


Why We Need a Guarantor of Last Resort

By Kathryn Judge



Guarantor of Last Resort: Is There a Better Alternative?

By Morgan Ricks



Emergency Guarantee Authority: Not Letting a Crisis Go to Waste

By Graham Steele

Editor-At-Large
Reynolds Holding

THE CLS BLUE SKY BLOG

COLUMBIA LAW SCHOOL'S BLOG ON CORPORATIONS AND THE CAPITAL MARKETS

Editorial Board
John C. Coffee, Jr.
Edward F. Greene
Kathryn Judge

Our
Contributors

Corporate
Governance

Finance &
Economics

M & A


Securities
Regulation

Dodd-Frank

International
Developments

Library &
Archives

What Do Mutual Fund Investors Really Care About?



By Itzhak Ben-David, Jiacui Li, Andrea Rossi and Yang Song April 5, 2019

Do investors behave rationally? Many researchers examine actively managed mutual funds to answer this question. The advantage of doing so setting is that we can observe the past performance of fund managers as well as the capital allocation decisions of

investors, both at relatively high frequency.

Yet, despite three decades of research on mutual funds, whether investors in mutual funds display rational behavior remains unsettled. The empirical facts are indisputable: While fund managers have a hard time beating their benchmarks and their performance is not consistent over time, investors continue to chase past returns. In a seminal study, Berk and Green (2004) propose that these facts are consistent with a model of rational investors looking for managers with high skill (i.e., chase past performance). However, because fund managers' investment strategies are not scalable, investor flows that go to the best performing funds compete profits away, thus leading to the observed lack of consistency. Because this model describes an equilibrium outcome, it is hard to test empirically and thus to refute.

An important assumption of the Berk and Green (2004) model is that investors rationally evaluate managerial skill and allocate capital accordingly. As such, they should be able to distinguish between performance that is a result of managers' investment skill and performance that is a result of exposure to systematic risk factors (e.g., market risk). Two recent celebrated studies (Barber, Huang, and Odean 2016 (BHO) and Berk and van Binsbergen 2016 (BvB)) use different empirical techniques to assess whether investors adjust past performance for risk. Both studies reach the same conclusion: Among the asset pricing models tested, investors appear to use the Capital Asset Pricing Model (CAPM) for benchmarking the performance of mutual fund managers. This means that investors understand the exposure of funds to market risk and assess managers' skill by adjusting their past performance for this exposure. While BHO interpret this result as indicating that retail investors are unsophisticated (they do not account for other known risk factors), BvB draw a broader and more general conclusion – that the CAPM is the “closest to the asset pricing model investors are actually using” (p. 2).

The results of BvB and BHO are so striking that they deserve a second look. In particular, the picture that emerges from the two studies is of investors who understand market risk and can calculate exposure to risk (beta). This is at odds with the profile of mutual fund investors (retail investors and households) and with the limited information that is readily available to them. Furthermore, these results are in sharp contrast with the dozens of studies showing that mutual investors behave in a naive fashion – e.g., they respond to simple signals, invest in high-fee funds, and time the market poorly. These results are also at odds with a battery of studies finding that investors respond to attention-grabbing and easy-to-process signals, such as external rankings (e.g., Morningstar, Wall Street Journal).

Motivated by the fact that households hold the vast majority of mutual fund assets, we test whether simple and readily available signals explain investor behavior better than asset pricing models. Specifically, we test whether Morningstar's star ratings as well as recent unadjusted returns explain mutual fund flows better than risk-adjusted returns. Morningstar ratings are the ideal candidate for our tests for several reasons. First, Morningstar is the leader of the U.S. fund rating industry, and its star ratings are often provided to investors by financial advisors, brokers, defined-contribution retirement plan sponsors, and fund companies themselves in their marketing materials. The ratings are also available for free on Morningstar's website. Second, Morningstar ratings do not adjust for fund exposure to any systematic risk factor. Third, these ratings are available for most U.S. equity mutual funds.

Our results show that ratings are the main determinant of capital allocation across mutual funds, followed by recent past returns. We find no evidence that investors allocate capital to funds based on the exposure to the market factor or to any other risk factor. Furthermore, while investors follow Morningstar

religiously, they do not appear to understand Morningstar's rating methodology. In sum, our results show that mutual fund retail investors do not have a well-defined benchmarking standard for evaluating fund managers.

In the first part of our study, we examine the ability of Morningstar to explain flows. We begin with raw statistics. Just by eyeballing the capital allocations over time (see figure below), one can easily observe that Morningstar ratings drive flows significantly better than any of the model-based alphas in any single year. Over the sample period, 5-star Morningstar funds (7.4 percent of fund-months) attracted a total of \$656 billion. Conversely, during these overweighted periods, the funds that performed best according to commonly used factor models attracted only \$274 billion to \$326 billion.

To evaluate more rigorously whether Morningstar ratings explain capital allocation better than other benchmarks based on asset pricing models, we adopt the diagnostic test proposed by BvB. This test measures the degree of agreement between the direction of net fund flows and various signals (e.g., the sign of a fund's alpha using different asset pricing models, or Morningstar ratings, in our case). We first replicate BvB's main finding. Consistent with their results, the sign of alphas from common asset pricing models agrees with the sign of fund flows in the narrow band of 57.0 percent to 59.9 percent of the time, and the CAPM dominates other models by a small margin (60.6 percent). Morningstar ratings, in contrast, predict the direction of flows materially better (up to 67.9 percent of the time).

The BvB test is only based on signs. To further sharpen the test, we also analyze the magnitude of the spread between flows to top and bottom funds ranked according to various asset pricing models or Morningstar ratings. Overall, ratings decisively outperform all asset pricing models considered. Moreover, when using either raw dollar flows or flows as a fraction of total net assets, rankings based on the CAPM no longer outperform those based on unadjusted fund returns in explaining flows.

Next, we look in depth at BHO's methodology and results. BHO decompose fund returns into components associated with a host of commonly used risk factors and an alpha. They then regress fund flows on different components of fund returns, controlling for time fixed effects. BHO find that while fund flows respond to all return components, flows respond less strongly to returns originating from exposure to the market factor. They conclude that investors care about market risk and therefore discount returns that originate from exposure to market risk.

Our analysis indicates that BHO's findings should be interpreted in a different way. Specifically, BHO's conclusion is based on a panel regression with time fixed effects. While this method is standard, in this specific case it overweights cross-sections with extreme market returns because the dispersion in the independent variable of interest – the market-related component of fund returns – is particularly high in those periods. Importantly, during the same periods, fund flows are significantly less responsive to fund performance, an empirical fact documented by earlier literature. As a result, even if investors do not differentiate between the sources of fund returns, the panel regression of BHO would still convey the impression that flows respond less to the market-related component of fund returns.

We perform the following test to show that the data, including BHO's evidence, is entirely consistent with fund flows not differentiating between return components. Specifically, we generate counterfactual flows that respond equally to all components of fund returns. Using these counterfactual flows in the BHO empirical framework generates results that are indistinguishable from the regression results based on the original data. In other words, flows appear to be less sensitive to fund returns attributable to market exposure, even though by construction the flows in this analysis are constrained to respond to all return components equally. Not only does the result using counterfactual data match BHO's "market-discounting" evidence, it also matches the degree of "discounting" of other factors found in the BHO test. Therefore, we conclude that BHO's findings are consistent with investors chasing unadjusted returns and not discounting any factor-related return component.

How do we interpret our finding that investors rely so heavily on Morningstar ratings? One may wonder whether investors do care about risk, but outsource risk adjustments to Morningstar. Morningstar ratings account for total fund return volatility and size and value style benchmarks. Hence, it is possible that investors rely on Morningstar ratings as a simple way to adjust for total fund return volatility, or to indirectly adjust for exposure to size and value risk factors.

We show that this is unlikely. First, we consider that flows appear to penalize funds with higher volatility. We show that flows are negatively correlated only with the part of variation in volatility that is related to having a different Morningstar rating. This represents only 3.6 percent of the dispersion in return volatility; investors do not adjust for the remaining 96.4 percent. Therefore, if investors indeed care about volatility, they should realize that relying on Morningstar ratings is a very ineffective way to account for volatility.

Second, we consider whether investors rely on Morningstar as a way to adjust for size and value exposure in order to assist in benchmarking managers' performance. To test this, we use the fact that Morningstar started ranking funds within size and value style categories only in June 2002. Before then, all U.S. equity funds were ranked together. If investors use Morningstar for style adjustment, then flows should have responded less to the ratings before the implementation. We do not find empirical support for this supposition. In fact, investors relied heavily on the ratings both before and after the methodology change. Furthermore, an event study around the transition date reveals that before Morningstar implemented the style benchmarking, investors not only failed to adjust for style exposure, but actively chased funds whose ratings were high due to style-related returns rather than because of managerial ability.

In summary, we find no evidence that retail mutual fund investors use the CAPM, or any other of the commonly used factor models, to allocate capital to mutual funds. Rather, they naively rely on external rankings and chase past winners. Our results suggest that the critical assumption of the Berk and Green (2004) model – that investors allocate capital across funds in an optimal and rational fashion – is not likely to hold in the data. It is hard to imagine that

investors who fail to distinguish between skill and risk factor exposure would be able to rationally learn about managerial skill and allocate capital optimally, leading to an equilibrium in which the expected adjusted return of all funds equals zero.

REFERENCES

Barber, Brad M., Xing Huang, and Terrance Odean, 2016, Which factors matter to investors? Evidence from mutual fund flows, *Review of Financial Studies* 29, 2600–2642.

Berk, Jonathan B., and Richard C. Green, 2004, Mutual fund flows and performance in rational markets, *Journal of Political Economy* 112, 1269–1295.

Berk, Jonathan B., and Jules H. van Binsbergen, 2016, Assessing asset pricing models using revealed preference, *Journal of Financial Economics* 119, 1–23.

This post comes to us from Professor Itzhak Ben-David at The Ohio State University and NBER; Jiacui Li, a PhD candidate at Stanford University; Professor Andrea Rossi at the University of Arizona; and Professor Yang Song at the University of Washington. It is based on their recent paper, “What Do Mutual Fund Investors Really Care About?” available [here](#).