How do Investors Measure Risk?

They use the CAPM.

Jonathan B. Berk[1] Jules H. van Binsbergen[2]

Abstract: We infer which risk model investors use by looking at their capital allocation decisions. We find that investors adjust for risk using the beta of the Capital Asset Pricing Model (CAPM). Extensions to the CAPM perform poorly, implying that they do not help explain how investors measure risk.

This article summarizes research originally published in Berk and van Binsbergen (2014).

There is no question more fundamental to the study of investments than the question of how investors measure risk. Indeed, one could argue that the study of investments only began when the first model of risk, the Capital Asset Pricing Model (CAPM), was developed in the early 1960's.[3] In the half century that has since elapsed, the ability of the model to accurately measure risk has been questioned. In response, a number of extensions to the original model have been proposed, and in some cases, adopted as improved measures of risk.[4] The principal empirical shortcoming that these extensions are designed to address is that much of the cross sectional difference in realized stock returns cannot be explained by cross sectional differences in the CAPM beta. The fact is that the relation between CAPM beta and return differences is weak in the full sample spanning 1926 to the present, and, importantly, sub periods of that sample exist, when the relation is appears to be absent altogether.

The implicit assumption underlying the literature that extends the CAPM is that a model that better explains cross sectional variation in returns necessarily better explains risk differences. But this assumption is problematic. To see why, consider the following analogy. Rather than look for an alternative theory, early astronomers reacted to the inability of the Ptolemaic theory to explain the motion of the planets by "fixing" each observational inconsistency. Just as modern financial economists added new risk factors, the early astronomers added epicycles to the theory. The net result was that by the time Copernicus proposed the correct theory that the Earth revolved around the Sun, the Ptolemaic theory had been fixed so many times it *better* explained the motion of the planets

[1]A. P. Giannini Professor of Finance, Graduate School of Business, Stanford University, Stanford, CA, 94305-7298 and National Bureau of Economic Research. Email: jbberk@stanford.edu Phone: 650-721-1280

[2]Associate Professor of Finance, The Wharton School, University of Pennsylvania, Philadelphia, PA, 19104 and National Bureau of Economic Research.

[3]This model was developed independently by Sharpe (1964), Lintner (1965), Mossin (1966) and Treynor (1961).

[4]The most notable are the factor specifications proposed by Fama and French (1993) and Carhart (1997) which are motivated by theoretical developments in Ross (1976) and Merton (1973).

than the Copernican system.[5] Similarly, although the extensions to the CAPM better explain the cross section of stock returns, it is hard to know, using traditional tests, whether these extensions represent true progress towards a better measure of risk or simply the asset pricing equivalent of an epicycle. To determine whether any extension to the CAPM better explains risk, one needs to confront the models with facts they were not designed to explain. That is the principal objective of this article.

To understand the basis of our new test, it is helpful to recall how prices and returns are determined in any risk model. All models of risk assume that investors compete with each other to find attractive investment opportunities. When investors find such opportunities, they react by submitting buy or sell orders and by doing so, the opportunity is removed. As a consequence of this competition, equilibrium prices are set so that the expected return of every asset is solely a function of its risk. Our key insight is that these buy and sell orders reveal the preferences of investors and therefore they reveal which risk model investors are using. By observing these orders we can infer whether investors price risk at all, and if so, which risk model they are using.

There are two criteria that are required to implement this idea. First, one needs a mechanism that identifies attractive investment opportunities. Second, one needs to observe investor reactions to these opportunities. We can satisfy both criteria if we implement the method using mutual fund data. Using this dataset we infer, from a set of candidate models, the model that is closest to the risk model investors are actually using.

Our results are surprising. We find that the CAPM is the closest model to the model investors use. None of the extensions that have been proposed better explain investor behavior. Importantly, the CAPM better explains investor behavior than no model at all, indicating that investors do price risk. Most surprisingly, the CAPM also outperforms a naive model in which investors ignore beta and simply chase any outperformance relative to the market portfolio. Investors' capital allocation decisions reveal that they adjust for risk using the CAPM beta. The poor performance of the extensions to the CAPM implies that although these extensions might better explain cross sectional variation in realized returns, they do not help explain how investors measure risk. In short, we are no closer to understanding the risk-return relation today than we were when the CAPM was originally developed more than half a century ago.

1 Methodology

In earlier work we explain how the mutual fund market equilibrates.[6] When investors find an investment in a particular mutual fund to be attractive, they invest capital in the fund. As the fund grows, the expected return of the fund declines as the fund manager's attractive investment ideas are exhausted. The flow of capital ceases when the expected return the mutual fund delivers to its investors is solely a function of the risk of the fund. That is, competition between investors drives the fund's net alpha to zero. What this implies is that the flows of capital in and out of mutual funds are the buy and sell orders mentioned in the introduction. Thus, the flow of funds reveals which investment opportunities mutual funds investors considered to be attractive.

Notice that when the market is in equilibrium, all mutual funds have a zero net alpha. Now consider what happens when new information arrives that allows investors to make a better inference about a fund's alpha. One example of new information is the fund's return itself. If the fund's return exceeds the risk adjusted return predicted by the risk model investors are using, investors will positively update their beliefs about the skill level of the fund's manager and infer that at the fund's current size, the alpha is positive. Similarly, if the fund's realized return is less than the risk adjusted return predicted by the risk model, investors will negatively update their beliefs about the skill level of the manager and infer that at the fund's current size, the alpha is negative. In short, the fund's realized return reveals attractive investment

[5] Copernicus incorrectly assumed that the planets followed circular orbits when in fact their orbits are ellipses.

[6] Berk and Green (2004), Berk (2005), Berk and van Binsbergen (2014) and Berk and van Binsbergen (2016).

opportunities, and the subsequent flow of funds reveals investor reactions to these opportunities.

We are now ready to describe our test. Each risk model we consider uniquely determines which funds outperform and which funds underperform. We then observe the subsequent flow of funds. The model for which outperformance best drives capital flows is the model that comes closest to the model that investors are actually using to price risk.

We implement this idea as follows. We compute the fraction of times we observe an inflow when the fund's realized return exceeds the risk adjusted return and the fraction of times we observe an outflow when the fund's realized return is less than the risk adjusted return, as defined by the risk model. Our measure of fit is the average of these two fractions. We show in Berk and van Binsbergen (2014) that this average can also be estimated by running a simple linear regression of the sign of flows on the sign of outperformance. The latter approach is preferable because, as we show in the same paper, the t-statistics of this regression is an accurate measure of statistical significance. In particular, if the coefficient using one risk model statistically significantly exceeds the coefficient using a second risk model, then we can say the first model is closer to the risk model investors are actually using.

2 Results

We use the mutual fund data set in Berk and van Binsbergen (2015). The data set spans the period from January 1977 to March 2011. We remove all funds with less than five years of data leaving 4275 funds.[7] Berk and van Binsbergen (2015) undertook an extensive data project to address several shortcomings in the CRSP database by combining it with Morningstar data, and we refer the reader to the data appendix of that paper for the details.

There are two practical issues that we need to confront in order to run this test. The first concerns what a flow actually is. A fund's assets under management changes for two reasons. Either the prices of the underlying stocks change or investors invest or withdraw capital. Although both mechanisms change assets under management, they are unlikely to equally affect the fund's alpha. For example, increases in fund sizes that result from inflation are unlikely to affect the alpha generating process. Similarly, the fund's alpha generating process is unlikely to be affected by changes in fund size that result from changes in the price level of the market as a whole. Consequently, in our empirical specification, we only consider capital flows into and out of funds net of what would have happened had investors not invested or withdrawn capital and had the fund manager adopted a purely passive strategy and invested in Vanguard index funds. That is, we measure the flow of funds as

$$\mathrm{SIGN}(q_{it} - q_{it-T}(1 + R_{it}^V)), \tag{1}$$

where q_{it} is the size of fund i at time t, and R_{it}^V is the cumulative return, over the horizon from $t-T$ to t, to investors of the collection of available Vanguard index funds that comes closest to matching the fund under consideration. Under this definition of capital flows, we are assuming that, in making their capital allocation decisions, investors take into account changes in the size of the fund that result from returns due to managerial outperformance alone. That said, all of our results are robust to replacing R_{it}^V with the funds own return in (1).

The second practical issue that we need to confront is the horizon length over which to measure the effects. For most of our sample funds report their AUMs monthly, however, in the early part of the sample many funds report their AUMs only quarterly. In order not to introduce a selection bias by dropping these funds, the shortest horizon we will consider is three months. If investors react to new information immediately, then flows should immediately respond to performance and the appropriate horizon to measure the effect would be the shortest horizon possible. But in reality, there is evidence that some investors do not respond immediately. For this reason, we also consider longer horizons (up to four years). The downside of using longer horizons is that longer horizons tend to put less weight on investors who update imme-

[7] We chose to remove these funds to ensure that incubation flows do not influence our results. Changing the criterion to two years does not change our results.

diately, and these investors are also the investors more likely to be marginal in setting prices.

We will consider the following models of risk. Because the market portfolio is not observable, we will test two versions of the CAPM that correspond to two different market proxies, the CRSP value weighted index of stocks and the S&P 500 index. We will also test the factor models proposed in Fama and French (1993), hereafter the FF factor specification and Carhart (1997), hereafter the FFC factor specification. In addition we will consider three "no model" benchmarks. The first uses the actual return of the fund, which corresponds to investors using no model at all. The second uses the return of the fund in excess of the risk free return. Investors would use this measure of risk if they were risk neutral. Finally, we will consider a model where the performance of the fund is just the fund's return minus the return of the market (as measured by the CRSP value weighted index). Although similar to the CAPM, in this model investors ignore beta. All they care about is outperformance relative to the market.

Which model best approximates the true asset pricing model? Table 1 reports the average of the fraction of times we observe an inflow when the fund's realized return exceeds the risk adjusted return and the fraction of times we observe an outflow when the fund's realized return is less than the risk adjusted return. If flows and outperformance are unrelated, we would expect this average to equal 50%. The first takeaway from Table 1 is that none of our candidate models can be rejected,[8] implying that regardless of the risk adjustment, a flow-performance relation exists. On the other hand, none of the models perform better than 64%. It appears that a large fraction of flows remain unexplained. Investors appear to be using other criteria to make a non-trivial fraction of their investment decisions.

Importantly, the CAPM with the CRSP value weighted index as the market proxy, performs best at all horizons. To assess whether the difference in performance between the CAPM and the other models is statistically significant, we report, in Table 2, the double-clustered (by fund and time) t-statistics. No model statistically outperforms the CAPM at any horizon.

To assess the relative performance of the models, we begin by first focusing on the behavioral model that investors just react to past returns without adjusting for risk, the column marked "Ret" in the table. By looking down that column in Table 2, one can see that the factor models all statistically significantly outperform this model at horizons of less than two years. For example, the t-statistic reported in Table 2 that the CAPM outperforms this no model benchmark at the 3-month horizon is 4.98, indicating that we can reject the hypothesis that the behavioral model is a better approximation of the true model than the CAPM. Based on these results, we can reject the hypothesis that investors just react to past returns. The next possibility is that investors are risk neutral. In an economy with risk-neutral investors, we would find that the excess return (the difference between the fund's return and the risk free rate) best explains flows, so the performance of this model can be assessed by looking at the columns labeled "Ex. ret." Notice that all the risk models nest this model, so to conclude that a risk model better approximates the true model, the risk model must statistically outperform this model. For horizons less than 2 years, all the risk models satisfy this criterion. Finally, one might hypothesize that investors benchmark their investments relative to the market portfolio alone, that is, they do not adjust for any risk differences (beta) between their investment and the market. The performance of this model is reported in the column labeled "Ex. mkt." The CAPM statistically significantly outperforms this model at all horizons — investors' actions reveal that they use betas to allocate resources.

Next, we use our method to discriminate between the risk models. Recall that both the FF and FFC factor specifications nest the CAPM (the first factor in each specification is the market), so to conclude that either factor model better approximates the true model, it must statistically significantly outperform the CAPM. The test of this hypothesis is in the columns labeled "CAPM." Neither factor model statistically outperforms the CAPM at

[8] The second column of Table 2 reports the double-clustered (by fund and time) t-statistics under the null that flows and performance are unrelated.

Model	Horizon					
	3-month	6-month	1-year	2-year	3-year	4-year
	Market models (CAPM)					
CRSP value weighted	**63.63**	**63.49**	**63.38**	**64.08**	**63.86**	**63.37**
S&P 500	62.52	62.26	61.61	62.20	61.40	60.92
	No model					
Return	58.55	59.77	57.72	59.76	60.83	61.20
Excess return	58.29	59.64	57.57	60.91	61.27	61.69
Return in excess of the market	62.08	61.99	61.19	62.45	62.05	61.76
	Multifactor models					
FF	63.14	62.84	63.05	63.62	63.59	62.43
FFC	63.25	62.92	63.09	63.59	63.46	62.35

Table 1: **Flow of funds outperformance relationship (1977-2011):** The table reports the average of the fraction of times we observe an inflow when the fund's realized return exceeds the risk adjusted return and the fraction of times we observe an outflow when the fund's realized return is less than the risk adjusted return. Each row corresponds to a different risk model. The first two rows report the results for the market model (CAPM) using the CRSP value-weighted index and the S&P 500 index as the market portfolio. The next three rows report the results of using as the benchmark return, three rules of thumb: (1) the fund's actual return, (2) the fund's return in excess of the risk-free rate, and (3) the fund's return in excess of the return on the market as measured by the CRSP value-weighted index. The next two rows are the FF and FFC factor specifications. The largest value in each column is shown in bold face.

Panel A: 3-Month horizon

Model	Prob.	Univ t-stat	CAPM	FFC	FF	CAPM SP500	Ex. mkt	Ret	Ex. ret
CAPM	63.63%	26.35	0.00	1.15	1.52	4.71	7.28	4.98	5.77
FFC	63.25%	28.64	-1.15	0.00	0.65	1.69	3.16	4.42	5.13
FF	63.14%	28.45	-1.52	-0.65	0.00	1.42	2.76	4.35	5.07
CAPM SP500	62.52%	21.25	-4.71	-1.69	-1.42	0.00	1.25	3.97	4.62
Excess market	62.08%	22.46	-7.28	-3.16	-2.76	-1.25	0.00	3.40	3.95
Return	58.55%	10.72	-4.98	-4.42	-4.35	-3.97	-3.40	0.00	1.18
Excess return	58.29%	10.11	-5.77	-5.13	-5.07	-4.62	-3.95	-1.18	0.00

Panel B: 6-Month horizon

Model	Prob	Univ t-stat	CAPM	FFC	FF	CAPM SP500	Ex mkt	Ret	Ex ret
CAPM	63.48%	21.11	0.00	1.08	1.23	3.24	4.64	2.63	3.17
FFC	62.92%	21.21	-1.08	0.00	0.35	0.95	1.47	2.21	2.64
FF	62.84%	22.40	-1.23	-0.35	0.00	0.79	1.38	2.09	2.49
CAPM SP500	62.26%	14.21	-3.24	-0.95	-0.79	0.00	0.50	1.78	2.09
Excess market	61.99%	16.03	-4.64	-1.47	-1.38	-0.50	0.00	1.47	1.73
Return	59.77%	8.44	-2.63	-2.21	-2.09	-1.78	-1.47	0.00	0.32
Excess return	59.64%	8.26	-3.17	-2.64	-2.49	-2.09	-1.73	-0.32	0.00

Panel C: 1-Year horizon

Model	Prob	Univ t-stat	CAPM	FFC	FF	CAPM SP500	Ex mkt	Ret	Ex ret
CAPM	63.38%	13.54	0.00	0.44	0.47	3.89	6.42	2.25	2.98
FFC	63.09%	14.30	-0.44	0.00	0.18	1.63	2.39	2.17	2.79
FF	63.05%	14.55	-0.47	-0.18	0.00	1.47	2.25	2.11	2.67
CAPM SP500	61.61%	8.31	-3.89	-1.63	-1.47	0.00	0.54	1.69	2.15
Excess market	61.18%	10.38	-6.42	-2.39	-2.25	-0.54	0.00	1.26	1.60
Return	57.72%	4.10	-2.25	-2.17	-2.11	-1.69	-1.26	0.00	0.17
Excess return	57.57%	4.00	-2.98	-2.79	-2.67	-2.15	-1.60	-0.17	0.00

Table continues on following page ...

Panel D: 2-Year horizon

Model	Prob	Univ t-stat	CAPM	FF	FFC	Ex mkt	CAPM SP500	Ex ret	Ret
CAPM	64.08%	12.80	0.00	0.80	0.97	5.73	3.81	1.45	1.42
FF	63.62%	16.17	-0.80	0.00	0.13	1.86	1.57	1.37	1.37
FFC	63.59%	16.46	-0.97	-0.13	0.00	2.06	1.72	1.31	1.33
Excess market	62.45%	10.89	-5.73	-1.86	-2.06	0.00	0.36	0.70	0.89
CAPM SP500	62.20%	8.16	-3.81	-1.57	-1.72	-0.36	0.00	0.60	0.84
Excess return	60.91%	7.09	-1.45	-1.37	-1.31	-0.70	-0.60	0.00	1.22
Return	59.76%	5.99	-1.42	-1.37	-1.33	-0.89	-0.84	-1.22	0.00

Panel E: 3-Year horizon

Model	Prob	Univ t-stat	CAPM	FF	FFC	Ex mkt	CAPM SP500	Ex ret	Ret
CAPM	63.85%	13.86	0.00	0.51	1.04	4.90	3.53	1.24	1.11
FF	63.59%	14.39	-0.51	0.00	0.43	2.54	2.41	1.21	1.09
FFC	63.46%	14.42	-1.04	-0.43	0.00	2.67	2.55	1.07	0.98
Excess market	62.05%	9.93	-4.90	-2.54	-2.67	0.00	0.84	0.37	0.46
CAPM SP500	61.40%	8.05	-3.53	-2.41	-2.55	-0.84	0.00	0.05	0.19
Excess return	61.27%	6.91	-1.24	-1.21	-1.07	-0.37	-0.05	0.00	0.51
Return	60.83%	5.85	-1.11	-1.09	-0.98	-0.46	-0.19	-0.51	0.00

Panel F: 4-Year horizon

Model	Prob	Univ t-stat	CAPM	FF	FFC	Ex mkt	Ex ret	Ret	CAPM SP500
CAPM	63.37%	13.02	0.00	1.81	1.95	4.76	0.79	0.90	3.93
FF	62.43%	11.77	-1.81	0.00	0.37	1.11	0.38	0.57	1.62
FFC	62.35%	11.61	-1.95	-0.37	0.00	0.96	0.32	0.50	1.58
Excess market	61.76%	9.70	-4.76	-1.11	-0.96	0.00	0.04	0.24	1.26
Excess return	61.69%	7.20	-0.79	-0.38	-0.32	-0.04	0.00	0.52	0.32
Return	61.20%	6.37	-0.90	-0.57	-0.50	-0.24	-0.52	0.00	0.11
CAPM SP500	60.92%	7.30	-3.93	-1.62	-1.58	-1.26	-0.32	-0.11	0.00

Table 2: **Tests of statistical significance:** The first column in the table reports the average of the fraction of times we observe an inflow when the fund's realized return exceeds the risk adjusted return and the fraction of times we observe an outflow when the fund's realized return is less than the risk adjusted return. The second column provides the t-statistic of the test of whether this average is significantly different from 50%. The rest of the columns provide the statistical significance of the pairwise test of whether the models are better approximations of the true asset pricing model. For each model in a column, the table displays the t-statistic of the test that the model in the row is a better approximation of the true asset pricing model. The rows (and columns) are ordered by the probabilities in the first column, with the best performing model on top. All t-statistics are double clustered by fund and time (see Thompson (2011)).

any horizon implying that the additional factors add no explanatory power for flows. Indeed, as Table 1 shows, at all horizons the CAPM actually outperforms all extensions to the model.

It is also informative to compare the tests of statistical significance across horizons. The ability to statistically discriminate between the models deteriorates as the horizon increases. This is what one would expect to observe if investors instantaneously moved capital in response to the information in realized returns. Thus, this evidence is consistent with the idea that capital does in fact move quickly to attractive investment opportunities.

3 Implication

The empirical finding that the CAPM does a poor job explaining cross-sectional variation in expected returns raises a number of possibilities about the relation between risk and return. The first possibility, and the one most often considered in the existing literature, is that this finding does not invalidate the neoclassical paradigm that requires expected returns to be a function solely of risk. Instead, it merely indicates that the CAPM is not the correct model of risk, and, more importantly, a better model of risk exists.

The second possibility is that the poor performance of the CAPM is a consequence of the fact that there is no relation between risk and return. That is, that expected returns are determined by non-risk based effects. The final possibility is that risk only partially explains expected returns, and that other, non-risk based factors, also explain expected returns. The results in this paper shed new light on the relative likelihood of these possibilities.

The fact that we find that the factor models all statistically significantly outperform our "no model" benchmarks implies that the second possibility is unlikely. That leaves the question of whether the failure of the CAPM to explain the cross section of expected stock returns results because a better model of risk exists, or because factors other than risk also explain expected returns. To conclude that a better risk model exists, one has to show that the part of the variation in asset returns not explained by the CAPM can be explained by variation in risk. This is what the flow of funds data allow us to do. If variation in asset returns that is not explained by the CAPM attracts flows, as is the case for the extensions of the CAPM we tested, then one can conclude that this variation is not compensation for risk.

4 Conclusion

The main contribution of this paper is a new way of testing the validity of a risk model. Instead of following common practice that relies on tests that use returns, we use mutual fund capital flow data. If the risk model under consideration correctly prices risk, then investors must be using it, and must be allocating their money based on that risk model. Consistent with this theory, we find that investors' capital flows in and out of mutual funds do reliably distinguish between asset pricing models. We find that the CAPM outperforms all extensions to the original model, which implies, given our current level of knowledge, that it is still the best method to use to adjust for risk. This observation is consistent with actual experience. Despite the empirical shortcomings of the CAPM, Graham and Harvey (2001) find that it is the dominant model used by corporations to make investment decisions.

We leave the question of what drives the fraction of flows that are unrelated to CAPM beta risk unanswered. A thorough investigation of what exactly drives these flows is likely to be highly informative about how risk is incorporated into asset prices.

References

Berk, J. B. (2005): "Five Myths of Active Portfolio Management," *Journal of Portfolio Management*, 31(3), 27–31.

Berk, J. B., and R. C. Green (2004): "Mutual Fund Flows and Performance in Rational Markets," *Journal of Political Economy*, 112(6), 1269–1295.

Berk, J. B., and J. H. van Binsbergen (2014): "Assessing Asset Pricing Models Using Revealed Preference," *Journal of Financial Economics*, Forthcomming.

——— (2015): "Measuring skill in the mutual fund industry," *Journal of Financial Economics*, 118(1), 1 – 20.

——— (2016): "Active Managers are Skilled," *Journal of Portfolio Management*, Forthcoming.

Breeden, D. T. (1979): "An intertemporal asset pricing model with stochastic consumption and investment opportunities," *Journal of Financial Economics*, 7(3), 265 – 296.

Carhart, M. M. (1997): "On Persistence in Mutual Fund Performance," *Journal of Finance*, 52, 57–82.

Fama, E. F., and K. R. French (1993): "Common Risk Factors in the Returns on Stocks and Bonds," *Journal of Financial Economics*, 33(1), 3–56.

Graham, J. R., and C. R. Harvey (2001): "The Theory and Practice of Corporate Finance: Evidence from the Field," *Journal of Financial Economics*, 60, 187–243.

Lintner, J. (1965): "The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets," *The Review of Economics and Statistics*, 47(1), 13–37.

Merton, R. C. (1973): "Optimum Consumption and Portfolio Rules in a Continuous-Time Model," *Journal of Economic Theory*, 3, 373–413.

Mossin, J. (1966): "Equilibrium in a Capital Asset Market," *Econometrica*, 34(4), 768–783.

Petersen, M. A. (2009): "Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches," *Review of Financial Studies*, 22(1), 435–480.

Ross, S. A. (1976): "The arbitrage theory of capital asset pricing," *Journal of Economic Theory*, 13(3), 341–360.

Sharpe, W. F. (1964): "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk," *Journal of Finance*, 19(3), 425–442.

Thompson, S. B. (2011): "Simple formulas for standard errors that cluster by both firm and time," *Journal of Financial Economics*, 99(1), 1–10.

Treynor, J. (1961): "Toward a Theory of the Market Value of Risky Assets," .