6 months to scale the exposure to momentum.

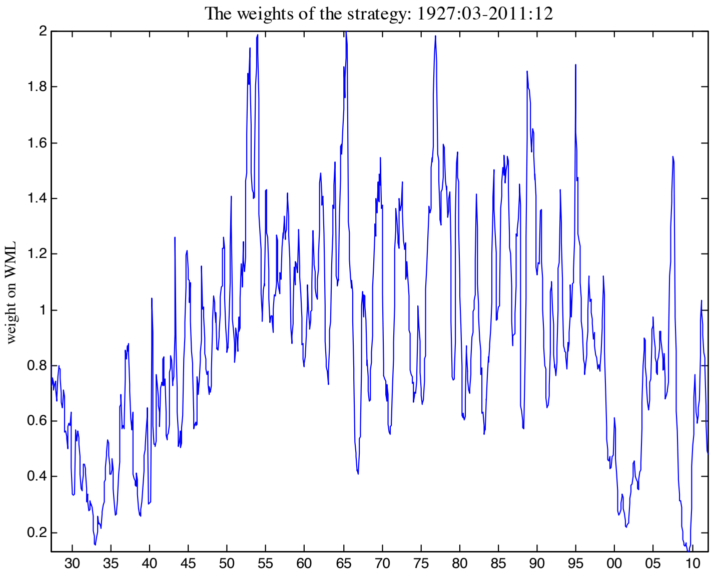

Figure 7. Weights of the scaled momentum. The risk-managed momentum uses the realized variance in the previous 6 months to scale the exposure to momentum.