WORKING PAPER SERIES

Lucia Milone, Paolo Pellizzari

Mutual funds flows and the
'Sheriff of Nottingham' effect

Working Paper n. 188/2009
June 2009

ISSN: 1828-6887

This Working Paper is published under the auspices of the Department of Applied Mathematics of the Ca' Foscari University of Venice. Opinions expressed herein are those of the authors and not those of the Department. The Working Paper series is designed to divulge preliminary or incomplete work, circulated to favour discussion and comments. Citation of this paper should consider its provisional nature.

Mutual funds flows and the “Sheriff of Nottingham” effect*

Lucia Milone
Dept. of Applied Mathematics
Dorsuduro 3285/E
Venice - Italy

Paolo Pellizzari
Dept. of Applied Mathematics
and SSAV
Venice - Italy

June 2009

Abstract

Investors in mutual funds appear to reward disproportionately the best performing funds with large inflows while, at the same time, avoid to withdraw similar amounts from the poorly managed funds. We show that this peculiar flat-convex shape of the flow-performance curve for mutual funds can be generally explained by a model where profit chasing customers punish the bad funds by switching a fraction of their wealth to the best ones (“Sheriff of Nottingham” effect). In the absence of external flows, the model provably produces a constant curve when the standard deviation of excess returns is much larger than the level of the returns. This for the most part explains the apparent insensitivity of flows to below-average returns. The introduction of exogenous injections of money invested in the top funds completes the model and provides a realistic increase in the flows of the funds yielding above-average returns. We finally show by simulation that our results are robust to variations in the values of the parameters of the model.

Keywords: Funds’ flows, flow-performance curve, agent based models.

JEL Classification Numbers: G11, G23, C63.

1 Introduction

Flows into and out of mutual funds have an interesting dynamics and exhibit some puzzling features. We focus in this paper on the shape of the flow-performance curve that shows an asymmetric reaction of households to past returns. Customers appear to chase (excess) profits, “flocking to funds

* Partially funded by MIUR under grant 2007TKLTSR, “Computational markets design and agent-based models of trading behavior”. Corresponding author: paolop@unive.it.

with the highest recent returns, though failing to flee from poor performers”, [Sirri and Tufano, 1998]. This results in flat-convex (*J*-shaped) curve suggesting that customers are pretty insensitive to low and average performances and appears to care about extremely positive returns only. In other words, the (positive) inflow into the best performing funds has a much bigger magnitude than the (negative) outflow from the poorly managed funds. More subtly, in the intermediate range of returns a rather good (but not exceptional) fund is likely to experience the same flow pattern of a rather bad (but not awful) fund, even though the difference in the (properly adjusted) performance is quite substantial.

The empirical evidence robustly supporting this finding is abundant and some recent papers offer possible explanations that go beyond the high search costs and lack of sophistication on the part of clients that were initially proposed. [Berk and Tonks, 2007] observe that there is a mechanical reason partly explaining the puzzle: only people who have invested in the worst funds can withdraw money from them while all investors have the option to invest in the best funds, producing a higher positive growth rate. In second place, they observe that the apparent smaller sensitivity to bad performance is due to the heterogeneous willingness to realize losses: only a part of the investors in the fund escape away after one poor yearly return but others are less responsive to past returns and funds with two consecutive negative returns have small outflows. The work in [Lynch and Musto, 2003] supports the view that some apparent indifference to bad returns is linked to the ambiguous inference that can be drawn from a negative outcome: from the one hand, it can be interpreted as a signal of poor managerial skills and disappointing future performances. Hence, investors can switch their capital away. On the other hand, poor performance is also leading to a change in the strategy of the manager that (exactly because of the past failures) wants to modify his trades and/or risk exposure. In this respect, inferior returns are not indicative of future prospects and clients may hold their shares in the funds.

[Del Guercio and Tkac, 2008] consider the 5-star Morningstar rating as a performance measure and show that most of the features of the flow dynamics are not strictly dependent on return-based rankings: a 5-star rating induces a disproportionate abnormal inflows, while other changes are less relevant (but still economically significant at times).

In this work, we describe a model of internal and exogenous flows among mutual funds that is driven by the simple principle “take from the poor (performer) to give to rich (one)”. This deliberate negation of the famous Robin Hood’s catchphrase suggested to name this behaviour the “Sheriff of Nottingham” effect. All the explanations provided so far in the literature indeed specialize in some clever way a general kind of behaviour where customers move part of their wealth away from the worst performers to invest in the best funds. The specific details or the motivation for this partial shift may

change and there have been various suggestions in the aforementioned literature but we offer some theoretical insight and extensive simulation to argue that any “Sheriff of Nottingham”-like behaviour can produce a remarkably flat flow-performance curve.

Our work does not fully model all the actors in the fund industry, namely customers and funds. In fact, we leave aside the customers to model a given number N of funds that experience flows due to the action of investors. If a specific fund ranks in the worst x funds, based on some performance measure in a reference period, then a fraction α of its net asset value (NAV) is withdrawn and moved (randomly) to one of the best x performing funds. The asymmetric reaction to bad performances can be easily justified from a psychological point as the disutility of a loss is larger than the utility of a similar gain, [Kahneman and Tversky, 1979]. Moreover, the idea that there is a large blurred intermediate area where all performances look the same is somewhat intrinsic in the structure of the basic information that is presented to funds’ customers. The ratings provided by Morningstar and other vendors, for example, assign stars on the basis of a subdivision of funds into “quintiles” of performances, somewhat suggesting that there are worse and better funds but also giving a rather coarse and indistinct picture.

The change in the wealth of the funds can be interpreted as the result of many individual agents switching to other and possibly more skilled managers. We do not keep track of the single investors but only study the funds and how their growth rates are changed by this behaviour of the agents. In the presence of internal flows only (i.e., if there are not external flows pouring into the mutual funds), the flow-performance curve turns out to be very flat for a wide set of parameters, provided that the standard deviation of the returns is relatively large (as it is the case in reality).

In order to capture the excessive growth rate of the best performing funds, we additionally assume that customers with fresh wealth invest among the best mutual funds. In detail, we assume that only the best x funds receive net inflows, amounting to a fraction β of their NAV. Our general model is able to robustly generate a flat flow curve for average and below-average funds and much larger growth rates for the top performing products in the industry. Putting things together, we reproduce the empirical observations with a general process, few principles and repeated interactions in terms of flow exchange and injections among the funds (leaving behind the curtains the individual investors that ultimately drive the whole dynamics).

The paper goes on as follows. Section 2 presents our ideas in a simplified probabilistic and tractable framework. We formally prove that flat curves can be obtained in the limit when the standard deviation of returns is large. Section 3 describes our model, gives the details of the switching mechanism and presents results using simulations. We conclude in Section 4.

2 A simple example and one analytical result

This section presents a simple example of rank-based flows among funds. Assume that households can invest in four mutual funds yielding excess returns (or alphas) distributed as normal independent random variables $X_{1t}, X_{2t}, X_{3t}, X_{4t}$ in year t . Letting $X_{jt} \sim N(\mu_j, \sigma)$, $\mu_1 \leq \mu_2 \leq \mu_3 \leq \mu_4$, we can interpret the mean of the excess return as the skill of the manager, so that the fourth manager is on average the best one. The clients, however, do not observe the μ_j s but only the yearly realizations of the four funds. For concreteness, the following figures are representatives of the actual universe of mutual funds split in quartiles, see [Kosowski et al., 2006]: $\mu_1 = 4.25\%$, $\mu_2 = 5.40\%$, $\mu_3 = 6.48\%$, $\mu_4 = 8.40\%$, with a common standard deviation $\sigma = 14.81\%$.

How would you allocate wealth in a similar situation? Waiting for a few decades to estimate the means with high precision does not seem realistic and one reasonable response would be a “Sheriff of Nottingham” behaviour that periodically takes from the poor fund to transfer to the rich one¹. This section shows how the flows between funds are affected by this commonsensical switching strategy.

Clients of the funds myopically try to optimize their excess return ranking the funds every year on the base of the realized $X_{1t}, X_{2t}, X_{3t}, X_{4t}$ and moving some wealth from the worst to the best performer. The precise details are described in the following section but the important thing to realize is that the probability of an in(out)-flow for a specific fund is depending on the probability to be ranked first or last in a year. Let $abcd$ be a permutation (ranking) of the set $\{1, 2, 3, 4\}$. We are interested by the probability P_{abcd} , namely

$$P_{abcd} = Pr(X_a \leq X_b \leq X_c \leq X_d),$$

where we omit t for notational simplicity. Defining $Y_{ij} = X_i - X_j$, $i \neq j \in \{1, 2, 3, 4\}$, the above probability can be rewritten as $P_{abcd} = Pr(Y_{ab} \leq 0, Y_{bc} \leq 0, Y_{cd} \leq 0)$. Each Y is the difference of independent normal variates and we have $E[Y_{ij}] = \mu_i - \mu_j$, $Var[Y_{ij}] = 2\sigma^2$. The density of the 3-dimensional random variable $Y = (Y_{ab}, Y_{bc}, Y_{cd})'$ is normal with covari-

¹For the sake of precision, clients divert money each year from the worst to the best performing fund.

	X_1	X_2	X_3	X_4
u	0.208	0.232	0.256	0.305
d	0.295	0.265	0.240	0.200

Table 1: Probabilities to experience an inflow, u , and outflow, d , for each of the four funds. Data are based on the excess returns provided in [Kosowski et al., 2006].

ance matrix Σ that can be easily derived. For example,

$$\begin{aligned}
Cov(Y_{ab}, Y_{bc}) &= E[Y_{ab}Y_{bc}] - E[Y_{ab}]E[Y_{bc}] \\
&= E[(X_a - X_b)(X_b - X_c)] - (\mu_a - \mu_b)(\mu_b - \mu_c) \\
&= (\mu_a\mu_b - \mu_a\mu_c - (\sigma^2 + \mu_b^2) + \mu_b\mu_c) - (\mu_a\mu_b - \mu_a\mu_c - \mu_b^2 + \mu_b\mu_c) \\
&= -\sigma^2, \\
Cov(Y_{ab}, Y_{cd}) &= E[Y_{ab}Y_{cd}] - E[Y_{ab}]E[Y_{cd}] \\
&= E[(X_a - X_b)(X_c - X_d)] - (\mu_a - \mu_b)(\mu_c - \mu_d) \\
&= 0.
\end{aligned}$$

Equipped with

$$\mu = (\mu_a - \mu_b, \mu_b - \mu_c, \mu_c - \mu_d)' \text{ and } \Sigma = \begin{pmatrix} 2\sigma^2 & -\sigma^2 & 0 \\ -\sigma^2 & 2\sigma^2 & -\sigma^2 \\ 0 & -\sigma^2 & 2\sigma^2 \end{pmatrix},$$

the probability P_{abcd} can be computed as

$$P_{abcd} = \Phi_{\mu, \Sigma}(0, 0, 0). \quad (1)$$

In turn, the probabilities u_j and d_j to be ranked best or worst and, as a result, to experience an inflow or outflow, can be computed for each fund ($j = 1, \dots, 4$) using (1) and summing over the appropriate subset of permutations. Table 1 shows for each fund the chances to have positive (u) and negative (d) flows, using the aforementioned values. The probabilities in Table 1 show that the chance for the less skilled manager X_1 to get a positive inflow are exceeding 20%. Moreover, intermediate funds X_2 and X_3 have virtually the same probabilities to be ranked at opposite extremes. This means that real mutual funds in the second and third quartiles of performance are likely to be very similar in terms of flows, despite considerable difference in returns. In other words, myopically switching funds from the worst to the best performer is consistent with the flat performance-flow curve that is empirically documented for mutual funds.

This result holds under general conditions about the size of μ and σ .

Theorem 1 *The probability P_{abcd} is constant (independent of the permutation) if*

$$\frac{\mu}{\sigma} \rightarrow 0.$$

Proof. We can write P_{abcd} using (1) as

$$P_{abcd} = Const \times \int_{-\infty}^0 \int_{-\infty}^0 \int_{-\infty}^0 \exp \left(-\frac{1}{2} (\mathbf{x} - \mu)' \Sigma^{-1} (\mathbf{x} - \mu) \right) d\mathbf{x},$$

where $\mathbf{x} = (x_1, x_2, x_3)$ and $\mu = (\mu_a - \mu_b, \mu_b - \mu_c, \mu_c - \mu_d)'$. Changing variables by the position $\mathbf{y} = (\mathbf{x} - \mu)/\sigma$ we have

$$P_{abcd} = Const \times \int_{-\infty}^{-\frac{\mu_a - \mu_b}{\sigma}} \int_{-\infty}^{-\frac{\mu_b - \mu_c}{\sigma}} \int_{-\infty}^{-\frac{\mu_c - \mu_d}{\sigma}} \exp \left(-\frac{1}{2} \mathbf{y}' \begin{pmatrix} 2 & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 2 \end{pmatrix}^{-1} \mathbf{y} \right) d\mathbf{y}.$$

Taking the limit for $\mu/\sigma \rightarrow 0$ and observing that the integrand does not depend on the permutation through $\mu_a, \dots, \mu_d$ proves the result. ■

Theorem 1 shows that in the limit case in which the standard deviation of the excess returns is very large relatively to the mean excess returns, then every rank is equally likely. Hence, the inflows and outflows are equally probable resulting in a (totally) flat performance curve. The result is insightful as the standard deviation measures the confusion arising in the process to pick the best manager: if σ is (infinitely) large, customers are unable to reward consistently the best manager and the flows are flat. The proof can be easily generalized to many funds along the same lines.

However, Theorem 1 crucially depends on extreme assumptions and does not provide ways to compute the switching probabilities for any actual and bounded set of values for μ and σ . Moreover, were customers switching every T years or according to a performance measure computed on a time window of length T , the mean-standard deviation ratio relative to T years would become $\mu T / \sigma \sqrt{T}$ which, for large T , could be bounded away from 0. The theorem offers in this case little guide, if any.

To overcome these limitations, we now describe a richer computational model of flows into and out of mutual funds.

3 A computational model

In this section we describe a micro-simulation model to explore some departures from the over-simplified setup that can be treated analytically. In this respect, this section is an effort to explain the empirical regularity of the flow-performance curve using an agent-based model, in the spirit of [Lettau, 1997]. As we made clear previously, we model N mutual funds that have identical and independent yearly excess return described by a normal distribution $N(\mu_i, \sigma_i), i = 1, \dots, N$. In the absence of any flow, the net asset value w_{it} of fund i at time t changes due to the return R_{it} :

$$w_{it} = w_{i,t-1}(1 + R_{it}),$$

where R_{it} is a random draw from $N(\mu_i, \sigma_i)$.

Every year all the funds are ranked according to the mean past excess return over T years. We assume that (due to the action of customers that are not modelled), each of the worst x funds loses a fraction α of its net asset value that is randomly transferred to one of the best x performing funds. The Sheriff of Nottingham behaviour ensures that all the x -worst funds lose some money but the x -best funds do not receive equal inflows and (unrealistic) symmetries are avoided due to the random picking of the recipient. Observe that different choices for x and α give rise to different situations. Setting, say, $x = N/4$ is resemblant of the setup described in Section 2, while $x = N/5$ is somewhat mimicking the 5-star Morningstar rating. Another important parameter in the model is α , the percentage of wealth that is removed from the bad funds to flow to the best ones. In general, we assume that α is much smaller than 100%, meaning that poor performing funds face some outflows but not massive withdrawals. Among the reasons there are, as seen in the Introduction, heterogenous willingness of investors to remove capital from the manager and the problematic inference about future returns that can be drawn. We think that another obvious reason is a simple diversification effect. Customers may indeed draw away only a part of the wealth invested low-return funds in order to stay invested in different mutual funds to avoid keep all the eggs in the same basket.

So far the model cares only about “internal” flows, that depends on x, α and on the distributions $N(\mu_i, \sigma_i)$. However, the growth rates of the value under management are obviously contingent on the flow of exogenous wealth that is injected in the funds. The magnitude of this external flow clearly impacts the growth rates and data show that a great part of the increment in the funds’ NAV is due to the entry of fresh capital. Hence we introduce external flows in the model and assume that (all of) the best x funds acquire new customers so that their NAV is increased by a factor $1 + \beta$. The fact that new customers select one of the best funds for their new investment is rather natural and the previously mentioned papers offer some guidance in the choice of the size of β , which is realistically in the range 20%-40%, see also [Ding et al., 2008]. Of course, letting $\beta = 0$ is *de facto* allowing for internal flows only and it is an important special case for our model. We compare this benchmark with the case $\beta > 0$ to better enhance the effects of external flow on the flow-performance curve.

3.1 Results

All our results are based on averages computed using 200 simulations for different values of the parameters. We do not adventure ourselves in the estimation of the parameters, but to keep some resemblance to real situations and to the literature, we consider $N = 500$ funds over a period of 20 years. Representative values for excess returns (that are periodically used

by customers to rank the funds) are listed, say, in [Kosowski et al., 2006]. The $\mu_i(\sigma_i)$ are obtained grouping funds in deciles based on the yearly performance (in percent): $\mu_1 = 3.36$ (15.28), $\mu_2 = 4.80$ (14.46), $\mu_3 = 5.04$ (14.27), $\mu_4 = 5.40$ (13.86), $\mu_5 = 5.52$ (13.75), $\mu_6 = 6.36$ (13.89), $\mu_7 = 6.48$ (14.13), $\mu_8 = 6.96$ (14.58), $\mu_9 = 7.92$ (15.55) and $\mu_{10} = 9.36$ (18.36). Observe that the observed σ 's are approximately constant and that excess returns are often an order of magnitude smaller than the standard deviation. The assumptions of Theorem 1 are very roughly met but there is clearly scope for simulations to explore the effects of departures from that setting.

The simulation of N funds when only deciles-based statistics are known requires some caution. We could replicate $N/10$ times the values listed above for μ_i and σ_i but this produces a very granular set of funds. Moreover, every simulation would have the very same set of values for the parameters while deciles do not determine a single composition of the set. Hence, we resort to a resampling algorithm described in [Thompson, 2000]. The method, called SIMDAT, allows for the creation of pseudo-samples of arbitrary size given a matrix of reference observations $\mathbf{X}$ (in this case $X_i = (\mu_i, \sigma_i), i = 1, \dots, 10$). The pseudo-samples are generated by first picking at random one of the X_i and then blending its components using linear combinations with its k -nearest neighbors. The shape of the resulting distribution can be tuned modifying k and, as extreme cases, we can sample from the singular distribution that gives equal non-null mass to the points X_i when $k = 1$ or from a unique multivariate density $N(E[\mathbf{X}], Cov(\mathbf{X}))$ when k approaches infinity. We take a reasonable compromise in the generation of the pseudo-samples and set $k = 5$ that produces a cloud of point visually approximating fairly well the deciles statistics plotted on a $\mu\sigma$ -plane, as shown in Figure 1. Additional details on the SIMDAT algorithm can be found in Thompson's book and the code can be freely downloaded at <http://lib.stat.cmu.edu/S/gendat>. We tried other values of k to check that our results do not change in any significant way.

The yearly growth rate of a fund i at time t is defined as:

$$gr_{it} = \frac{w_{it} - w_{i,t-1}(1 + R_{it})}{w_{i,t-1}} \quad (2)$$

Each year, funds are ranked based on their mean excess return over the previous T years and flows are generated as described before (because of internal movements driven by α and exogenous streams related to β). It is difficult to estimate the length T for the revision window used by investors. Most probably there is notable heterogeneity among different kinds of mutual funds' customers but there is evidence that long horizons are sometimes used. Financial media, for example, often report also 3- and 5-year returns for funds' investments. The average holding period of funds' shares may be related to the revision frequency, even though they may not be necessarily equal. Informal evidence is supportive of average holding periods of several

Figure 1: Simulated excess returns and standard deviations for 100 funds (circles). The points are a pseudo-sample of the 10 deciles statistics quoted in the text (solid circles) and are obtained using SIMDAT with $k = 5$.

years and [Sirri and Tufano, 1998] reports time-spans as long as 7 years. We investigate values for T in the set $\{1, 3, 5\}$.

We select $\alpha = 10\%, 15\%, 20\%$ so that only a portion of the invested monies are switched away in a single year after a poor outcome. The previous discussion also motivates our choice to pick x in $\{25, 75, 125\}$. As $N = 500$ this means that being ranked in the worse 5% to 25% of funds causes some outflow by the invested agents.

We present our results having in mind a somewhat representative benchmark configuration where $(\alpha^*, x^*, T^*, \beta^*) = (15\%, 75, 3, 30\%)$. We then move one parameter at a time, keeping the others fixed to investigate the individual effects and the sensitivity of our findings.

Figure 2 depicts the arithmetic average growth rate of the funds as a function of the mean excess return. Different values for β are considered and we clearly see the flat and convex shape of the flow-performance curve. There are two interesting features in the graph. The left part of the curves is nearly constant, demonstrating that the flow pattern of a bad fund is pretty similar to the one of the median performer. The right part of the curve is, instead, seen to be rising steeply. In stark contrast with the low μ zone, this means that the best funds enjoy much higher relative inflows than the relative outflows of the worst ones. The picture shows that this effect is depending on the size of β , that has also the unsurprising effect to raise the level of the whole curve brought by increased external flows infused in the economy.

It is also worthwhile noticing that the position $\beta = 0$ produces a remarkably flat curve. This situation devoid of any external flow is very close to the assumptions of Theorem 1 and indeed the theoretical result, predicting

Figure 2: Flow-performance curves for $\beta \in \{0, 20\%, 30\%, 40\%\}$. The graphs are obtained averaging the results of 200 simulations.

a constant curve, is a decent portrait of the simulations. A similar, almost constant outcome is found for other values of the parameters, provided that β is held at 0.

The right panel of Figure 3 show the flow-performance curves for $\alpha = 10\%, 15\%, 20\%$. The sensitivity with respect to α is relatively low and the same realistic shape is robustly generated. The level of the curves increases with α and this holds even in the right part of the μ -axis. Two effects are competing here: a bigger α leads to bigger outflows from the worst funds and this reduce the average growth rate. At the same time, however, a fund with low μ occasionally ranks high and receives a relatively large and even outlying inflow. This lucky event has non-negligible probability in view of the insights given by Theorem 1 and more than compensate the more commonly observed negative growth rates when taking averages.

The left panel of Figure 3 shows the results of our simulations when $x = 25, 75, 125$. The value of x measure the extent of the region where below-average funds are considered “bad” and consequently experience outflows that are gained by symmetrically positioned good funds. The flow-performance curves relative to $x = 25, 75$ are both convex and show the usual shape with a flatter region followed by a increasing section on the right. Setting $x = 25$, i.e. allowing only 5% of the funds to be faced with withdrawals, produces a mild curvature and probably $x = 75$ is a better choice that enhances rather appropriately the rise in the growth rate of the best performers. Increasing x has some interesting effects. On the one hand the set of funds that suffers outflows or enjoys inflows is much larger. Due to the sizeable standard deviations of the returns used to rank the funds, this results in (relatively) many cases where bad to average funds rank in the top part of the standing. On the other hand, the U-shaped graph shows

Figure 3: Flow-performance curves for $\alpha \in \{10\%, 15\%, 20\%\}$, on the left, and $x \in \{25, 75, 125\}$, on the right. The graphs are obtained averaging the results of 200 simulations.

that the previously described effect for α enters the stage here as well. Even if inferior funds lose wealth more often than not, nonetheless they are some occasions when they get inflows that are huge in relative terms. As there is no evidence of a U-shaped curve in the empirical literature, the examination of the whole picture lead to the belief that investors are probably escaping from a fund when the performance is extremely negative (worse 5 or 15% of the industry), implicitly tolerating without any action a rank close to the 3-rd quartile.

Figure 4: Flow-performance curves for different values of $T = 1, 3, 5$ years. The graphs are obtained averaging the results of 200 simulations.

The effect of varying T in the set $\{1, 3, 5\}$ is depicted in Figure 4. Only

the left part of the flow-performance curve is visibly affected. Low values for T produce volatile periodic rankings and, being the T -period average excess return smaller than the corresponding dispersion, the dynamics predicted by Theorem 1 is at work (flat curve), with the usual distortion on the right due to the external inflows. The economic intuition is clear: investors found their decision on short term performance and, hence, easily mistake the good fund with the bad. If, conversely, a longer period is used for the evaluation, the ranking is more consistent because more observations are used. Moreover, for many funds μT is of the same order of magnitude of $\sigma\sqrt{T}$ when $T = 5$ and indeed the flow-performance curve almost linearly penalizes the poor-performers in favour of the funds with higher values of μ . Again this fits quite well with the basic intuition that careful investors, taking some years to assess the ability of their managers, can better pick and reward funds with average superior returns. The finding that slower and more cautious agents are less prone to generate puzzling outcomes is not surprising and was reported in famous agent-based papers like in [Lettau, 1997], where longer T cut down the size of the risk exposure taken by the agents, or in [Arthur et al., 1997] where infrequent activation of the learning device suffices to obtain prices not far from the rational-expectations equilibrium.

4 Conclusion

We have shown that a “Sheriff of Nottingham” behaviour of profit-chasing clients of mutual funds can robustly produce a fairly flat performance-flow curve, in the absence of external flows. The main ingredients of the model are moderate levels of x and α , the number of funds that are rated as bad/good and the fraction of invested wealth that is moved away, respectively. Our model can robustly generate flow-performance curve that are qualitatively close to the empirically observed ones, if some exogenous flow is taken into account. Various papers in the literature offer interesting (and distinct) explanations for the flat-convex shape of the flow-performance curve and all the proposed mechanisms are in a sense versions of this model. We have proved a proposition that shows that a constant curve is indeed produced (for any x and α) when the ratio between average returns and standard deviation goes to zero. This analytical result stresses that investors are incapable to reward skilled managers in a confused world where average excess returns are contaminated by high randomness due to the sizeable standard deviation of the same returns.

We show that a sufficient condition to bend up the right part of the curve is to assume that fresh funds are invested in the best x -funds that have their net asset value increased by a factor $1 + \beta$. The parameter β , combined with x , controls the amount of exogenous flows and affects the level and curvature of the curve.

By simulations and comparative static exercises, we demonstrate that our results are robust and hold for large intervals of the parameters. Most notably, large values for x are disruptive in that they cause unrealistic U-shaped curves. Finally, we show that the length T of the time-span used by agents to evaluate the average returns of the funds has some impact. Allowing for longer T gives rise to an almost linear curve, while shorter horizons are more likely to produce realistic flow-performance curvature. This behaviour is an indirect test of Theorem 1 as an increase in T makes the required assumptions more and more inaccurate.

References

- [Arthur et al., 1997] Arthur, W., Holland, J., LeBaron, B., and Taylor, P. (1997). Asset pricing under endogenous expectations in an artificial stock market. In Arthur, W., Lane, D., and Durlauf, S., editors, *The economy as an evolving complex system II*. Addison-Wesley.
- [Berk and Tonks, 2007] Berk, J. B. and Tonks, I. (2007). Return persistence and fund flows in the worst performing mutual funds. NBER Working Papers 13042, National Bureau of Economic Research, Inc.
- [Ding et al., 2008] Ding, B., Getmansky, M., Liang, B., and Wermers, R. R. (2008). Investor Flows and Share Restrictions in the Hedge Fund Industry. *SSRN eLibrary*.
- [Del Guercio and Tkac, 2008] Del Guercio, D. and Tkac, P. (2008). Star power: The effect of Morningstar ratings on mutual fund flow. *Journal of Financial and Quantitative Analysis*, 43(04):907–936.
- [Kahneman and Tversky, 1979] Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–91.
- [Kosowski et al., 2006] Kosowski, R., Timmermann, A., Wermers, R., and White, H. (2006). Can mutual fund "stars" really pick stocks? New evidence from a bootstrap analysis. *Journal of Finance*, 61(6):2551–2595.
- [Lettau, 1997] Lettau, M. (1997). Explaining the facts with adaptive agents: the case of mutual fund flows. *Journal of Economic Dynamics and Control*, 21:1117–1147.
- [Lynch and Musto, 2003] Lynch, A. and Musto, D. (2003). How investors interpret past fund returns. *Journal of Finance*, 58(5):2033–2058.
- [Sirri and Tufano, 1998] Sirri, E. and Tufano, P. (1998). Costly search and mutual fund flows. *Journal of Finance*, 53(5):1589–1622.

[Thompson, 2000] Thompson, J. (2000). *Simulation: a Modeler's Approach*. Wiley.