Re-Evaluation of the Low-Risk Anomaly in Finance via Matching

GOV2001 Final Paper

Yang Lu [*], Daniel Wu [†], Kwok Yu[‡]

April 29, 2013

Abstract

The long-term success of low-risk stocks over high-risk stocks runs contrary to the basic finance principle that risk is compensated with higher expected returns. Our paper examines this low-risk anomaly using Coarsened Exact Matching to balance high and low-risk stock portfolios on industry, company size, and trading volume. After matching, we find that the low-risk anomaly still exists but has a more muted effect than in previous studies, especially when beta is used as a measure of risk. We also find moments in which the low-risk anomaly does not hold, most notably during the dot-com bubble. To our knowledge, we are the first to apply matching techniques to the study of the low-risk anomaly, and our findings complicate various previous explanations of this phenomenon.

[*]yang.lu2014@gmail.com
[†]danielwu@fas.harvard.edu
[‡]kwok_yu@harvard.edu

1 Introduction

The low-risk anomaly is one of the longest-standing puzzles in finance. Contrary to basic principles of finance, low-risk stocks have long outperformed high-risk stocks. Risk of a stock or portfolio is typically measured using volatility (standard deviation of returns) or beta (slope of regression line between asset returns and market risk premium).[1] According to the Modern Portfolio Theory, market participants are assumed to be risk averse and investing is a tradeoff between risk and return. Thus, assets with a higher expected return are generally riskier. Similarly, the Sharpe-Lintner Capital Asset Pricing Model (CAPM) that is commonly used to price securities assumes a positive correlation between risk and return. However, historical stock returns show that both high-volatility and high-beta stocks consistently underperform low-volatility and low-beta stocks.

Behaviorial finance models have offered two primary explanations of the low-risk anomaly: irrational market participants and limits to arbitrage. First, irrational investors who exhibit overconfidence, representativeness or extrapolation biases, or those who have a preference for lotteries can drive up demand and prices for high-risk stocks. Second, institutional benchmarks and leverage constraints discourage investment in low-risk stocks and can lead fund managers to exacerbate the anomaly rather than arbitrage it away.

Additionally, some have questioned the assumptions of CAPM (Fama and French, 1992; Jensen et al., 1972) while others argue that problems of sample bias and spurious correlations from data mining are responsible for the appearance of an anomaly (Bernstein and Fabozzi, 1998).

We contribute to the literature methodologically by introducing the use of matching techniques to this particular topic. Using Coarsened Exact Matching (CEM) (Iacus et al.,

[1] Market risk premium = expected market returns - risk-free rate.

2011), we match low-risk portfolios with high-risk portfolios on industry, company size, and trading volume to create a more balanced data set in an attempt to isolate the relationship between risk and return. Replicating results from the influential piece titled *Benchmarks as Limits to Arbitrage: Understanding the Low-Volatility Anomaly* (Baker et al., 2011), we use monthly stock-level data from 1968 through 2008 to demonstrate the existence of the low-risk anomaly.[2] We then pre-processed the lowest and highest-risk portfolios with CEM, improving the balance by almost 10%, and re-calculated the returns on the pruned portfolios. Although the low-risk anomaly still exists after matching, we find that the spread of cumulative returns between the two portfolios narrows, especially when using beta as a measure of risk.

Our work ultimately builds on the work of others in the field (Baker et al., 2011; Blitz et al., forthcoming; Vliet et al., 2011), who have contributed methodologically by subjecting the low-risk anomaly to further controls informed by theory. In contrast to others' work, we found a moment, during the dot-com bubble, in which the low-risk anomaly disappears. During this period, instead of finding that high-risk stocks underperform, we find the opposite. This complicates the existing propositions by scholars including those from (Fama and French, 1992; Baker et al., 2011). Our study opens up questions about the effect of certain time periods on the low-risk anomaly (e.g. October 1999 through June 2000) and sheds some light on the muddy relationship between risk and return, ultimately suggesting to scholars the possibility of alternative explanations to the anomaly.

[2]See Appendix for the replication results.

2 Data

The data set used in the paper comes from the Center for Research in Security Prices (CRSP). It includes all stocks traded in the United States from January 1968 through December 2008 on a monthly basis. The stock-level data include the month-end close price, the total monthly return, the monthly trading volume, the number of outstanding shares as well as the major industry group to which a stock belongs based on the Standard Industrial Classification (SIC) codes available on CRSP.

We calculated the volatility and beta for each stock in the data set using up to 60 months of trailing data (the data set thus went back to January 1963). Volatility was calculated by taking the standard deviation of returns for each stock. Beta was calculated by regressing stock returns on excess market returns over the risk-free rate. Any stock that did not have at least 24 months of data within the 60-month trailing window was removed for that particular month. Market returns and risk-free rates were obtained from Kenneth French's data library at Dartmouth.[3]

3 Methodology

We used matching techniques to re-evaluate the low volatility and low beta anomalies. Matching is typically used in social science research to improve balance between treatment and control groups by eliminating observations that do not have a counterpart in the comparison group (Rubin, 1973, 1974; Ho et al., 2007). However, we believe that it can also be useful to pre-process and prune two stock portfolios so that the relationship between returns and risk can be more easily isolated. For example, in a given month, if the lowest volatility

[3]See `http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html` for more information.

portfolio contains a handful of mega-cap, high-volume oil companies, but no such companies exist in the highest volatility portfolio, those stocks can be removed from the study to ensure balance between the two portfolios. It is important to note that removing stocks from the study does not bias portfolio return estimates if they are removed based on confounders (major industry groups, company size, trading volume) or "treatment assignment" (volatility or beta) but not on outcomes (stock returns).

Matching has also been shown to reduce model dependence (Ho et al., 2007; Iacus et al., 2011). In the context of this study, improving balance between the highest and lowest volatility or beta portfolios could reduce a portfolio's over-reliance on particular industries or types of stocks in generating gains or losses. Although matching does not allow us to prove causality in our study, we use these techniques in an attempt to gain some additional insight into the relationship between risk and returns.

Various matching techniques were considered in the study include nearest neighbor propensity score matching (PSM), nearest neighbor Mahalanobis distance matching, and CEM. We chose CEM because it provides the best balance (as measured by the L1 statistic) between the two portfolios. Balance or the similarity of stock features is our primary concern as we aim to re-evaluate the low-risk anomaly by controlling for the background covariates of the portfolios.

4 Empirical Analysis

For each of the 492 months in the data set, stocks were sorted by their trailing volatility and assigned to one of five volatility portfolios (low to high) weighted by market capitalization. The same procedure was followed to create five beta portfolios (low to high). We ignored

transaction costs when rebalancing the portfolios each month.

We then used the CEM algorithm to match stocks in the lowest volatility or beta portfolio (quintile 1) with stocks in the highest volatility or beta portfolio (quintile 5) each month based on major industry group, company size, and trading volume.

- **Industry Group:** The classification into various industry groups is based on the SIC code. Each stock is assigned to one of 83 major industry groups based on the SIC codes. Matching on a stock to which it belongs is important in that different sectors may have different trends and are susceptible to various economic conditions. Imagine an average stock in the mining sector would behave differently from one in the agriculture sector.

- **Trading Volume:** Trading volume reflects the liquidity of a stock in the market. The higher the volume, the easier the stock is to be purchased. Liquidity can be a factor in the risk-return relationship. Trading volume is measured using the natural logarithm of trading volume for the prior month and is coarsened into quartiles. Matching on volume quartiles prevents comparing stocks of very different liquidity and helps to see whether other factors exist that contribute to the aforementioned anomaly.

- **Company Size:** Company size is defined by its market capitalization, which is the product of price per share and the number of outstanding shares. Company size is measured using the natural logarithm of market capitalization for the prior month and is coarsened into quartiles. Stocks of different sizes tend to behave differently. Large caps are followed by more analysts and institutions; thus chances of mispricing are lower. Matching on size quartiles is important in that it minimizes the chances of comparing stocks with vastly different degrees of price efficiency.

Table 1 details the various statistics of both the original and the matched portfolios from January 1968 through December 2008. In Table 1 and Panel A of Figure 1, one dollar invested in the matched lowest volatility portfolio in January 1968 grows to $41.86 by December 2008. One dollar invested in the matched highest volatility portfolio during the same time period was worth $0.21 by December 2008. As a comparison, the original lowest and highest volatility quintile portfolios were worth $62.72 and $0.40, respectively. Thus, the low volatility anomaly still exists even after matching on major industry group, market cap, and trading volume. However, in many months, the spread between the lowest and highest volatility portfolios has been reduced compared to the original quintile portfolios.

In Panel B of Figure 1, one dollar invested in the matched lowest beta portfolio in January 1968 grows to $22.99 by December 2008. One dollar invested in the matched highest volatility portfolio during the same time period was worth $7.08 by December 2008. As a comparison, the pre-matched lowest and highest beta portfolios were worth $59.84 and $6.25, respectively. Thus, the spread between the lowest and highest beta portfolios shrinks dramatically after matching, although the lowest beta portfolio still produces a superior return. What is even more noteworthy is that during the period from October 1999 through June 2000, the matched quintile 5 portfolio outperformed the matched quintile 1 portfolio.

Figure 1: Post-Matching Returns by Volatility and Beta Quintile, Jan. 1968-Dec. 2008.

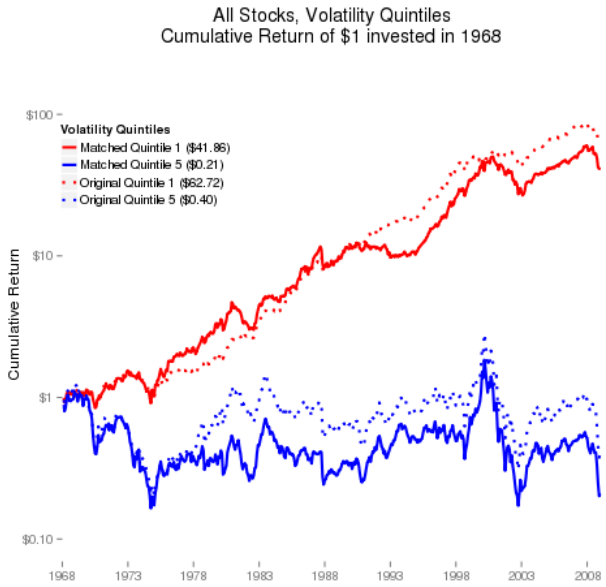

(a) Panel A

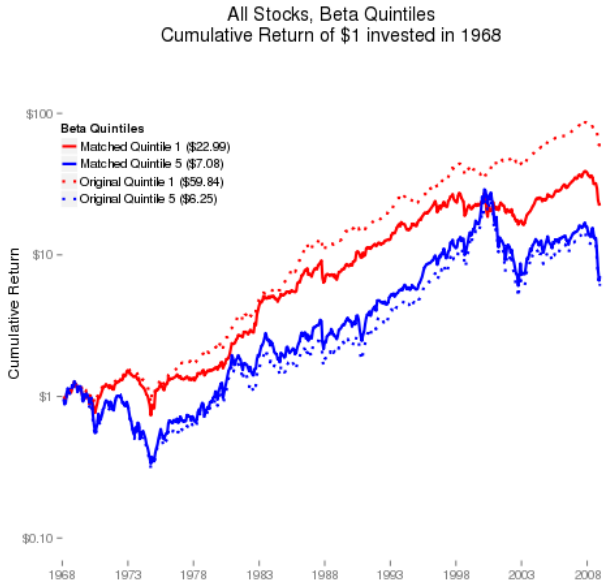

(b) Panel B

Notes: For each month, we use Coarsened Exact Matching to match stocks in the lowest volatility/beta portfolio to stocks in the highest volatility/beta portfolio based on major industry group, company size, and trading volume. Each stock is assigned to one of 83 major industry groups based on the SIC codes. Company size is measured using the natural logarithm of market capitalization for the prior month and is coarsened into quartiles. Trading volume is measured using the natural logarithm of trading volume for the prior month and is coarsened into quartiles. In January 1968, $1 is invested into each portfolio according to capitalization weights and matching weights, and the cumulative value of each portfolio is calculated each month until December 2008. We estimated volatility and beta by using up to 60 months of trailing returns (i.e. return data starting as early as January 1963) obtained from CRSP. At the end of each month, we rebalanced each portfolio, excluding all transaction costs. The original (pre-matched) returns shown in Figure 1 are represented with dotted lines for comparison purposes.

Table 1: Returns by Volatility and Beta Quintile, Original vs. Matched, January 1968-December 2008

	Volatility				Beta			
	Quintile 1		Quintile 5		Quintile 1		Quintile 5	
	Original	Matched	Original	Matched	Original	Matched	Original	Matched
Cumulative value in 2008 of $1 invested in 1968	$62.72	$41.86	$0.40	$0.21	$59.84	$22.99	$6.25	$7.08
Geometric Average $R_p - R_f$	4.43%	3.43%	-7.99%	-9.57%	4.32%	2.00%	-1.25%	-0.94%
Average $R_p - R_f$	5.29%	4.93%	-2.74%	-3.84%	5.06%	3.51%	2.54%	2.84%
Standard Deviation	13.09%	17.32%	31.86%	33.33%	12.20%	17.54%	27.16%	27.15%
Sharpe Ratio	0.40	0.28	-0.09	-0.12	0.42	0.20	0.09	0.10
Average $R_p - R_m$	1.11%	0.75%	-6.93%	-8.02%	0.88%	-0.67%	-1.64%	-1.34%
Tracking Error	7.02%	9.67%	20.37%	22.0%	10.15%	11.37%	14.08%	14.63%
Information Ratio	0.16	0.08	-0.34	-0.36	0.09	-0.06	-0.12	-0.09
Beta	0.75	0.91	1.68	1.73	0.60	0.85	1.56	1.53
Alpha	2.17	1.11	-9.77	-11.09	2.57	1.11	-4.00	-11.09

Notes: For each month, we formed matched portfolios based on the original quintile 1 and 5 portfolios, respectively. Stocks in all portfolios are cap-weighted. We estimated volatility and beta by using up to 60 months of trailing returns (i.e., return data starting as early as January 1963). The return on the market, R_m, and the risk-free rate, R_f, are from Ken French's data library at Dartmouth. The information ratio uses the market return for the relevant universe, all stocks in the first five columns and the top 1,000 stocks in the last five columns. Average returns are monthly averages multiplied by 12. Standard deviation and tracking error are monthly standard deviations multiplied by the square root of 12.

Table 2: Monthly Imbalance and Portfolio Size, January 1968-December 2008

	Mean L1	Std. Dev. of L1	Mean N	Std. Dev. of N
Original (Volatility)	0.982	0.007	1792.5	754.5
Matched (Volatility)	0.822	0.080	480.0	308.7
Original (Beta)	0.863	0.049	1781.7	765.9
Matched (Beta)	0.755	0.082	941.0	582.8

Notes: For each month from January 1968 to December 2008, we calculate L1 as a measure imbalance between the lowest and highest volatility/beta portfolios. An L1 value of 1 indicates total imbalance, with lower L1 values indicating less imbalance. The mean and standard deviation of the monthly L1 are displayed for the original data set and the matched data set for both volatility and beta. The mean and standard deviation of the monthly sample size (N) is also shown.

As shown in Table 2, CEM matching reduces imbalance from a monthly mean L1 value of 0.982 (1 indicates total imbalance) in the original volatility data set to 0.822 in the matched volatility data set. As a consequence of matching, the mean monthly sample size was reduced from 1792.5 to 480.

Similarly, matching reduces imbalance from a monthly mean L1 value of 0.863 in the original beta data set to 0.755 in the matched beta data set. As a consequence of matching, the mean monthly sample size was reduced from 1781.7 to 941.

5 Discussion and Conclusion

To our knowledge, this study is the first of its kind in applying matching techniques to analyze and evaluate the low-risk anomaly, arguably one of the most interesting puzzles in finance. Matching, especially CEM used in the study, allows us to control for features other than risk before forming portfolios.

Both sets of results largely confirm the existence of the low-risk anomaly in a more rigorous statistical setting. However, the reduced spread of cumulative returns between low

and high-risk portfolios after matching shows that at least a small portion of the low-risk anomaly can be explained by the fact that stocks in the original quintile 1 and quintile 5 portfolios did not share similar features (i.e. major industry group, company size, and trading volume). Interestingly, the magnitude of the anomaly was reduced more dramatically when beta was used as a measure of risk rather than volatility. The reason behind this phenomenon requires further research.

As mentioned, various matching techniques were considered, and CEM was selected because it gave us the portfolios with the best balance. However, we observed similar results when applying other matching techniques (including PSM and Mahalanobis distance matching).

We made two assumptions in the study, each of which could be relaxed to further explore the issue. The first assumption is the absence of the omitted variable bias. We matched stocks in quintile 1 with those in quintile 5, or vice versa, based on industry group, company size, and trading volume. We believe that these three variables best capture the most fundamental feature of a stock and the impact of adding other matching variables would be marginal at best.

We also assumed that volatility or beta of a stock is independent of its industry group, company size, and trading volume. By matching on the three variables, we assumed that the differences in both monthly return and cumulative return between the original and the matched portfolios are due to volatility or beta. Any potentially correlations between volatility or beta and any of the three variables are assumed to be zero.

Finally, the crossing-over of cumulative returns of the matched quintile 5 beta portfolio and the matched quintile 1 beta portfolio, which was not observed in previous studies, is interesting and worthy of further investigation. During the dot-com bubble when irrational

exuberance ruled the market, one would expect the effect of the anomaly be to greater in magnitude, yet the contrary was observed. This finding complicates the existing explanations of the anomaly. From the perspective of an efficient market, one has to explain why the effect of certain risk factors and cognitive biases were not as strong during that period. This short disappearance of the low-risk anomaly opens up questions about the overwhelming effect of certain time periods.

Overall, this study re-confirms the existence of the low-risk anomaly with added statistical rigor but raises new questions about the relationship between risk and return. Our research also serves a starting point for scholars to further explore the low-risk anomaly in extraordinary circumstances such as the dot-com bubble.

6 Appendix

Replication

Similar to Baker et al. (2011), we tracked stock returns from January 1968 to December 2008 using data from CRSP. For each month, we calculated the volatility and beta for each stock in the data set using up to 60-months of trailing data (the data set thus went back to January 1963). Volatility was calculated by taking the standard deviation of returns for each stock. Beta was calculated by regressing stock returns on excess market returns over the risk-free rate (market return minus the risk-free rate). Any stock that did not have at least 24-months of data within the 60-month trailing window was removed for that particular month. Market returns and risk-free rates were obtained from Ken French's data library at Dartmouth. For each of the 492 months in the study, stocks were sorted by their trailing volatility and assigned to one of five volatility portfolios (low to high) weighted by market capitalization. The same procedure was followed to create five beta portfolios (low to high). We ignored transaction costs when rebalancing the portfolios each month. We also conducted the same exercise while restricting the CRSP data to the top 1,000 stocks by market capitalization. The results from both studies appear in Figure 2.

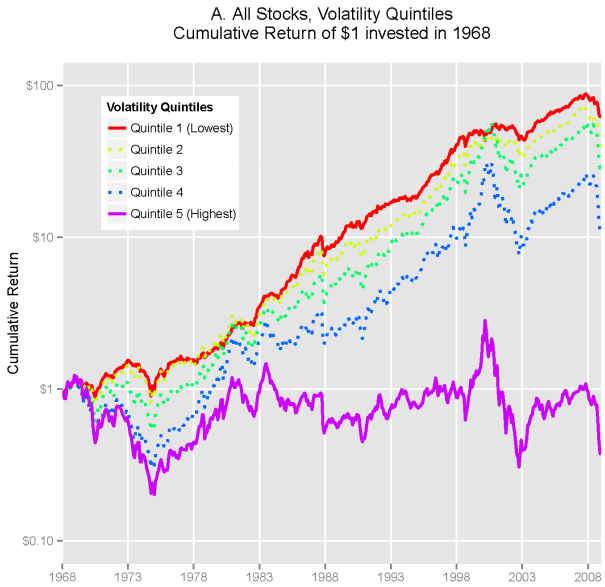

(c) Panel A

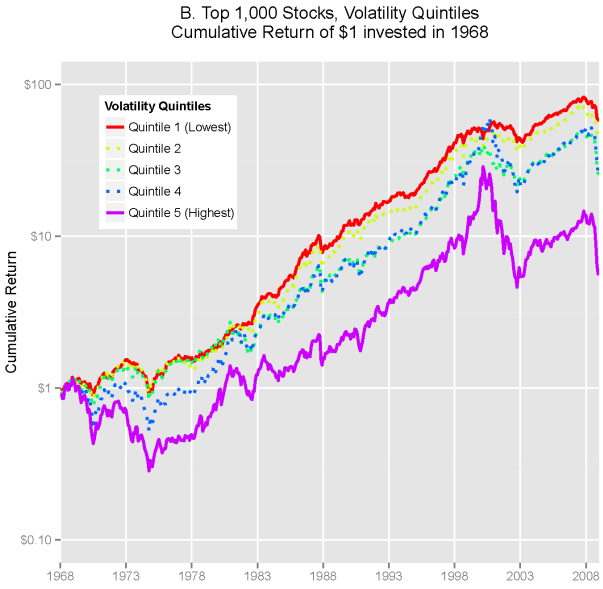

(d) Panel B

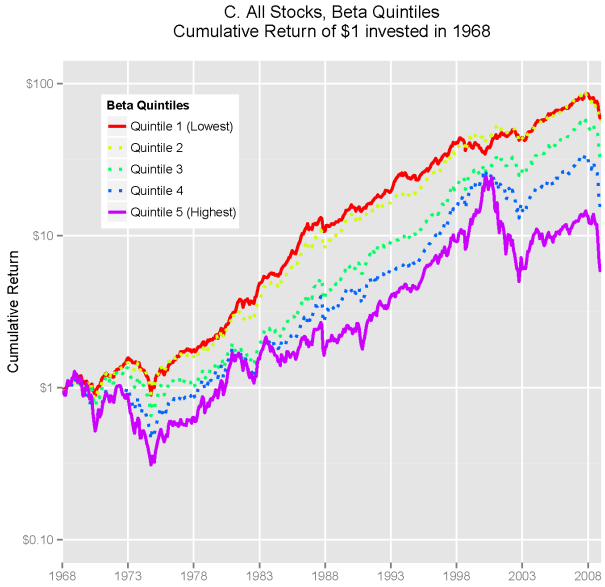

(e) Panel C

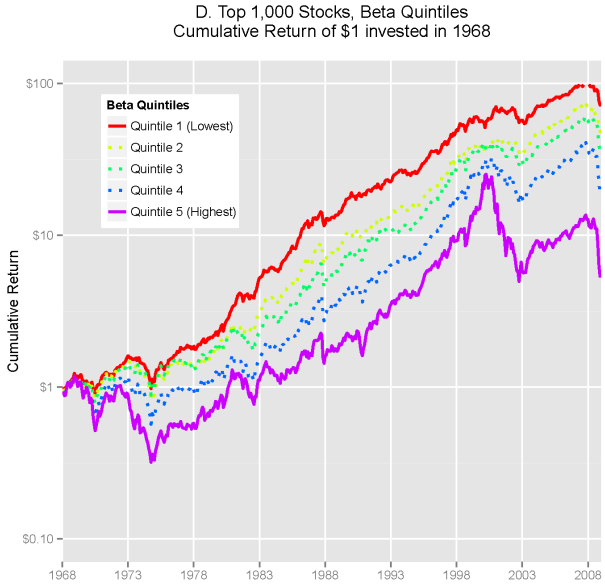

(f) Panel D

Figure 2: For each month, we sort all publicly traded stocks (Panels A and C) and the top 1,000 stocks by market capitalization (Panels B and D) tracked by CRSP (with at least 24 months of return history) into five equal quintiles according to trailing volatility (standard deviation) and beta. In January 1968, $1 is invested into each portfolio according to capitalization weights. We estimate volatility and beta by using up to 60 months of trailing returns (i.e. return data starting as early as January 1963). At the end of each month, we rebalance each portfolio, excluding all transaction costs.

	All Stocks					Top 1,000 Stocks				
	Low	Quintile 2	Quintile 3	Quintile 4	High	Low	Quintile 2	Quintile 3	Quintile 4	High
A.Volatility Sorts										
Cumulative value in 2008 of $1 invested in 1968	$62.72	$40.54	$29.32	$11.60	$0.40	$58.70	$45.02	$26.93	$27.44	$5.86
Geometric Average $R_p - R_f$	4.43%	3.35%	2.55%	0.26%	-7.99%	4.27%	3.61%	2.34%	2.39%	-1.40%
Average $R_p - R_f$	5.29%	4.79%	4.92%	4.02%	-2.74%	5.08%	4.79%	3.93%	4.66%	2.42%
Standard Deviation	13.09%	16.75%	21.44%	27.04%	31.86%	12.72%	15.26%	17.54%	21.00%	27.28%
Sharpe Ratio	0.40	0.29	0.23	0.15	-0.09	0.40	0.31	0.22	0.22	0.09
Average $R_p - R_m$	1.11%	0.61%	0.74%	-0.16%	-6.93%	0.90%	0.61%	-0.26%	0.48%	-1.76%
Tracking Error	7.02%	4.97%	8.25%	14.58%	20.37%	8.12%	6.08%	4.84%	8.35%	14.87%
Information Ratio	0.16	0.12	0.09	-0.01	-0.34	0.11	0.10	-0.05	0.06	-0.12
Beta	0.75	1.01	1.28	1.53	1.68	0.70	0.89	1.06	1.24	1.53
Alpha	2.17	0.57	-0.42	-2.36	-9.77	2.17	1.06	-0.52	-0.51	-3.98

	All Stocks					Top 1,000 Stocks				
	Low	Quintile 2	Quintile 3	Quintile 4	High	Low	Quintile 2	Quintile 3	Quintile 4	High
B.Beta Sorts										
Cumulative value in 2008 of $1 invested in 1968	$59.84	$57.14	$33.39	$16.14	$6.25	$72.61	$48.85	$36.12	$19.70	$5.68
Geometric Average $R_p - R_f$	4.32%	4.20%	2.87%	1.08%	-1.25%	4.80%	3.81%	3.07%	1.57%	-1.48%
Average $R_p - R_f$	5.06%	5.11%	4.24%	3.14%	2.54%	5.56%	4.82%	4.41%	3.46%	1.89%
Standard Deviation	12.20%	13.47%	16.39%	19.97%	27.16%	12.40%	14.15%	16.24%	19.18%	25.58%
Sharpe Ratio	0.42	0.38	0.26	0.16	0.09	0.45	0.34	0.27	0.18	0.07
Average $R_p - R_m$	0.88%	0.93%	0.06%	-1.05%	-1.64%	1.38%	0.64%	0.23%	-0.72%	-2.29%
Tracking Error	10.15%	7.40%	5.66%	6.32%	14.08%	9.65%	6.90%	5.09%	5.53%	12.34%
Information Ratio	0.09	0.13	0.01	-0.17	-0.12	0.14	0.09	0.04	-0.13	-0.19
Beta	0.60	0.76	0.97	1.21	1.56	0.63	0.81	0.97	1.17	1.49
Alpha	2.57	1.95	0.18	-1.92	-4.00	2.94	1.45	0.34	-1.43	-4.34

Table 3: For each month, we formed portfolios by sorting all publicly traded stocks (first five columns) and the top 1,000 stocks by market capitalization (second five columns) tracked by CRSP into five equal-sized quintiles according to trailing volatility (standard deviation) for Panel A and trailing beta for Panel B. We estimated volatility and beta by using up to 60 months of trailing returns (i.e., return data starting as early as January 1963). The return on the market, R_m, and the risk-free rate, R_f, are from Ken French's data library at Dartmouth. The information ratio uses the market return for the relevant universe, all stocks in the first five columns and the top 1,000 stocks in the last five columns. Average returns are monthly averages multiplied by 12. Standard deviation and tracking error are monthly standard deviations multiplied by the square root of 12.

Our substantive conclusions are in line with those of Baker et al. (2011) in that low volatility and low beta portfolios outperformed high volatility and high beta portfolios, respectively.

In Panel A of Figure 2, one dollar invested in the lowest volatility portfolio in January 1968 grows to $62.72 ($10.66 in real terms after adjusting for inflation) by December 2008. One dollar invested in the highest volatility portfolio during the same time period was worth $0.40 ($0.07 in real terms) by December 2008. The results are similar when only considering the 1,000 largest stocks by market capitalization as shown in Panel B of Figure 2. one dollar invested in the lowest volatility portfolio grows to $58.70 ($9.98 in real terms) by December 2008 as compared to $5.86 ($1.00 in real terms) for the highest volatility portfolio.

In Panel C of Figure 2, one dollar invested in the lowest beta portfolio in January 1968 grows to $59.84 ($10.17 in real terms) by December 2008. One dollar invested in the highest beta portfolio during the same time period was worth $6.25 ($1.06 in real term's) by December 2008. In Panel D of Figure 2, only the 1,000 largest stocks are considered. One dollar invested in the lowest beta portfolio grows to $72.61 ($12.34 in real terms) by December 2008 as compared to $5.68 ($0.97 in real terms) for the highest beta portfolio.

References

Acharya,Viral V., and Lasse Heje Pedersen. Asset pricing with liquidity risk. *Journal of Financial Economics*, 77 (2): 375-410, Aug. 2005 URL http://www.sciencedirect.com/science/article/pii/S0304405X05000334

Ang, Andrew, Robert J. Hodrick, Yuhang Xing, and Xiaoyan Zhang. The Cross-Section of Volatility and Expected Returns. *Journal of Finance*, 61 (1): 259-299, Feb. 2006

Ang, Andrew, Robert J. Hodrick, Yuhang Xing, and Xiaoyan Zhang. High Idiosyncratic Volatility and Low Returns: International and Further U.S. Evidence. *Journal of Financial Economics*, 91 (1): 1-23, Jan. 2009

Baker, Malcolm, Brendan Bradley, and Jeffrey Wurgler. Benchmarks as Limits to Arbitrage: Understanding the Low-Volatility Anomaly. *Financial Analysts Journal*, 67 (1): 40–54, Jan. – Feb. 2011.

Baker, Malcolm, Brendan Bradley, and Ryan Taliaferro. The Low Risk Anomaly: A Decomposition into Micro and Macro Effects. *SSRN* Jan 23, 2013. URL http://ssrn.com/abstract=2210003

Bernstein, Peter L., and Frank J. Fabozzi. Streetwise: The Best of the Journal of Portfolio Management. *Princeton University Press*, Jan. 19, 1998

Blitz, David C., and Pim van Vliet. The Volatility Effect: Lower Risk without Lower Return. *Journal of Portfolio Management*, 34 (1): 102-113, Fall 2007 URL http://ssrn.com/abstract=980865

Blitz, David C., Juan Pang, and Pim van Vliet. The Volatility Effect in Emerging Markets. *Emerging Markets Review*, forthcoming URL http://www.sciencedirect.com/science/article/pii/S1566014113000150

Fama, Eugene F., and Kenneth R. French. The Cross-Section of Expected Stock Returns. *Journal of Finance*, 47(2):427-465, June 1992

Ho, Daniel, Kosuke Imai, Gary King, and Elizabeth Stuart. Matching as Nonparametric Preprocessing for Reducing Model Dependence in Parametric Causal Inference. *Political Analysis*, 15: 199-236, 2007 URL http://j.mp/jPupwz

Iacus, Stefano M., Gary King, and Giuseppe Porro. Causal Inference Without Balance Checking: Coarsened Exact Matching. *Political Analysis*, Aug. 2011 URL http://j.mp/iUUwyH

Jensen, Michael C., Fischer Black, and Myron S. Scholes. The Capital Asset Pricing Model: Some Empirical Tests. *Michael C. Jensen, STUDIES IN THE THEORY OF CAPITAL MARKETS*, 1972 URL http://ssrn.com/abstract=908569

Li, Xi, Rodney Sullivan and Luis Garcia-Feijoo. The Limits to Arbitrage Revisited: The Low-Risk Anomaly. *Financial Analysts Journal* Forthcoming. URL http://ssrn.com/abstract=1738316

Rubin, Donald. The Use of Matched Sampling and Regression Adjustment to Remove Bias in Observational Studies. *Biometrics*, 29:185-203, 1973

Rubin, Donald. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 6:688-701, 1974

van Vliet, Pim, David C. Blitz, and Bart van der Grient. Is the Relation between Volatility and Expected Stock Returns Positive, Flat or Negative? *SSRN*, July 2011 URL http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1881503