Journal of Financial Economics 108 (2013) 506–528

Contents lists available at SciVerse ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/jfec

Mutual fund risk and market share-adjusted fund flows ☆

Matthew Spiegel [a,*], Hong Zhang [b,c,1]

[a] Yale School of Management, P.O. Box 208200, New Haven CT 06520, USA
[b] CKGSB, 20/F Tower E2 Oriental Plaza, 1 East Chang An Avenue, Beijing 100738, China
[c] INSEAD, 1 Ayer Rajah Avenue, Singapore 138676, Singapore

ARTICLE INFO

Article history:
Received 28 September 2010
Received in revised form
19 April 2012
Accepted 18 May 2012
Available online 20 December 2012

JEL classification:
G11
G1

Keywords:
Mutual funds
Fund flows
Investment flows
Convexity
Market share

ABSTRACT

Several papers use a fractional specification (net inflow/ assets under management) to infer a convex relation between flow and past performance. However, heterogeneous linear response functions combined with the pooled analysis commonly used in these studies can yield false convexity estimates. We show that such heterogeneity obtains in practice. Along these same lines, the paper also finds that several previously unexamined implications of a convex flow-performance relation fail to hold. Moreover, convexity with fractional flows (which we confirm) largely disappears in a conditional analysis that controls for heterogeneity. Market shares offer an alternative specification for flow that is more resilient to heterogeneity. Using this alternative specification, we again find no evidence of convexity in the flow-performance relation. We conclude that the widely held belief that the flow response function is convex is due solely to misspecification of the empirical model. The flow-return relation is linear.

1. Introduction

Numerous studies show a convex relation between mutual fund flows and past returns, including those by Chevalier and Ellison (1997), Sirri and Tufano (1998), Fant and O'Neal (2000), and Huang, Wei, and Yan (2007). These findings have since been used as a foundation for theoretical models of inter-fund competition (Carpenter, 2000; Lynch and Musto, 2003; Taylor, 2003; and Basak and Makarov, 2012). Convexity plus Jensen's inequality naturally leads these articles to conclude that increasing a fund's risk increases its expected capital inflows. While these papers differ in many ways, they share a common element: Returns are assumed to map directly into flows divided by assets under management (AUM). While this fractional flow specification may appear innocuous, it is not. If these models are properly specified, then aggregate flows should be linked to the cross-sectional distribution of fund returns. Our tests, however, indicate otherwise, suggesting that the standard fractional flow model is misspecified. Market shares offer an alternative, more resilient specification and show no evidence of convexity in the fund growth-performance relation. As the paper demonstrates, misspecification can fully account for the previously empirically documented convex flow-return relation. In reality, the relation appears to be linear.

A simple example shows how the fractional flow model links the distribution of individual fund returns to aggregate flows. Consider an economy with two funds: one has $100 under management, the other, $10. The fractional flow

☆ We would like to thank seminar participants at Columbia University, the University of Texas at Austin, University of Western Ontario, the Fourth Singapore International Conference on Finance, and Vanderbilt University for helpful comments. Hong Zhang thanks the INSEAD Alumni Fund for financial support.

* Corresponding author. Tel.: +1 203 981 1184.
E-mail addresses: matthew.spiegel@yale.edu (M. Spiegel), hong.zhang@insead.edu (H. Zhang).
[1] Tel.: +65 97879206.

0304-405X/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jfineco.2012.05.018

model states that the better performing fund will see an inflow of 10%, the other, 0%. If the large fund does better, the aggregate flow equals \$10, and if the small one does better, the aggregate flow is \$1. This relation between fund returns across different size groups and aggregate flows is, in fact, a general implication of the standard fractional flow model: When large funds do relatively well, aggregate flows should be larger. However, our tests yield little evidence that this is the case. Aggregate flows are seemingly determined by economy-wide events, such as the overall market return, not by whether large or small funds have recently done relatively well. A model based on market share changes provides a simple way around this problem. By construction, market share changes add to zero. Thus, in the absence of additional specification, a market share model does not link aggregate flows to the distribution of individual fund returns.

Any empirical specification is necessarily misspecified to some degree. What is important is its robustness to such errors—in the case of mutual funds, an ability to handle unaccounted for heterogeneity in the cross section and across time. In practice, this can be difficult to accomplish with a fractional flow-return model as the (nonexhaustive) example depicted in Fig. 1 illustrates.

In Fig. 1, the dashed blue lines represent the relation between period t flows (f_t) and period $t-1$ returns (r_{t-1}) within a period for two different fund types. They are labeled hot and cold money funds. The picture has three important elements. First, heterogeneity exists in the cross section of functions mapping flows to returns, which the empirical model has not fully controlled for. (In this example, the researcher does not estimate separate functions for the hot and cold money funds.) Second, when a fund's inflows are poised to be relatively high (meaning a relatively high cross on the y-axis), the flow-return relation steepens. Third, hot money funds have more volatile returns than cold money funds. As a result, the cold funds have relatively more observations close to the y-axis and the hot ones relatively more further away, as the red ovals indicate. Importantly, in all cases the relation between flows and returns is linear, implying volatility will not

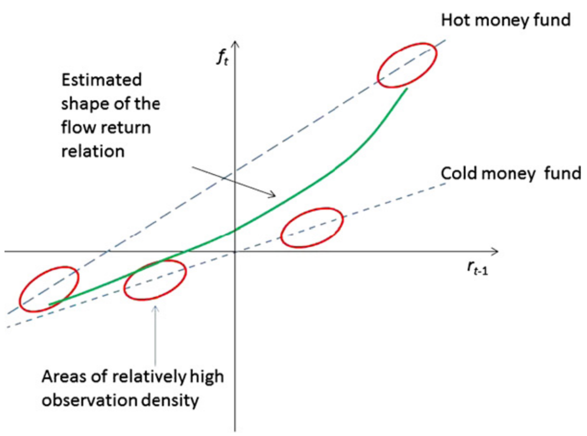

Fig. 1. Spuriously estimated convexity example. Fractional flows are graphed against relative returns.

affect a fund's expected asset accumulation. Nevertheless, a regression with fractional flows as the dependent variable and returns as the explanatory variable will produce the convex green line. The convexity is entirely due to misspecification error. In this case, the curve tries to minimize the sum of squared errors by moving closer to the relatively densely populated areas, i.e., closer to the red ovals.

To see how Fig. 1 might come about in actual data, consider a scenario in which one fund is growing (hot) while another has essentially stabilized at its current asset level (cold). To fix the general idea, the stable fund might be old, large, and housed within a well-established mutual fund family. While the cold fund's many long-term investors might not look to it as the place to put their money when seeking out the latest and greatest, neither are they likely to pull out if recent returns are subpar. This could arise from simple inertia (Choi, Laibson, Madrian, and Metrick, 2002, 2004a, 2004b; Duflo and Saez, 2002; and Choi, Laibson and Madrian, 2009) or if the fund's assets derive largely from essentially automated 401(k) retirement flows (Choi, Laibson, Madrian, and Metrick, 2002, 2004b; and Mitchell, Mottola, Utkus, and Yamaguchi, 2006). These and other possible factors produce for the older (cold money) fund the flow-return relation described by the lower blue short dashed line. The high long dashed blue line in Fig. 1 can be thought of as potentially representing young (hot money) funds with relatively few assets under management. Due to their size, these hot money funds can easily hold portfolios that, compared with large cold money ones, are relatively undiversified. Consequently, the hot money funds have a higher return volatility than the cold money funds. The important point is that the red ovals, indicating where data for the cold money funds are relatively dense, lie closer to the origin than those for the hot money funds.

Finally, consider the impact of the age of the hot and cold money funds. If a hot money fund's investors have been with it for only a short period of time, its flow will likely vary dramatically in response to its recent returns, at least relative to the cold money fund. Combining the impact of heterogeneity in the unconditional flow, return volatility, and investor responsiveness to past returns yields the flow-return relation displayed in Fig. 1. A typical flow-return model will now produce the convex solid green line. Can controls fix the problem? Yes, if one knows how to properly divide up the funds.

In this paper, we show that separating out funds that are both young and small from the rest yields a pair of linear relations. These young-small funds have all the relative properties depicted in Fig. 1: volatile returns, high relative unconditional growth, and a particularly steep flow-return relation. Thus, even accounting for just this one source of cross-sectional heterogeneity, the convexity found by running a pooled regression largely disappears.

While we show that the empirical problems that lead to false convexity can be mitigated, it requires knowing exante what cross-sectional controls are needed. But any number of scenarios also yield the pattern depicted in Fig. 1, so there is no guarantee that any single set will be sufficient. Consider that the relation depicted in Fig. 1 could arise dynamically. High aggregate market returns