

DISCUSSION PAPER SERIES

DP11896

PICKING FUNDS WITH CONFIDENCE

Niels Groenborg, Asger Lunde, Allan Timmermann
and Russ Wermers

FINANCIAL ECONOMICS



PICKING FUNDS WITH CONFIDENCE

Niels Groenborg, Asger Lunde, Allan Timmermann and Russ Wermers

Discussion Paper DP11896

Published 09 March 2017

Submitted 09 March 2017

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **FINANCIAL ECONOMICS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Niels Groenborg, Asger Lunde, Allan Timmermann and Russ Wermers

PICKING FUNDS WITH CONFIDENCE

Abstract

We present a new approach to selecting active mutual funds that uses both holdings and return information to eliminate funds with predicted inferior performance through a sequence of pairwise comparisons. Our methodology determines both the number of skilled funds and their identity; funds identified ex-ante as being superior earn substantially higher risk-adjusted returns than top funds identified by conventional alpha ranking methods. Importantly, we find strong evidence of variation in the breadth of the set of funds identified as superior, as well as fluctuations in the style and industry exposures of such funds over time and across different volatility states

JEL Classification: G2, G11, G17

Keywords: Fund confidence set, equity mutual funds, risk-adjusted performance

Niels Groenborg - ngruenborg@econ.au.dk
University of Aarhus

Asger Lunde - alunde@econ.au.dk
University of Aarhus

Allan Timmermann - atimmerm@ucsd.edu
UCSD and CEPR

Russ Wermers - rwermers@rhsmith.umd.edu
University of Maryland

Acknowledgements

We thank conference participants at the 2016 SoFiE conference and seminar participants at Erasmus University, Rotterdam, Georgetown University, University of Maastricht, Tilburg, and CReATES for comments on the paper. Grønborg gratefully acknowledges support from the Danish Council for Independent Research, Social Sciences (4091-00190/FSE)

Picking Funds with Confidence*

Niels S. Grønborg

Aarhus BSS and CReATES

Allan Timmermann

UCSD

Asger Lunde

Aarhus BSS and CReATES

Russ Wermers

University of Maryland

February 23, 2017

Abstract

We present a new approach to selecting active mutual funds that uses both holdings and return information to eliminate funds with predicted inferior performance through a sequence of pair-wise comparisons. Our methodology determines both the number of skilled funds and their identity; funds identified ex-ante as being superior earn substantially higher risk-adjusted returns than top funds identified by conventional alpha ranking methods. Importantly, we find strong evidence of variation in the breadth of the set of funds identified as superior, as well as fluctuations in the style and industry exposures of such funds over time and across different volatility states.

Key words: Fund confidence set; equity mutual funds; risk-adjusted performance

JEL codes: G2, G11, G17

*We thank conference participants at the 2016 SoFiE conference and seminar participants at Erasmus University, Rotterdam, Georgetown University, University of Maastricht, Tilburg, and CReATES for comments on the paper. Grønborg gratefully acknowledges support from the The Danish Council for Independent Research, Social Sciences (4091-00190/FSE)

1. Introduction

Each year, large sums of money and vast amounts of effort are dedicated towards selecting the best mutual funds.¹ Despite such interest, and despite decades of academic research, identifying skilled active mutual funds largely remains an elusive goal. While many papers have singled out different characteristics of individual funds that are correlated with performance, it is unclear how to effectively identify the *set* of funds that is most likely to outperform in the future.²

Two main issues plague this identification. First, members of this set can change substantially as broad economic conditions evolve, resulting in the set dramatically expanding or shrinking—as well as containing different component funds—during different years.³ Specifically, the nature of a fund’s information—and the fund’s strategies for acting on such information—can depend on the competitiveness in the fund’s peer-group as well as the state of the economy, both of which evolve and can lead to changes in the set of funds that outperform.⁴ For example, the ability of a fund to outperform may be short-lived because evidence of successful strategies is bound to attract more competition (Hoberg, et al., 2015) as well as additional investor flows.

A second major issue is whether an investor can, *ex-ante*, locate the number and identity of the set of active funds that will outperform in the future (with a high level of probability). Estimates of fund risk-adjusted returns (alphas) tend to be surrounded by large sampling errors, as outperforming fund managers tend to carry a greater level of idiosyncratic risk (Kacperczyk, et al., 2005). In addition, the estimation error in using past fund data renders any fund-by-fund analysis as resulting in a large number of “false discoveries” of funds with disappointing future performance.⁵ The model of Kacperczyk et al. (2016) implies that dispersion in return performance across active fund managers increases during periods with high uncertainty, which both heightens the need for an efficient identification of the best funds and makes it more statistically challenging.⁶

¹As of 2015, roughly \$8 trillion was invested in U.S. equity mutual funds, with the majority of this money invested in an actively managed mandate. A large industry of investment advisors and consultants is engaged in advising retail and institutional clients on how to select funds. According to the 2015 Cerulli RIA Marketplace 2015 report, there were more than 56,000 registered investment advisors in the U.S. in 2014. Foerster et al (2015) study how financial advisors affect the decisions (and performance) of a sample of Canadian retail investors.

²See, e.g., Cremers and Petajisto (2009), Kacperczyk, et al. (2005, 2008), and Wei, et al. (2015). For a survey of the literature on the characteristics that predict active fund performance, see Jones and Wermers (2011).

³Recent studies suggest that the ability of individual mutual funds to outperform varies over time. For example, Kacperczyk et al. (2014) find that the investment strategies of mutual funds, as well as their ability to outperform, depend on whether the economy is in an expansion or in a recession state. Similarly, Glode, et al. (2011) find evidence of strong (weak) predictability of mutual fund returns following periods of high (low) market returns. Ferson and Schadt (1996), Avramov and Wermers (2006), and Banegas et al. (2013) show that macroeconomic state variables can be used to better predict the future performance of individual equity mutual funds.

⁴Kacperczyk, et al. (2016) provide a model that delivers predictions motivating empirical tests of the changing size of this set. Their model predicts larger outperformance for skilled funds during stressed states in the economy (e.g., high market volatility), as compared to more stable states.

⁵This estimation error can be especially keen during stressed states, due to the above-noted higher variability in performance of active funds in such states.

⁶During periods with high aggregate volatility, individual fund managers lean more on their private signals, and

Mutual funds’ fleeting competitive advantage—and the challenge to identifying transient skills—has been discussed in both empirical and theoretical studies. Empirically, Carhart (1997) finds that the performance of top-ranked funds reverts towards the mean after about one year, while Bollen and Busse (2004) find abnormal performance that lasts for one quarter, but disappears at longer horizons. Berk and Green (2004) propose a theoretical model that leads to fund alphas that converge toward zero. Mamaysky, et al. (2008) develop a model in which managers observe private information signals which revert toward being uninformative. Similarly, Glode et al. (2011) present a flow-based model in which diseconomies-of-scale at the fund level remove any abnormal performance over time as investors allocate more money to small funds with high past alphas and allocate less money to large funds with negative past alphas.⁷

This paper introduces an efficient approach to identifying (ex-ante) the set of funds with superior performance, as well as identifying how large the selected set of funds is, what types of investment strategies they adopt, and how this set of funds (along with its risk-adjusted performance) evolves over time. Our approach to identifying the set of “best” (or superior) funds requires not only that we compare each fund’s performance against a single benchmark (or a set of risk factors, as is common practice), but that we also conduct a large set of pairwise comparisons of all funds in existence to eliminate any funds whose performance is dominated by at least one other fund. This comparative process is crucial in providing the highest certainty of capturing the most skilled fund managers, as it eliminates funds from the choice set that are likely to reduce the composite performance of the set.

Besides providing valuable information for investors, our approach also sheds new light on the nature of competition in the mutual fund industry. Specifically, by singling out the set of funds with superior performance and showing how this set evolves through time, we can analyze how many funds are in this set, for how long they can maintain their competitive edge, and how rapid the turnover is among funds included in the set. Moreover, by cross-checking fund performance against quarterly holdings, we gain information into which strategies the funds pursue in order to obtain superior performance.⁸

Conventional approaches in the finance literature are not well-designed to handle such comparisons, nor do they control for the “size” of the test, i.e., the probability of wrongly eliminating truly superior funds.⁹ To deal with such issues, our analysis adopts a new approach for the selection

less on common, aggregate signals, which increases the cross-sectional heterogeneity in fund-manager beliefs and, as a consequence, in their investment strategies and portfolio returns.

⁷Conversely, the analysis in Garleanu and Pedersen (2016) suggests that, due to differential search costs, funds populated by the most sophisticated investors can be expected to outperform in equilibrium.

⁸Our search for funds with superior performance uses only simple and widely available data on past returns and publicly reported holdings—both of which are available at the time a decision on whether to invest in the fund is made.

⁹While the “False Discovery Rate” procedure used by Barras, Scaillet, and Wermers (2010) provides an efficient approach to estimating the Type I error (size) associated with identifying a set of outperforming funds, it says little about the error rate in identifying individual funds within a given set that are expected to outperform. Indeed, Barras

of mutual funds that makes use of the Model Confidence Set (MCS) methodology of Hansen et al. (2011), which is designed to select the most accurate prediction models from a large set of candidate models. Hansen, et al. (2011) show that a step-wise bootstrap approach can be used to determine critical values for the elimination of models from the MCS in a way that bounds the probability of wrongly eliminating a truly superior model by the size of the test while eliminating inferior models with a probability that asymptotically approaches one.

In our context, the set of “candidate models” is the list of mutual funds in existence at a given point in time, and their performance is measured through risk-adjusted returns by a particular econometric factor equation (e.g., the four-factor model of Carhart, 1997). Like the MCS methodology, our approach undertakes a series of pair-wise tests to sequentially eliminate funds with inferior performance. If at least one fund with significantly inferior performance is identified, the fund with the “worst” performance is eliminated, and the elimination process is repeated on the reduced set of funds. The procedure continues until no further funds with inferior performance can be identified and eliminated. We label the set of funds remaining at the end—the funds identified to have superior performance—the “Fund Confidence Set” (FCS).

The FCS approach uses estimates of individual funds’ risk-adjusted returns as a way to rank and compare their performance, and, thus, the results from applying this methodology will depend on how good the underlying performance model is at extracting information about fund manager skills. We apply the FCS approach to three different performance models which assume (i) constant (or slowly evolving) alphas; (ii) time-varying alphas, using the latent skill approach of Mamaysky et al. (2008); (iii) time-varying alphas with alpha estimates extracted using both fund returns and holdings data. For each performance model, we find that the FCS approach can be used to select funds and form portfolios of funds with considerably better performance than portfolios containing a fixed proportion (e.g., 5% or 10%) of top ranked funds. The performance results are particularly strong for the model that combines data on returns and holdings to obtain a sharper estimate of fund alphas. Using this model, an equal-weighted portfolio of funds included in the top FCS generates a four-factor alpha exceeding 50 bps/month (which is also highly statistically significant); for comparison, the four-factor alpha estimate of a constant allocation to the top decile of funds is only about 2 bps/month and is statistically insignificant.

Empirically, we find that the FCS approach can be used to identify a narrow set of funds with superior performance. In fact, because the performance of most funds is estimated with large sampling errors, the approach is found to work best if we use a relatively stringent criterion for inclusion of funds in the FCS, resulting in a reduced probability of wrongly including funds with inferior performance. Using this approach, the set of funds identified as superior fluctuates considerably over time. Sometimes, only a single fund is selected for a long stretch of time; at other

et al. (2010) apply a simple approach that examines the set of funds with the best estimated alphas (or t-statistics), with no mechanism to winnow out individual funds within this set that are unlikely to be skilled.

times, the approach identifies a much broader set of funds as being superior. Moreover, the breadth of the set of funds with (expected) superior performance is found to be significantly correlated with a range of macroeconomic state variables, as the fraction of superior funds is larger in expansions, and smaller in recessions, suggesting that funds' ability to outperform is state-dependent.¹⁰

We also find that dispersion in fund performance is greater during times with high levels of market volatility, leading to the best funds—as identified by our FCS approach—performing better during periods with high market volatility, relative to calmer times. This is consistent with the model of Kacperczyk, et al. (2016) which suggests that the rewards to skilled investors from processing information are higher when market risk is high. Interestingly, we find no significant following-year outperformance of top decile funds during high risk periods. This result highlights the efficacy of our FCS approach, which carefully eliminates funds that mainly contribute idiosyncratic risk—and not alpha—to the set of superior funds.

In a deeper analysis of the portfolio holdings data of the mutual funds, we find that the changing set of funds that are identified as superior results in a change in industry concentration and a significant shift in risk loadings. For example, the superior FCS funds tilt toward growth stocks from 1994 to 2003 (a period of great return variability in the cross-section of growth stocks), as well as from 2009 to 2011. Conversely, they overweight value stocks from 2003 to 2009 (a period of great return variability in the cross-section of value stocks). Superior funds also overweight small stocks from 2001 to 2003 and 2009 to 2012, but over-weight large capitalization stocks from 2005 to 2008. Superior funds have above-average exposures to the momentum factor only during brief spells. Turning to industry concentration, funds in the top FCS overweight computer and electronic equipment stocks after 2005, and also overweight business services and machinery stocks during shorter spells. Conversely, they underweight retail stocks and, in particular, banking stocks throughout most of the sample. All of these results indicate that the identification of superior funds is both time-varying and complex; our modeling procedure uniquely provides an approach to such identification.

The FCS methodology can also be used to successfully identify funds with inferior performance. When applied to select inferior funds, we find that the set of “worst” funds is somewhat wider than the equivalent set of superior funds. This reflects the greater persistence of factors giving rise to underperformance, such as high trading costs and management fees. Again, we find that the funds in the bottom FCS portfolio produce substantially worse performance than a portfolio consisting of a fixed proportion of alpha-ranked funds, such as the bottom 5-10% of funds, indicating that the FCS effectively identifies funds with poor skills rather than funds that have experienced bad luck.

Our analysis differs from previous studies in several important dimensions. Kosowski, et al.

¹⁰This result highlights the difficulty in identifying, ex-ante, superior funds in a recession setting—even though the average performance of active funds is better during recessions—due to the short time-series nature of recessions. For instance, there are only 34 months of NBER-identified recessions during our entire sample period.

(2006) ask whether there exist star fund managers, i.e., if the single best manager can outperform some benchmark, such as the four-factor model of Carhart (1997), or, alternatively, if a single fund at a particular ranked percentile (i.e., the fund at the top 10 percentile point), can outperform. However, their methodology cannot be used to endogenously determine the size of the *set* of funds with superior performance or the identity of the individual funds. In fact, a strategy of only investing in the fund whose alpha is deemed highest can sometimes backfire, because such a fund might have been “lucky,” and its performance could reflect very high idiosyncratic risk taking. Barras, et al. (2010) develop an approach that controls for funds whose high alpha estimates are due to “luck.” However, unlike us, they do not address the identification of *individual* funds deemed to have the highest (positive) alphas. This is a highly relevant question from an investor’s perspective, because not all funds, even those with truly positive alphas, should be expected to perform equally well.¹¹ In addition, investors typically wish to invest in (or have access to) only one or a few funds, and not a larger set. Thus, the identification of the set of superior funds is of great value to such investors. Further, Barras, et al. (2010) find evidence that the set of funds with positive alphas has been shrinking over time, making it more difficult to identify truly superior fund managers. This result highlights the importance of identifying individual funds that are predicted to outperform in order to form a precise set, rather than starting with a set and conducting inference on the entire set (equal-weighted) as is done by Barras, et al. (2010).¹²

We show empirically that our approach, which endogenously builds the set of funds (starting with pair-wise individual fund comparisons) expected to produce superior performance, is far better at identifying truly skilled managers than the conventional portfolio sorting approach. Thus, our approach has two major advantages over past methods. First, it helps to discard individual funds that lie in the top past-performance fractile due to luck. And, second, our approach helps to adjust the search for superior active managers in accordance with the breadth of active management skills during a particular period of time, broadening or narrowing the set of chosen funds as market conditions evolve. We show that these two advantages translate into superior fund alphas, as well as better identification of inferior funds.

The outline of our paper is as follows. Section 2 introduces the fund confidence set methodology, while Section 3 describes our data and the models used to measure the performance of individual funds. Section 4 presents performance results for the FCS portfolios and Section 5 provides details on which funds get selected to be among the superior or inferior funds and how this set varies through time. Section 6 performs an attribution analysis which decomposes the performance of the FCS funds using holdings data on industry concentrations and stock-level estimates of individual

¹¹In fact, we show that an approach similar to that of Barras et al. (2010) can be used to screen the initial set of candidate funds from which the fund confidence set is selected.

¹²Harvey and Liu (2016) estimate individual fund alphas using a random effects panel approach that pools time series and cross-sectional information from the population of funds. This approach does not endogenously identify the set of funds with superior performance, however.

funds' exposures to common risk factors. Section 7 concludes.

2. The Fund Confidence Set

A large empirical literature in finance explores whether it is possible to, *ex ante*, identify funds with superior risk-adjusted performance. The most widespread practice used in the literature is, first, to rank individual funds based on their expected (predicted) performance, then, second, to form decile portfolios based on such rankings and, finally, to track the subsequent risk-adjusted performance of the portfolios of funds. While the practice of allocating individual funds into decile portfolios is simple to perform and intuitive to interpret, it only addresses whether the *average* risk-adjusted performance of the top 10% of funds (or a similar proportion) is positive. This question is very different from the more relevant and interesting question of whether we can identify a particular set of individual funds with superior performance, regardless of whether they reside in a particular past-performance ranking quantile.¹³ For example, if the set of funds capable of producing positive risk-adjusted performance varies over time, and is sometimes far narrower than 10%, then we may well find, empirically, that the top 10% of funds do not generate positive performance, on average, even though there exists a set of funds with positive (ex-ante) risk-adjusted performance. Or, suppose that the top 10% of funds produce positive risk-adjusted returns, but that the top 5% produce much higher (and more precisely estimated) risk-adjusted returns. In such a case, an investor would be well-served to invest only in the top 5% of funds in order to maximize her chances of generating an economically significant level of outperformance.

This section introduces an alternative approach to fund selection that does not fix the proportion of funds deemed capable of delivering superior performance but, rather, determines this proportion endogenously, as part of the process used to estimate the performance of individual funds. Our econometric approach reflects how a sophisticated investor with access to historical data on individual funds' returns and security holdings would go about identifying funds with skills and thus mimics the search process underlying the analysis in papers such as Garleanu and Pedersen (2016). We first describe the approach in broad terms and characterize its properties, then provide details on how we implement the approach on our mutual fund data.

¹³An especially salient example of the failure of simple ex-post ranking approaches is the “stars” system of Morningstar. Kinnel (2010) shows that expense ratios, by themselves, are better predictors of future fund performance than the Morningstar star rating—which includes the effect of both pre-expense past returns and current expense ratios. Morey (2003) finds that Morningstar five-star rated funds experience a severe drop in performance over the following three years.

2.1. Predictive Alpha

For a fund to be attractive to investors, it must have a high expected risk-adjusted performance, relative to the benchmark model that those investors use. This requires that the fund's performance is at least modestly predictable.¹⁴ Our objective is, therefore, to identify funds which, with some confidence, we can predict will produce positive future risk-adjusted returns. Following common practice, we compute a fund's risk-adjusted return by adjusting the fund's excess returns, net of the T-bill rate, $R_{i,t}$, for its exposure to a set of risk factors, \mathbf{z}_t :

$$R_{i,t} = \alpha_{i,t} + \boldsymbol{\beta}'_{i,t} \mathbf{z}_t + \varepsilon_{i,t}. \quad (1)$$

Here i refers to the fund and t refers to the time period; $\varepsilon_{i,t} \sim (0, \sigma_{\varepsilon_i}^2)$ is the fund's idiosyncratic return, $\boldsymbol{\beta}_{i,t}$ measures the fund's exposure to the common risk factors, while $\alpha_{i,t}$ measures its risk-adjusted (abnormal) performance, often referred to as the fund's "alpha".

The model in equation (1) is quite general as it allows both α_{it} and β_{it} to vary over time. If a fund's alpha is constant over time, the fund's average historical performance can be used to compute its expected future performance. Conversely, if a fund's abnormal performance changes over time, we need to model how such changes occur.

Given an estimate of the expected value of fund i 's alpha in period $t+1$ based on information available at time t , $\hat{\alpha}_{i,t+1|t}$, we need to determine if $\hat{\alpha}_{i,t+1|t}$ can predict whether the fund subsequently outperformed. To this end, we consider the product of $\text{Max}(\hat{\alpha}_{i,t+1|t}, 0)$ and the sign of the fund's actual risk-adjusted performance, $\text{sign}(R_{i,t+1} - \hat{\boldsymbol{\beta}}'_{i,t+1} \mathbf{z}_{t+1})$:

$$P_{i,t+1} = \text{Max}(\hat{\alpha}_{i,t+1|t}, 0) \times \text{sign}(R_{i,t+1} - \hat{\boldsymbol{\beta}}'_{i,t+1} \mathbf{z}_{t+1}). \quad (2)$$

Here the sign function $\text{sign}(\bullet)$ equals +1 if the argument is positive, -1 if the argument is negative, or 0 if the argument is zero. To motivate the objective in (2), note that $P_{i,t+1}$ will be large for funds with large, positive predicted alphas ($\hat{\alpha}_{i,t+1|t} > 0$) whose subsequent risk-adjusted returns are positive ($R_{i,t+1} - \hat{\boldsymbol{\beta}}'_{i,t+1} \mathbf{z}_{t+1} > 0$). Conversely, the objective in (2) penalizes funds with large positive alpha estimates but negative future risk-adjusted returns. Thus, the "predictive alpha" measure in (2) accounts for both the magnitude and sign of the predicted performance (through $\text{Max}(\hat{\alpha}_{i,t+1|t}, 0)$), and for the success of the forecast (through the product of the two terms). By multiplying the predicted alpha by the sign of realized excess returns, the predictive alpha preserves

¹⁴To see the importance of this point, consider a fund that takes high idiosyncratic risks and has produced a high average (historical) risk-adjusted performance due to a single high monthly return ("outlier"). This would not instill much confidence in the fund's ability to produce high future returns. Contrast this with another fund with a lower average risk-adjusted return but consistent good performance; this fund is likely to be more attractive to investors, particularly if the periods when it outperforms can be predicted.

scale, i.e., it is measured in units of returns (basis points). Finally, note that funds with negative predicted alphas ($\hat{\alpha}_{i,t+1|t} < 0$) get excluded from consideration, since $P_{i,t+1} = 0$ for such funds and so (2) is intended for identifying superior performance. We explain how to identify inferior performance in Section 4.5 below.

To select funds with the highest expected value of the predictive alpha, for each fund i we measure the past success in predicting this fund's performance up to time t , using the sample estimate of the average value of (2):

$$\bar{P}_{i,t} = \frac{1}{t - t_{i0}} \sum_{\tau=t_{i0}+1}^t P_{i,\tau} = \frac{1}{t - t_{i0}} \sum_{\tau=t_{i0}+1}^t \text{Max}(\hat{\alpha}_{i,\tau|\tau-1}, 0) \text{sign}(R_{i,\tau} - \hat{\beta}'_{i,\tau} \mathbf{z}_\tau). \quad (3)$$

Here $\hat{\alpha}_{i,\tau|\tau-1}$ is the forecast of $\alpha_{i,\tau}$ based on information available at time $\tau - 1$, and $\hat{\beta}'_{i,\tau}$ are least-squares estimates of β'_i , using only data up to time τ . And, $t_{i0} + 1$ is the starting point of the sample used to estimate the i th fund's predictive alpha performance up to time t , $\bar{P}_{i,t}$.

2.2. Fund Confidence Set

Our data contain more than 2,000 funds whose performance needs to be pair-wise compared at each point in time. This introduces a complicated multiple hypothesis testing problem, which we address by applying the model confidence set (MCS) approach of Hansen, et al. (2011). That approach is designed to choose the set of “best” forecasting models from a larger set of candidate models. Because the approach is developed for the selection of forecasting models, we need to modify it for our setting. Most obviously, the object of interest in our analysis is not a model, but a model applied to a fund's return performance, and, so, we label our approach the Fund Confidence Set (FCS). We next describe how the approach works.

Our goal is to select a set of funds which, at a certain level of confidence, contains the best fund—or set of funds, if multiple funds are believed to have identical performance. The approach relies on an equivalence test and an elimination rule. Let $\mathcal{F}_t^0 = \{F_{1t}, \dots, F_{nt}\}$ be the initial set of funds under consideration while P_{it} , given by equation (2), measures fund i 's performance in period t . The difference between the performance of funds i and j at time t is then

$$d_{ij,t} = P_{i,t} - P_{j,t}, \quad i, j \in \mathcal{F}_t^0. \quad (4)$$

Defining $\mu_{ij} = E[d_{ij,t}]$ as the expected difference in the performance of funds i and j , we prefer fund i to fund j if $\mu_{ij} > 0$; both funds are judged to be equally good if $\mu_{ij} = 0$. The set of “superior” funds at time t , \mathcal{F}_t^* , consists of those funds that are not dominated by any other funds in \mathcal{F}_t^0 , i.e., $\mathcal{F}_t^* = \{i \in \mathcal{F}_t^0 : \mu_{ij} \geq 0 \text{ for all } j \in \mathcal{F}_t^0\}$. The FCS approach identifies \mathcal{F}_t^* by means of a

sequence of tests, each of which eliminates the worst fund—the fund deemed to be worst relative to another fund in the current set of surviving funds, \mathcal{F}_t —provided that this difference is statistically significant. Each round of this procedure tests the null hypothesis of equal performance

$$H_{0,\mathcal{F}_t} : \mu_{ij} = 0, \text{ for all } i, j \in \mathcal{F}_t \subset \mathcal{F}_t^0, \quad (5)$$

against the alternative hypothesis that the expected performance differs for at least two funds:

$$H_{A,\mathcal{F}_t} : \mu_{ij} \neq 0 \text{ for some } i, j \in \mathcal{F}_t. \quad (6)$$

Following Hansen et al. (2011), we define the Fund Confidence Set (FCS) as any subset of \mathcal{F}_t^0 that contains \mathcal{F}_t^* with a certain probability, $1 - \lambda$, where λ is the size of the test.

With these definitions in place, we next explain how the algorithm for constructing the FCS works. The first step sets $\mathcal{F}_t = \mathcal{F}_t^0$, the full list of funds under consideration at time t . The second step uses an equivalence test to test $H_{0,\mathcal{F}_t} : E[d_{ij}] = 0$ for all $i, j \in \mathcal{F}_t^0$ at a critical level λ . If H_{0,\mathcal{F}_t} is accepted, the FCS is $\hat{\mathcal{F}}_{1-\lambda,t}^* = \mathcal{F}_t$. If, instead, H_{0,\mathcal{F}_t} gets rejected, the elimination rule ejects one fund from \mathcal{F}_t , and the procedure is repeated on the reduced set of funds. The procedure continues until the equivalence test does not reject, and, so, no additional funds need to be eliminated. The remaining set of funds in this final step is $\hat{\mathcal{F}}_{1-\lambda,t}^*$.

Elimination of individual funds is based on their relative sample performance, using measures such as $\bar{P}_{i,t}$ from (3). Specifically, from (3) and (4) we can estimate the performance of fund i relative to fund j as $\bar{d}_{ij} = t^{-1} \sum_{\tau=1}^t \widehat{d}_{ij,\tau}$. To obtain a better behaved test statistic, we divide this measure by its standard error, $\sqrt{\widehat{\text{var}}(\bar{d}_{ij})}$, to obtain¹⁵

$$t_{ij} = \frac{\bar{d}_{ij}}{\sqrt{\widehat{\text{var}}(\bar{d}_{ij})}}, \quad (7)$$

As in Hansen, et al. (2011), we base the test of H_{0,\mathcal{F}_t} on the range statistic

$$T_{\mathcal{F}_t} = \max_{i,j \in \mathcal{F}_t} |t_{ij}|, \quad (8)$$

which finds the largest t —statistic chosen among the many pairwise t —tests in (7).¹⁶

Under assumptions listed in Hansen, et al. (2011), the set of pair-wise t —tests, t_{ij} , are asymptotically pivotal and can be expected to have better sampling properties.

¹⁵Also, by scaling \bar{d}_{ij} by its standard error, the resulting test statistic is pivotal and can be expected to have better sampling properties.

¹⁶Because we are maximizing performance, we need to alter the range elimination rule to $\arg \min_{i \in \mathcal{F}_t} \sup_{j \in \mathcal{F}_t} t_{ij}$ so that the the fund (i) with the worst estimated performance measured relative to some other fund (j) is the candidate for elimination.

totically joint normally distributed with unknown covariance matrix, Ω . Because so many pairwise test statistics are being compared and Ω is unknown, the resulting test statistic, T_{R,\mathcal{F}_t} , has a non-standard asymptotic distribution whose critical values can be bootstrapped using the approach of White (2000). From these draws, the sequential elimination rule is used to purge any fund whose performance looks sufficiently poor relative to that of at least one other fund currently included in the FCS.¹⁷

Because a random number of possibly dependent tests are carried out, it is far from trivial to control the coverage probability of this step-wise procedure. Notably, if each round conducts a test at a fixed critical level, λ , then the final FCS will have a very different coverage probability than $1 - \lambda$. A key contribution of Hansen et al. (2011) is to design a sequential procedure that can be used to control the coverage probability, $1 - \lambda$, of the FCS. Specifically, Theorem 1 in Hansen et al. (2011) establishes conditions under which the probability that a truly superior fund is captured in the estimated FCS is greater than or equal to $1 - \lambda$, while the probability of wrongly including an inferior fund in $\hat{\mathcal{F}}_{1-\lambda,t}^*$ asymptotically goes to zero.

2.3. Choosing λ

As with any inference problem, the FCS approach requires us to trade off type I and type II errors. Type I errors (false positives) are incorrect rejections of a true null, i.e., wrongly eliminating funds whose performance is as good as that of the best fund. Type II errors, conversely, are failures to reject a false null hypothesis, i.e., failing to exclude a poor fund from the FCS. How these errors are traded off gets regulated by the choice of the size of the equivalence test (λ) used by the FCS approach, which, therefore, is an important parameter.

Setting λ high means reducing the probability of wrongly including inferior funds (i.e., increasing the power of the equivalence test), but also implies that we stand a reduced chance of including funds with truly superior performance. Conversely, setting λ low means increasing the probability of including both truly inferior and truly superior funds as we become more cautious about eliminating individual funds, and the algorithm becomes less selective.

If the estimated performance of many of the funds is quite noisy, tests of the equivalence hypothesis in (5) may not be very powerful, and the algorithm will eliminate too few funds, resulting in a bloated set that includes many inferior funds.¹⁸ This would simply reflect that the data are not sufficiently informative to distinguish between the performance of different funds. We can easily imagine economic states in which this would plausibly be the case, such as the economic crisis of 2008/2009, where market volatility and fund idiosyncratic volatility were quite high. Conversely,

¹⁷Specifically, if the FCS p -value for the fund identified by (8) is smaller than the λ -quantile of the bootstrapped distribution, then this fund is deemed inferior to at least one other fund and gets eliminated.

¹⁸In other words, funds can avoid being eliminated from the FCS either because they have a high average performance which is precisely estimated, or, alternatively, because their performance is imprecisely estimated, if their alpha is surrounded by large standard errors which reduces t_{ij} in (7).

when the data are informative and allow for sharper inference, the equivalence tests first eliminate the poor funds before questioning the superior funds.

In trading off these effects, our goal is to create a set of funds with the highest level of future composite risk-adjusted performance, with a constraint of reducing, to the largest extent possible, the risk of underperformance. We believe that this goal is especially useful for retirement plan fiduciaries, who seek to provide, with a high level of statistical confidence, a menu of the most promising funds for their investors, as well as fiduciaries who select managers for pension funds, endowments, sovereign wealth funds, and other trusts.

We, therefore, opt for a relatively high value of λ , choosing $\lambda = 0.90$ as our benchmark. However, to illustrate how sensitive our results are to our particular choice of λ , in robustness tests, we also consider two alternative values ($\lambda = 0.50, 0.10$) which result in fewer funds being eliminated. We refer to the three sets of λ -values as tight ($\lambda = 0.90$), medium ($\lambda = 0.50$), and wide ($\lambda = 0.10$).

2.4. Choice of Candidate Set of Funds

We implement the FCS approach as follows. First, because there are more than 2,000 funds in our sample, it is not feasible to conduct all possible pairwise performance comparisons.¹⁹ To handle this issue, we propose a pre-screening step that first (before implementing the FCS) excludes funds with poor or unpredictable performance—either is detrimental from an investment perspective. This step greatly reduces the complexity of the subsequent FCS computation step. Specifically, our baseline results restrict the initial set of funds being considered for inclusion in the FCS (\mathcal{F}_t^0) to the top 15% of funds as ranked by their alpha estimates, $\hat{\alpha}_{i,t+1|t}$. Because this cutoff is somewhat arbitrarily chosen, for robustness, we also consider cutoffs that use either 25% or 5% of the funds ranked by alpha estimates in the initial round.²⁰

In addition, we consider a strategy that endogenously chooses the initial set of funds (\mathcal{F}_t^0) by identifying funds whose predictive alpha estimates are significantly higher than zero. This is accomplished using the step-wise procedure of Romano and Wolf (2005) which aims to identify as many funds with genuinely positive alphas as possible, while controlling the familywise error (FWE) rate. The null and alternative hypotheses tested by the Romano-Wolf approach are $H_{0i} : \alpha_i \leq 0$ vs. $H_{1i} : \alpha_i > 0$, and the FWE rate is defined as the probability of wrongly rejecting the null for at least one fund, i.e., $FWE = \text{prob}(\text{rejecting at least one } H_{0i} \text{ for which } \alpha_i > 0)$. The challenge is to design an approach whose asymptotic FWE is no greater than some critical level, while accounting for the multiple hypothesis testing problem arising from comparing so many mutual funds.²¹

¹⁹For example, a cross-section of 2,200 funds observed over 270 months gives rise to just over 650 million pairwise comparisons.

²⁰We estimate the fund confidence set using the MulCom 3.0 package for Ox; see Hansen and Lunde (2010) and Doornik (2006).

²¹To implement the Romano-Wolf procedure, define a test statistic, $test_i$, for testing H_{0i} vs. H_{1i} . Suppose we have renumbered the funds $i = 1, \dots, n$ by the magnitude of their individual test statistics, $test_n \leq test_{n-1} \leq \dots \leq test_1$. A

3. Performance Measurement Models

This section introduces the models used to estimate the risk-adjusted performance of individual mutual funds. The performance measurement models play an important role for the results. As we show below, a conditional alpha approach that uses both returns and holdings data can provide sharper inference about alphas than an approach which uses only returns data. Sharper inference on alphas translates into an improved ability to discriminate between truly skilled and truly unskilled fund managers.

In common with much of the existing literature on mutual fund performance, we use a four-factor model that, in addition to the market factor, adjusts for the size and value factors of Fama and French (1992) and the momentum factor of Carhart (1997). However, we generalize this model in two ways. First, following Mamaysky et al. (2007, 2008), we assume that managers receive information (unobserved to the econometrician) that is correlated with future returns. As we show below, such information gives rise to a time-varying component in fund performance. Second, we show how to generalize this framework to combine information from past returns with holdings data to more accurately extract an estimate of fund performance.

3.1. Benchmark Model

Our benchmark specification is a four-factor model with constant alpha and constant loadings on the risk factors. Recall that R_{it} is the monthly return on fund i , measured in excess of a 1-month T-bill rate. Similarly, let $\mathbf{z}_t = (R_{mt}, SMB_t, HML_t, MOM_t)'$ denote the values of the four risk factors, where R_{mt} is the return on the market portfolio in excess of a 1-month T-bill rate, SMB_t and HML_t are the small-minus-big equity market capitalization and value-minus-growth factors of Fama and French (1992), and MOM_t is the momentum factor of Carhart (1997), constructed as the return differential on portfolios comprising winner versus loser stocks, tracked over the previous 12 months. The benchmark four-factor model takes the form

$$R_{it} = \alpha_i + \beta'_i \mathbf{z}_t + \varepsilon_{it}. \quad (9)$$

Following common practice in the finance literature, we obtain estimates of $\boldsymbol{\theta}_i = (\alpha_i \ \boldsymbol{\beta}'_i)'$ using a rolling 60-month estimation window. Such estimates account for slowly-evolving shifts in mutual

critical value, c_1 , is then determined such that the set of n_1 funds with test statistics $test_{R_1} \geq c_1$ has a fixed coverage probability. Only funds with higher test statistics (i.e., funds numbered $1, \dots, n_1$) are set aside in this step. Next, the procedure is repeated on the remaining $n - n_1$ funds, resulting in a new critical value, c_2 , and we set aside those funds whose test statistics exceed c_2 . The procedure is repeated until we can no longer reject the null of no superior performance for any remaining funds, at which point the algorithm stops and consolidates the funds set aside in previous steps. We use the studentized StepM method (Algorithm 4) of Romano and Wolf (2005) in the analysis. Like the MCS approach, the Romano-Wolf method uses the White (2000) bootstrap to calculate the critical values used for eliminating funds.

fund performance and risk exposures.²²

3.2. Time-varying Skills

Different approaches have been suggested for capturing time variation in fund alphas. In this section, we follow the approach of Mamaysky et al. (2007, 2008) and show how it can be generalized—and improved—to take advantage of information from mutual fund holdings data.

Individual funds' alphas are value-weighted averages of their stock-level alphas, so our analysis starts with the performance of individual stocks. Specifically, we decompose the excess return of each stock, r_{jt} , into a risk-adjusted return component, α_{jt} , a systematic return component obtained as the product between a set of risk exposures, β_j and factor returns, \mathbf{z}_t , and an idiosyncratic return component, ε_{jt} . We stack these return components into $N_t \times 1$ vectors $\boldsymbol{\alpha}_t$ and $\boldsymbol{\varepsilon}_t$ and an $N_t \times 4$ matrix of betas, $\boldsymbol{\beta}$, where N_t is the number of stocks in existence at time t . Notice that the individual stock alphas are allowed to vary over time, reflecting that any abnormal returns are likely to be temporary.²³ Following Mamaysky et al. (2008), we assume that stock betas are constant, although this assumption can be relaxed.

Mutual fund returns can be computed by summing across individual stock returns, \mathbf{r}_t , weighted by the fund's ex-ante portfolio weights at the end of the previous period, $\boldsymbol{\omega}'_{t-1}$. Using the decomposition of individual stock returns described above, the excess return of an individual mutual fund (i) net of the risk-free rate, transaction costs, fees and expenses (k_i) is given by:

$$\begin{aligned} R_{it} &= \boldsymbol{\omega}'_{it-1}(\boldsymbol{\alpha}_t + \boldsymbol{\beta}\mathbf{z}_t + \boldsymbol{\varepsilon}_t) - k_i, \\ &\equiv \alpha_{it} + \beta'_{it}\mathbf{z}_t + \varepsilon_{it}, \end{aligned} \tag{10}$$

where $\alpha_{it} = \boldsymbol{\omega}'_{it-1}\boldsymbol{\alpha}_t - k_i$, $\beta_{it} = \boldsymbol{\omega}'_{it-1}\boldsymbol{\beta}$, and $\varepsilon_{it} = \boldsymbol{\omega}'_{it-1}\boldsymbol{\varepsilon}_t$.²⁴

Mamaysky et al. (2008) represent skills through a persistent, latent process which reflects the fund's ability to process private information.²⁵ Suppose that fund i receives a private signal, F_{it} , which follows an autoregressive process,

$$F_{it} = \nu_i F_{it-1} + \eta_{it} \quad \text{for } \nu_i \in [0; 1]. \tag{11}$$

where η_{it} is independently and identically distributed (i.i.d.) with zero mean. The innovations, ε_{it} ,

²²Empirically, we find that five-year rolling window estimates are slightly better at identifying and predicting fund performance than estimates based on an expanding estimation window. The results reported below are, however, robust to using an expanding estimation window.

²³Liu and Timmermann (2013) develop a theoretical model with temporary abnormal returns (alphas) in the context of convergence trading.

²⁴As in Mamaysky et al. (2008) these are assumed to be proportionate to the fund's assets under management.

²⁵Kosowski (2011) models manager skills as a latent variable driven by a regime-switching process. As in Mamaysky et al. (2008), this gives rise to a filtering problem, although the filter is non-linear in this case.

in (10) and η_{it} in (11) are assumed to be mutually independent and normally distributed. Following Mamaysky et al. (2008), assume that (i) fund portfolio weights are linear in the private signal, $\omega_{it-1} = \bar{\omega}_i + \gamma_i F_{it-1}$; (ii) there is a constant mapping $\bar{\alpha}_i$ from fund i 's private signal into its ability to identify abnormal performance, $\alpha_{it} = \bar{\alpha}_i F_{it-1}$. These assumptions and equation (10) imply that fund i 's beta and alpha are linear and quadratic functions of the signal, F_{it-1} , respectively:

$$\begin{aligned}\alpha_{it} &= \bar{\omega}'_i \bar{\alpha}_i F_{it-1} + \gamma'_i \bar{\alpha}_i F_{it-1}^2 - k_i \\ &\equiv a_i F_{it-1} + b_i F_{it-1}^2 - k_i,\end{aligned}\tag{12}$$

$$\begin{aligned}\beta_{it} &= \bar{\omega}'_i \beta + \gamma'_i \beta F_{it-1} \\ &\equiv \bar{\beta}_i + \mathbf{c}_i F_{it-1}.\end{aligned}\tag{13}$$

The fund manager's signal, F_{it-1} , is unobserved to the econometrician, but an estimate of it can be obtained from the fund's observed returns. To this end, we put the model into state space form:

$$\begin{aligned}R_{it} &= a_i F_{it-1} + b_i F_{it-1}^2 - k_i + (\bar{\beta}'_i + \mathbf{c}'_i F_{it-1}) \mathbf{z}_t + \varepsilon_{it} \\ F_{it} &= \nu_i F_{it-1} + \eta_{it}.\end{aligned}\tag{14}$$

We refer to this as the latent skill (LS) model. As explained by Mamaysky et al. (2008), the parameters of this model can be estimated fund-by-fund using an extended Kalman Filter that accounts for the presence of the squared value of the state variable, F_{it-1} , in equation (14).²⁶

3.2.1. Introducing Information from Fund Holdings

Conventional approaches to measuring mutual fund performance base their inference on time-series estimates using past returns, which can be very noisy. This limits the ability of return-based methods—such as that of Mamaysky et al. (2008)—to identify funds with superior performance.

One way to address this issue is by augmenting the model to account for information from portfolio holdings which, as is clear from equation (12), can be used to track how a fund's alpha evolves through time. Building on this idea, we next generalize the methodology in Mamaysky et al. (2008) and show that holdings-based information can be added in the form of an additional measurement equation in the state space representation of the model (14). Moreover, this model can be estimated by means of an extended Kalman filter.

Data on fund holdings allow us to perform risk-adjustment at the individual stock level by matching each stock to a portfolio of stocks with similar characteristics in terms of their sensitivity to book-to-market, market capitalization, and price momentum factors. The difference between an individual stock's return and the return on its characteristics-matched portfolio can be used as a

²⁶Normalizing one of the elements of \mathbf{c}_i , the four-factor model requires 13 parameters to be estimated.

measure of that stock's abnormal return. Weighting individual stock abnormal returns across all stock positions held by a fund, we obtain the fund-level characteristic-selectivity (CS) measure of Daniel, et al. (1997),

$$CS_{it} = \omega'_{it-1}(\mathbf{r}_t - \mathbf{r}_{bt}). \quad (15)$$

Here, \mathbf{r}_t and \mathbf{r}_{bt} are vectors of excess returns on stocks (\mathbf{r}_t) and benchmark portfolios (\mathbf{r}_{bt}), respectively. Benchmark portfolios are chosen to match, as closely as possible, the characteristics of the individual stocks and so $\beta_{it} = \beta_{bt}$. Because the characteristic-matched stocks are chosen mechanically, the average stock can be expected to have zero alpha, $\alpha_{bt} = \mathbf{0}$. Using (12) and (15) we therefore have

$$\begin{aligned} CS_{it} &= \omega'_{it-1} (\alpha_t + \beta' \mathbf{z}_t + \varepsilon_t - (\alpha_{bt} + \beta'_b \mathbf{z}_t + \varepsilon_{bt})) \\ &= \omega'_{it-1} (\alpha_t - \alpha_{bt}) + \omega'_{it-1} (\beta - \beta_b)' \mathbf{z}_t + \omega'_{it-1} (\varepsilon_t - \varepsilon_{bt}) \\ &= \alpha_{it} + k_i - \omega'_{it-1} \alpha_{bt} + (\beta_{it} - \beta_{bt})' \mathbf{z}_t + \varepsilon_{it} - \varepsilon_{bt} \\ &= a_i F_{it-1} + b_i F_{it-1}^2 + \varepsilon_{it} - \varepsilon_{bt}. \end{aligned} \quad (16)$$

Since the CS measure does not depend on estimated risk factor loadings obtained over some prior historical period, it has the potential to generate a more accurate estimate of fund performance and, thus, to improve inference based on the return-only models.²⁷

The generalized latent skill, holding-based (LS-Holding) model can be summarized as

$$\begin{pmatrix} R_{it} \\ CS_{it} \end{pmatrix} = \begin{pmatrix} a_i F_{it-1} + b_i F_{it-1}^2 - k_i \\ a_i F_{it-1} + b_i F_{it-1}^2 \end{pmatrix} + \begin{pmatrix} \bar{\beta}_i + \mathbf{c}_i F_{it-1} \\ 0 \end{pmatrix} \mathbf{z}_t + \begin{pmatrix} \varepsilon_{it} \\ \varepsilon_{it} - \varepsilon_{bt} \end{pmatrix} \quad (17)$$

$$F_{it} = \nu_i F_{it-1} + \eta_{it}$$

Compared to the model of Mamaysky et al. (2008) in (14), the LS-Holding model uses the additional information contained in CS_{it} , which has the potential to make estimation and extraction of the fund's private information component, F_{it} , more precise. We can again estimate the parameters of (17) using the extended Kalman filter, now using two measurement equations; Appendix A provides further details.²⁸

²⁷The CS measurement equation is “idealized,” in the sense that the CS performance measure will not capture intra-quarter trading returns. Note that the CS and return-based measures can be viewed as different estimates of the same underlying fund performance, observed with different measurement errors.

²⁸We note an additional important difference between our approach and that of Mamaysky, et al. (2008). Before forming portfolios based on the conditional alpha estimates, Mamaysky, et al. (2008) trim the set of funds. By contrast, in our approach, funds that are eligible for inclusion at a given point in time are assigned to an “active pool”, while excluded funds are assigned to a “passive pool.” Funds can enter the passive pool at any point in time and return to the active pool again. The funds are allocated to the pools following two steps. First, funds whose

4. Performance Results

This section introduces our data and establishes a performance benchmark based on the popular decile sorting methodology widely used in academic studies. Next, we analyze the empirical performance of portfolios based on funds included in the fund confidence sets.

4.1. Data

Our empirical analysis uses monthly returns on a sample of U.S. equity mutual funds over the 32-year period from 1980:06 to 2012:12. Individual fund returns are taken from the CRSP mutual fund data base, and are net of transaction costs and fees.

To construct our estimate of the CS measure, we use quarterly holdings data from Thomson Financial.²⁹ We use these quarterly holdings data to construct a three-month estimate of CS . We merge the returns and holdings data using the MFLINKS files of Wermers (2000), which have been updated by the Wharton Research Data Services, and allows us to map the Thomson holdings data to CRSP returns using the funds' WFICN identifier. We require each fund to have at least six months of contiguous returns data. In total, we have returns and holdings data on 2,480 actively managed U.S. domiciled equity mutual funds, but we exclude 255 sector funds, and, so, end up with 2,225 funds.³⁰ The number of funds included in the analysis peaks at above 1,700 in 2009 before declining to around 1,500 in 2012.

For each fund, we obtain an alpha estimate using time-series data on the fund's historical returns. Funds with a very short return record tend to generate noisy alpha estimates. To avoid that our analysis becomes dominated by such funds, we require funds to have a return record of at least five years. Table 1 provides summary statistics for the cross-sectional distribution of individual fund (ex-post) alphas, using each of the three performance models described in the previous section. The median fund has a negative alpha, ranging from -61 bps/year to -74 bps/year across the three models. The finding that the median fund underperforms on a risk-adjusted basis is consistent with previous academic studies (e.g., Carhart, 1997). The bottom 5% of funds, ranked by alpha performance, have a negative alpha estimate of about -35 bps/month, or just under -4%/year—a number that, again, does not vary much across the three model specifications. The top 5% of funds have alpha estimates ranging from 21 bps/month to 25 bps/month, approximately 3%/year; these

alpha forecast for the previous month had the same sign as the fund's return in excess of the return on the market portfolio during that month, stay in the active pool for the current period. Second, funds with alpha forecasts less than -200 bps/month or greater than 200 bps/month, or funds whose predicted betas are less than zero or greater than two are moved to the passive pool. Our approach does not require that we assign funds to such active or passive pools.

²⁹Prior to 2004, funds were only required to report holdings every six months. However, the majority of funds disclosed holdings every three months to Thomson. See Wermers (1999) for details.

³⁰"Sector funds" are defined as funds whose R^2 is less than 70% in a four-factor regression. For such funds, the simple four-factor risk-adjustment approach is not appropriate, and these funds are, therefore, excluded.

estimates are again quite similar across the three models.

4.2. Performance of Decile-sorted Portfolios

To establish a reference point for our FCS results, we first follow the common practice of ranking individual funds' alphas and forming decile portfolios. This approach can be used to see if funds that are expected to have the highest alphas do indeed produce better subsequent performance than lower-ranked funds. Specifically, each month, t , we rank funds based on their expected alphas $\hat{\alpha}_{i,t+1|t}$. We then form ten equal-weighted decile portfolios with the first portfolio (P1) containing the bottom 10% alpha-ranked funds, the next decile containing funds ranked in the second-lowest 10%, continuing up to the top 10% of alpha-ranked funds (P10). To obtain more detailed results for the bottom and top funds, we also divide P1 into the bottom 5% alpha-ranked funds and funds ranked between the bottom 5% and 10% (labeled P1A and P1B, respectively); we use a similar split for the P10 portfolio (P10A and P10B). Finally, we record the returns on each of these portfolios over the subsequent month. Each month, as new data arrive, we repeat this sorting routine, form equal-weighted portfolios based on the funds' updated alpha estimates, then, again, record their next-month returns. We use five years of data to initiate the portfolio sorts, and another five years to obtain an estimate of the predictive alpha and, thus, generate a time series of portfolio excess returns, R_{pt} , over the 21-year period from July 1, 1990 to December 31, 2012. This approach, while requiring 10 years of returns history for each fund under consideration, does not induce substantial survival bias, as we do not require funds to survive beyond a single month into the future.³¹

To evaluate the performance of the portfolios, we follow conventional practice and estimate four-factor alphas on the (out-of-sample) portfolio returns,

$$R_{pt} = \alpha_p + \beta'_p \mathbf{z}_t + \varepsilon_{pt}, \quad t = 1, \dots, T. \quad (18)$$

The resulting estimates, $\hat{\alpha}_p$, can be interpreted as the “average” portfolio alphas.

Table 2 presents alpha estimates for the decile portfolios. We find strong evidence of negative and statistically significant alpha estimates for the bottom three ranked decile portfolios (P1-P3). The underperformance of these decile portfolios ranges from -8 bps/month to -21 bps/month, and is qualitatively similar across the three different models used to rank funds. The alpha estimates of the top-ranked decile portfolio (P10) range between 2 and 6 bps/month, and are statistically insignificant, indicating that the conventional portfolio sorting approach fails to identify funds with abnormal positive performance. This conclusion carries over from the top 10% to the top 5% portfolio of funds (P10B), which only perform marginally better than the P10 portfolio.

Despite these shortcomings, the portfolio sorting approach does succeed in differentiating be-

³¹See also Carhart (1997) for a discussion of survivorship bias.

tween the best and worst performing funds as the estimated differential in alphas between the P10 and the P1 portfolios is large and positive (18-26 bps/month) and highly statistically significant. This finding is, however, driven mostly by the identification of funds with inferior performance. As an alternative test of the value of the ranking information in the portfolio sorts, the last line in Table 2 reports the MR test for a monotonic pattern in the alphas proposed by Patton and Timmermann (2010). The null is that there is a flat or declining pattern in the alphas, while the alternative is that there is a monotonically increasing pattern as we move from $P1$ to $P10$. Small p -values for this test are evidence that a model succeeds in ranking the future risk-adjusted performance of the funds. The test statistic generates a p -value of 0.07 when applied to the portfolios ranked by the constant-alpha benchmark model, and is statistically significant for the LS and LS-Holdings models. Hence, these performance models contain valuable ranking information despite their failure to identify funds with large positive alphas.

We conclude from these findings that the conventional portfolio sorting approach can be used to identify a broad set of funds with inferior performance, but is not well suited for identifying funds with superior future performance.

4.3. Performance of Superior FCS Funds

We next analyze the performance of our FCS approach, which forms portfolios from the set of funds identified as top performers. Figure 1 shows how the FCS approach selects funds that stand out even among funds with positive forecasts of alpha. The figure plots the distribution of predictive alpha estimates obtained using the LS-holdings model at a single point in time (July 2006). The black curve shows the distribution of predictive alphas for the full set of funds in existence at that point, i.e., the population of funds. This curve is centered to the left of zero, and has a wide dispersion. We also show the distribution of predictive alphas for funds with positive predicted alphas (green), and funds whose alpha estimates are positive and statistically significant using the Romano-Wolf approach (red). Finally, the blue line captures the distribution of predictive alphas for funds included in the FCS. Note that the FCS curve lies to the right of the distribution curve for the more inclusive set of funds with significantly positive predictive alphas, highlighting that the FCS methodology is far more discriminating than an approach that simply selects funds with (significantly) positive alphas.

Turning to the formal analysis, Table 3 presents alpha estimates for a set of portfolios formed by equal- or value-weighting the funds selected each month by applying the FCS approach to the available set of candidate (top fractile ranked) funds. As in Table 2, the (out-of-sample) period goes from 1990:07 to 2012:12. We consider a range of values for the significance level of the test, $\lambda = \{0.10, 0.50, 0.90\}$, corresponding to a wide, medium, and tight FCS, and present results for the three performance models introduced in Section 3. Our baseline analysis (shown in Panel A)

assumes a tight FCS, and uses the top 15% of funds as the initial set of candidate funds. We consider both equal- and value-weighted portfolios of funds, but focus our discussion on the former. Results for the alpha-ranked top decile of funds (P10) are shown for comparison and serve as a benchmark given the widespread use of this approach.

For the constant-alpha benchmark model, moving from the top decile (P10) portfolio to the tight FCS portfolio increases the average alpha from 6 bps/month to 30 bps/month, although the latter estimate is not significantly different from zero. Similar, if somewhat stronger, results hold for the LS model estimated only on returns data whose FCS alpha estimate is 25 bps/month, i.e., 20 bps/month higher than that of the top decile (P10) portfolio.

We obtain very strong performance results for the latent skill model that combines return and holdings data (LS-Holdings). For this model, the tight FCS portfolio produces a statistically significant alpha estimate that averages 54 bps/month.

As we reduce λ from 0.90 to 0.50 or 0.10, and so move from the tight to the medium and wide set of funds, fewer funds get eliminated and the performance of the FCS portfolios is markedly reduced. This suggests that the additionally included funds tend to perform worse than the funds identified by the more selective approach. This finding is unsurprising, because many funds have alphas whose estimates are surrounded by large standard errors. For these funds, choosing a small value of λ means that the equivalence test used to eliminate inferior funds has insufficient power to reject the null that the funds' alphas are identical; we analyze this finding in more detail below.

We obtain similar results if we use value weights rather than equal weights to form the top FCS portfolios. For example, the alpha estimates generated by the tight FCS portfolio based on the three performance models change to 36, 25, and 52 bps/month (previously 30, 25, and 54 bps/month using equal weights).

Recall that we restrict the initial set of funds (\mathcal{F}_t^0) used in the FCS analysis in order to limit the number of pair-wise comparisons conducted by our approach. The results in Panel A assume that we apply the FCS approach to the top 15% of funds ranked on their alpha estimate. To explore the sensitivity of our results to starting either with a narrower (5%) or wider (25%) set of funds, panels B and C in Table 3 report results for these alternative initial ranked sets of funds. For the LS-Holdings models we see, again, that the performance is better for the FCS portfolios than for the simple ranked portfolios, and that the performance tends to get better as we narrow the set of funds in the FCS, i.e., apply a higher value of λ .

Panel D in Table 3 considers the performance of the strategy that first uses the step-wise bootstrap method of Romano and Wolf (2005) to identify funds with significantly positive alphas, followed by the use of the FCS approach to select the set of top performers among these funds. Again, we see that the FCS portfolios generate higher alpha estimates than the top decile (P10) portfolio. Improvements are again particularly large for the LS-Holdings model whose FCS portfolio alpha (48 bps/month) is similar to that in Panel A. The portfolio of funds with significantly positive

alphas, selected by the Romano-Wolf approach, generates a four-factor alpha of 15 bps/month—far below that obtained from applying the FCS approach to the same set of funds (48 bps/month). This shows that there is considerable heterogeneity in performance among funds with significantly positive alphas.

We conclude the following from these results. First, the FCS approach is capable of selecting funds whose performance, when combined into equal- or value-weighted portfolios, is far better than that achieved by funds included in the conventional rank-based top-decile portfolio. Second, the performance of the FCS approach is best when we allow the approach to eliminate more funds and become more selective. This is consistent with our finding from the analysis of the decile portfolio’s performance that the ability to outperform is not very widespread among the funds in our sample. Third, the FCS approach works particularly well for the LS-Holdings model which allows for time-varying fund performance and combines returns and holdings data to provide sharper inference about alpha estimates.

4.4. Time Variation in Superior Funds’ Performance

The alpha estimates reported in Tables 2 and 3 show the average performance of a set of decile-sorted and FCS portfolios, respectively. Such estimates do not reveal if the performance of different funds is concentrated in certain states or occur over certain periods of time. Addressing this point can help provide us with further insights into the nature of the performance of the selected funds. Accordingly, we next explore this issue using a non-parametric methodology designed to capture “local” time variation in portfolio alphas. Appendix B provides further details of this approach.

Figure 2 shows the evolution through time in the risk-adjusted performance (alpha) of the superior funds included in the FCS portfolio (in blue), along with the performance of the traditional top decile portfolio (P10, in purple). The red line tracks the performance of the initial set of funds, selected using the top-15% ranked funds. Finally, the green line tracks the performance of the set of funds with positive predictive alphas, while the black line tracks the performance of a portfolio containing all funds in existence at a given point time.

From 1993-1995, the local alpha estimates of the FCS portfolio are either close to zero or positive. Starting in 1996, the performance of the FCS portfolio increases dramatically, before peaking at about 200 bps/month in 1999-2001. Over the same five-year period the performance of the P10 decile portfolio hovers around zero. The FCS portfolio only performs poorly during brief spells in 1990 and around 2004. After 2007, the FCS portfolio produces positive alphas of about 25-50 bps/month.

4.5. Technology Factor

Our results up to this point use a four factor model to compute the fund-level alpha forecasts, compare fund performances using the FCS methodology, and, finally, evaluate fund performance. One concern that is commonly raised when assessing results based on factor models is that they could be affected by an omitted risk factor. As we show in Section 5.1, there are times when the FCS of superior funds is quite concentrated and loads up on technology-heavy funds, so a concern here might be that our results are not robust to the presence of a technology factor.

To address this issue, and as a robustness check, we therefore include a fifth, technology index, factor in the FCS fund selection step ((3)) as well as in the performance evaluation step.³² We construct this fifth technology factor from the Fama-French 12-industry classification, using returns on Computers, Software, and Electronic Equipment. Funds that are selected using the five factor model in the FCS step go on to produce a four-factor alpha estimate of 34 bps (t -stat of 1.92) and a five factor alpha estimate of 36 bps (t -stat of 1.87).

As an alternative robustness exercise, we computed the non-parametric characteristics selectivity (CS) performance measure for the FCS portfolio that uses the five factor specification in the fund selection step. The average value of the CS measure is 33.5 bps per month with an associated p -value of 0.0001.

These results all suggest that including a technology factor has a relatively modest effect on our results.

4.6. Performance of Inferior FCS Funds

Next, we analyze whether the FCS approach can be used to identify funds with inferior performance. To this end we need to redefine the objective function. We do so by maximizing a “predictive alpha” loss function which focuses on funds in the left tail of the performance distribution:

$$L_{i,t} = \text{Min}(\hat{\alpha}_{i,t|t-1}, 0) \text{sign}(R_{i,t} - \beta'_{i,t} \mathbf{z}_t). \quad (19)$$

Recall that the sign function $\text{sign}(\bullet)$ equals -1 if its argument is negative. Hence, $L_{i,t}$ will be large for funds expected to perform poorly (through $\text{Min}(\hat{\alpha}_{i,t|t-1}, 0)$) with negative risk-adjusted returns, i.e., $\text{sign}(R_{i,t} - \beta'_{i,t} \mathbf{z}_t) = -1$. Funds with zero or positive expected alphas ($\hat{\alpha}_{i,t|t-1} > 0$) get excluded from consideration since $L_{i,t} = 0$ for such funds.

³²We do not, however, include the technology factor in the alpha forecasting step which takes 1-2 weeks of computation time due to the use of the extended Kalman filter.

Again, we base our analysis on a sample estimate of the mean performance for fund i ,

$$\bar{L}_{i,t} = \frac{1}{t - t_{i0}} \sum_{\tau=t_{i0}+1}^t \text{Min}(\hat{\alpha}_{i,\tau|\tau-1}, 0) \times \text{sign}(R_{i,\tau} - \hat{\beta}'_{i,\tau} \mathbf{z}_\tau). \quad (20)$$

As a first illustration of the inferior funds selected by the FCS approach, the bottom panel of Figure 1 shows the distribution of predictive alphas for the funds that were in existence (and had a minimum of five years of data) on July 1st, 2006. The distribution of predictive alphas for funds in the inferior FCS (blue line) is centered around -0.5%/year, and is much further to the left than the distribution curve for funds with negative alpha estimates (green line), funds with significantly negative alphas (red line), or for the population of funds as a whole (black curve). As in the case of funds with superior performance (top panel), the FCS approach is more selective in identifying funds with inferior performance than an approach that simply considers funds with negative expected performance.

Table 4 reports the four-factor risk-adjusted performance of the set of inferior funds identified by the FCS approach. The bottom decile of funds, as ranked by the constant-alpha benchmark model, generates a negative alpha estimate of -19 bps/month. The alpha estimate for the tight ($\lambda = 0.90$) equal-weighted FCS portfolio chosen from the top 15% alpha-ranked funds (Panel A) more than doubles in magnitude to -43 bps/month, which is highly statistically significant. Moreover, we find large and statistically significant (negative) alpha estimates for the worst funds across a broad range of values for the tightness parameter (λ).

These results show that the performance estimates for the funds with inferior performance are not as sensitive to the choice of λ as are our estimates for funds with superior performance. It is easy to explain this finding. As we reduce λ from 0.90 to 0.50 or 0.10, we widen the set of funds included in the FCS portfolios. Because poor performance is so widespread—Table 2 suggests that more than 30% of funds underperform—this has little effect on the portfolio comprising the worst funds. In contrast, good performance is less widespread, and, so, requires constructing a more concentrated portfolio (high λ) to be successfully identified.

Turning to the fund rankings based on the latent skill model, we find equally strong evidence that the FCS methodology succeeds in identifying funds with strongly negative performance. Specifically, the average alpha goes from -21 bps/month for the bottom decile of funds to -36 bps/month for the FCS portfolio with $\lambda = 0.90$. In both cases, the t -statistics are greater than three in absolute value, and, so, the results are both economically and statistically significant. A similar finding holds when we apply our approach to the LS-Holdings model. As shown in the last columns of Table 4, the average alpha performance drops from -16 bps/month for the bottom decile of funds to -32 bps/month for the tight FCS portfolio ($\lambda = 0.90$). Again, the results are quite robust to value-weighting, as opposed to equal-weighting, the funds in the FCS portfolio; using value weights and

$\lambda = 0.90$, the alpha estimates for the three performance models are -41, -27, and -26 bps/month, respectively (-43, -36, and -32 bps/month under equal weights).

The remaining panels in Table 4 show performance results when the FCS methodology is applied to the bottom 5% or 25% of alpha-ranked funds (Panels B and C, respectively). When applied to an even narrower (5%) initial set of funds, the FCS portfolio of inferior funds produces even lower performance than in Panel A, whereas the performance is slightly better for the wider set of initial funds (25%). Applying the FCS approach to an initial set of funds selected based on the Romano-Wolf stepwise procedure (Panel D), we obtain similar results to those in Panel A.³³

Note that the FCS approach works slightly better for the constant-alpha model than for the LS or LS-Holdings models applied to the bottom-ranked funds. Bottom funds tend to underperform because of high fees and transaction costs (Carhart, 1997) and not so much because they can systematically select poorly performing stocks (Kosowski, et al. (2006)). Hence, for inferior funds, there is not so much information to be learned from the holdings data.

Figure 3 plots the time-series evolution in the local alpha estimates for the inferior FCS portfolio. For this portfolio, we see substantially worse performance during the first five years of the sample followed by slightly better performance during 1998-2002 and, again, in 2012.

4.7. Fund Performance in High and Low Volatility States

Kacperczyk et al. (2016) analyze mutual funds' allocation of attention in a setting with multiple risky assets. Provided that the uncertainty about the aggregate risk factors is sufficiently large, their analysis shows that an increase in the variance of the aggregate risk factors (i) increases the dispersion in fund portfolio holdings; (ii) increases the dispersion in fund excess returns (Proposition 3 in Kacperczyk et al. (2016)); and (iii) increases the excess returns for skilled funds (Proposition 5). The intuition underpinning these results is that rewards from processing information are higher during times when asset payoffs are highly volatile, leading to stronger incentives for skilled investors to allocate extra attention to their portfolio decisions during such times. Conversely, unskilled investors can either not increase their attention (noise traders) or face higher costs from doing so, with the result that the differential between skilled and unskilled funds' return performance increases during periods of high volatility.

To test the first prediction (wider dispersion in portfolio holdings and returns), we compute the cross-sectional dispersion in industry exposure across the funds in our sample. At any point in time, the dispersion measure is computed relative to the average industry exposure across all funds. We then sort the sample into three groups, according to whether the aggregate (market)

³³A portfolio of funds with significantly negative alphas (selected using the Romano-Wolf approach) yields a four-factor alpha of -15 bps/month compared to the -32 bps/month figure obtained from applying the FCS approach to this same set of funds. Thus, there is considerable heterogeneity among funds with significantly negative performance.

volatility is low, medium or high, assigning one-third of the total sample to each volatility bin.³⁴

Panel A in Table 5 (left column) shows that the dispersion in industry concentration—computed using holdings in 48 industries—across funds is twice as high during periods with high aggregate risk compared to periods with low aggregate risk (1.023 vs. 0.464).³⁵ Moreover, the difference in mean industry concentration across states with high and low aggregate volatility is statistically significant with a *p*-value of 0.023.

Next, consider the second prediction implied by the analysis in Kacperczyk et al. (2016), namely that excess returns should be more disperse across funds during times of high volatility. Consistent with the prediction, the right column of Panel A in Table 5 shows that during periods with high aggregate risk the dispersion in the CS measure is two and a half times greater (6.09 versus 2.29) than during periods with low aggregate risk. Once again, this difference in dispersion across the high and low volatility states is statistically significant. Since the CS measure is constructed from funds' individual security holdings, this finding is consistent with the cross-sectional dispersion increasing during times of high volatility not simply due to greater differences in funds' industry allocations, but due to more dispersion in individual stock holdings.

Our FCS approach is ideally suited for testing the third prediction—that excess returns should be higher among the most skilled funds during times with high volatility—as it seeks to identify the set of funds with superior performance and thus, as we have seen, provides a direct estimate of the breadth of skill among mutual funds at a given point in time. Even so, it can be challenging to identify, ex-ante, outperforming active equity funds during periods of high stock market volatility. Although the greater separation in fund alphas implied by the theory should enhance our ability to identify superior funds, test procedures are hindered by a larger level of noise in each point estimate during periods with high volatility. Which of these two influences dominates depends on how much the skilled fund managers' performance increases during periods with higher volatility.

Panel B1 in Table 5 reports four-factor alpha estimates computed for the low, medium and high risk samples. The FCS alphas are notably higher (103 bps per month) during periods with high aggregate volatility (left columns) than during periods with low volatility (one bp/month), and their difference is highly statistically significant. An even stronger result holds when we sort by idiosyncratic volatility (right columns), where the alpha of the FCS portfolio is essentially zero during low-risk episodes but 118 bps/month during periods with high idiosyncratic risk. Interestingly, the portfolio of the top-10 percent of fund managers (P10) only weakly, and only when considering idiosyncratic risk, detects this relation between volatility and performance—in large part because it fails to account for the effect of high volatility on the breadth of performance skills.

Similar, if even stronger, results hold when we use a 36-month rolling window to determine

³⁴Our empirical tests of the relation between fund performance and aggregate market volatility all use a “local” measure of volatility, created using a 12-month rolling window volatility estimate and based on daily returns.

³⁵1.023 corresponds to an average deviation of an industry's weight in a given fund of about 14%, relative to the industry's average portfolio weight across funds, while 0.464 corresponds to a deviation of only about 10%.

the underlying volatility state. Using a 36-month window produces a less noisy estimate of the underlying volatility state, although it may also fail to detect short-lived durations of unusually high or low volatility. Using the longer 36-month window (Panel B2), we find that performance increases monotonically from 5 bps/month in the low aggregate risk state to 111 bps/month in the high (aggregate) volatility state or from 7 bps/month to 146 bps/month when sorting on idiosyncratic volatility.

We conclude from this evidence that performance among superior funds is notably higher during periods with high aggregate or idiosyncratic volatility and that this effect outweighs any increased noise that is introduced to estimated alphas during such periods. The FCS efficiently locates the set of outperforming active equity funds during the periods when they produce their greatest levels of alpha. In contrast, the conventional decile-sorting approach falls notably short at identifying the improved performance in high-volatility states.³⁶

5. Selection of Superior and Inferior Funds

The previous section shows that the FCS approach can be used to successfully identify funds with superior or inferior risk-adjusted performance. Moreover, the approach works particularly well for the LS-Holdings model that combines holdings and returns data to admit a sharper inference on fund alpha estimates. The FCS approach is fundamentally different from existing methods, such as decile sorting which keeps the proportion of “top” funds fixed through time. In contrast, the FCS approach endogenously determines how many funds to include. This section provides details of both the number of funds selected by the FCS approach and their identify.

5.1. Identifying Superior Funds

We can gain important economic insights by studying how the set of superior funds evolves over time. Specifically, the turnover in the set of funds identified to be superior is related to the skill level and duration of individual funds’ ability to outperform and the speed with which the competition catches up.

The top panel of Figure 4 shows the number of funds included in the FCS portfolio at each point in time. Because this set depends on the size of the test, λ , we present results for the narrow, medium, and wide sets, $\lambda = \{0.9, 0.5, 0.1\}$, but focus on the narrow set ($\lambda = 0.90$) in most of our discussion. The number of funds included in the top FCS fluctuates considerably over time. Around

³⁶Kacperczyk et al. (2016) compare fund performance during recession and expansion periods, noting that recessions tend to be periods with higher stock market volatility. We also considered splitting the data sample using the NBER recession indicator, but found that fund performance was imprecisely estimated, partly due to the small number of recession months in our sample (34), which made the estimated recession alpha imprecise. Instead we find that our volatility measure provides a better proxy for high versus low volatility periods.

5-10 funds get selected during the two-year period from 1992-1994, followed by a narrower FCS up to 2000. As shown in Figure 2, this period coincides with the best risk-adjusted performance of the FCS portfolio, suggesting that, at least for this historical period, our methodology succeeds in identifying a narrow set of funds with high subsequent performance. After 2000, the FCS undergoes considerable change with a broad array of funds entering and leaving the FCS up to 2007. During the last 18 months of the sample, the top FCS grows in size, and includes around 40 funds. On average, the tight FCS ($\lambda = 0.90$) includes 8 funds, while the medium ($\lambda = 0.50$) and wide ($\lambda = 0.10$) FCS include 16 and 22 funds, respectively.³⁷

Because the number of funds in our data increases over time, it is useful to also consider the percentage of funds identified by the FCS approach as being top performers. The bottom panel in Figure 4 provides a time-series plot of this proportion for the three different values of λ considered above. The percentage of top performers varies greatly over time. It fluctuates between 0.5% and 4% in the early part of the sample (assuming $\lambda = 0.90$), and is quite low up to 2000. After this period, 2% or more of the funds regularly get selected, and the proportion of included funds peaks at around 4% in 2011-2012. For smaller values of λ , the percentage of funds included is, of course, a bit higher, but it rarely exceeds 6%. On average, the tight FCS includes 1.3% of the funds, while the medium and wide FCS portfolios include 3.0 and 4.2% of the funds, respectively.

Figure 5 presents more detail on the individual funds identified as superior by the FCS approach. Specifically, for each fund that gets selected at least once during the sample, the figure shows when this particular fund is chosen. (The labeling on the y-axis is arbitrary, but maps one-to-one to the fund ID.) A single fund—Fidelity Select Technology—gets selected for most of the five-year period from 1995-2000. Between 2008 and 2011, a single fund—T Rowe Price Media and Telecommunications—is again selected. For most of the remaining periods a broader array of funds is selected, consistent with Figure 4.³⁸

To interpret these findings, note that the number of funds identified as being superior at a given point in time depends on the ability of the underlying performance model to accurately estimate individual funds' risk-adjusted returns and, thus, discriminate between funds. For example, during periods where individual fund alphas are estimated with large sampling errors, we would not expect *t*-tests such as equation (7) to be able to eliminate many funds from the FCS which, thus, might include many funds. In contrast, during times where individual funds' track records are more informative and alpha estimates are sharper, we would expect the FCS approach to be more discriminating, resulting in a narrower set of funds being included.

³⁷Recall that the λ parameter should be chosen to reflect the trade-off between an investor's utility from correctly including a genuinely superior fund versus her disutility from wrongly including an inferior fund.

³⁸The ability of the FCS to identify funds with superior performance can also be explored by studying the average value of the *t*-test used to eliminate inferior funds through time. This value hovers between -2 and -6 through most of our sample. Periods when the *t*-test is less negative (closer to zero) are times when a broader set of funds have more similar levels of expected performance; periods with more negative *t*-tests coincide with a narrow set of funds that are expected to outperform (and, thus, other funds are easy to eliminate through the sequence of *t*-tests).

Figure 6 illustrates this point by showing the predictive alpha of the fund(s) included in the FCS (marked in red) at three points in time (January 1997, 2002, and 2010) along with the alphas of those funds excluded (marked in blue). In the top and bottom plots a single fund gets selected by the FCS due to its markedly superior performance. In contrast, no single fund stands out in the middle panel, giving rise to a wider fund confidence set.

We conclude the following from our analysis. First, our ability to identify superior funds fluctuates significantly over time—sometimes this set is quite broad, containing up to 5% of the funds, at other times, the set is very narrow and contains less than a handful of funds—or even just a single fund. Second, the set of superior funds is almost always much smaller than the 10% figure assumed in decile ranking studies, and has been decreasing over time. This means that a procedure that forms portfolios based on a fixed fraction of the total number of funds—such as the conventional decile sorting approach—is likely to grossly underestimate the possibility of identifying funds with superior performance.

5.2. Identifying Inferior Funds

Figures 7 and 8 present plots similar to those presented in figures 4 and 5, but now applied to the set of funds identified to have inferior performance. Figure 7 shows that the set of funds with inferior performance fluctuates considerably through time, peaking at close to 5% for the narrow FCS portfolio ($\lambda = 0.90$), but often containing 1% or fewer of the funds. On average, the narrow FCS ($\lambda = 0.90$) includes 11 funds, while the medium ($\lambda = 0.50$) and wide FCS ($\lambda = 0.10$) portfolios on average include 25 and 38 funds, respectively. These figures correspond to 1.7%, 4.5% and 6.7% of the total number of funds, respectively. Thus, across all values of λ , the set of inferior funds is broader and includes more funds than the equivalent FCS for superior funds.

Figure 8 shows the inclusion plot for individual funds with inferior performance. We continue to see periods during which only a single inferior fund is selected by the tight FCS approach. For example, between 1993 and 1995, only the Centurion Growth Fund gets included in the tight FCS, while, between 2001 and 2004, the tight FCS of inferior funds predominantly consists of the Phoenix Oakhurst Strategy F fund.

Interestingly, passive index funds, such as the Vanguard S&P500, are never included in either the superior or inferior FCS portfolios, suggesting that, at each point in time, there always exist funds with genuinely better or genuinely worse ex-ante expected performance than a passive low-cost strategy.

An alternative way to address how difficult it is for mutual funds to generate abnormal performance and how this varies through time is by studying how many funds have significantly positive or negative alpha estimates. The top panel in Figure 9 plots smoothed kernel estimates of the proportion of funds with a statistically significant alpha estimate, identified using the methodology

of Romano and Wolf (2005), which is ideally suited to address this question. The number starts at 12% in 1990 (approximately 25 funds), drops to less than 5% in 1998 (approximately 15-20 funds), and ranges between 5% and 8% for the remainder of the sample, corresponding to between 40 and 75 funds. This shows that there is significant time variation in the scope for individual funds to outperform. The downward trend from 1990 through 1997 and, again, from 2002 until the end of the sample (2012) is consistent with competition intensifying through time (Pastor, et al. (2015)).

The bottom panel of Figure 9 plots the proportion of funds with significantly negative alpha estimates. On average, about 20-25% of funds have significantly negative alpha estimates, however, the scope for underperformance varies substantially through time. Fewer than 15% of funds had significantly negative alphas in 1994 and from 2010-2011, whereas more than 25% of funds had significantly negative alpha estimates during much of the decade from 1999-2009. At the end of the sample, three times as many funds had significantly negative alphas as funds with significantly positive alphas.

5.3. Turnover and Durations of Portfolios

As a way to illustrate how often individual funds enter and exit the FCS, we next compute the turnover in the set of funds included in the tight FCS portfolio, and compare this to the turnover among funds in the top and bottom decile portfolios. Specifically, we compute

$$Turnover = \frac{1}{2} \sum_{i=1}^{N_t} |\Delta w_{it}|, \quad (21)$$

where w_{it} is the portfolio weight on fund i at time t , and $\Delta w_{it} = w_{it} - w_{it-1}$ measures the change in the portfolio weight on fund i from month $t-1$ to month t .

The average monthly turnover for the top FCS portfolio is 0.19, while the corresponding figure for the bottom FCS portfolio is 0.26. For comparison, the monthly turnover of the top and bottom decile portfolios are 0.12 and 0.11, respectively. Thus, the turnover is somewhat higher for the FCS portfolios, which is what we would expect, given that these portfolios typically contain fewer funds than the more broadly diversified decile portfolios.

To help explain what generates the higher turnover in the FCS portfolios, Figure 10 provides details on the average duration of funds in the FCS (left column) and decile portfolios (right column) with top performers in the upper panels and bottom performers in the lower panels. Focusing on top-ranked funds, the figure shows that the main reason for the high turnover in the top FCS portfolios is that approximately 40% of all funds remain in the FCS portfolio for only one or two months, while another 20% of the funds get selected for three or four consecutive months. In comparison, about 45% of funds remain in the top decile portfolio for four months or less. There are also fewer funds with very long durations in the top FCS portfolio, compared to the top

decile portfolio. Similar conclusions hold for the portfolios of bottom-ranked funds shown in the lower panels of Figure 10. These short durations are in part a result of estimation errors affecting individual funds' alpha estimates but may also be indicative of strong competition which makes it more difficult for individual funds to be repeatedly included in the set of superior funds.

5.4. Business Cycle Variation in the Set of Superior and Inferior Funds

Our empirical analysis up to this point shows that the width of the set of funds identified to have superior performance varies a great deal over time. To see whether the proportion of funds identified to have either superior or inferior performance depends on the state of the economy, we next perform a simple analysis that regresses the percentage of funds in either the top or the bottom FCS on a range of state variables commonly used to capture the state of the economy. Specifically, we use the 1-month T-bill rate, the term spread, the default spread (i.e., the difference between the yield on portfolios of (Moody's) Baa- and Aaa-rated bonds) and the dividend yield as our state variables. These variables feature prominently in the literature on return predictability for the broad stock market (Welch and Goyal, 2008), and have also been used to model state-dependence in mutual fund performance (Ferson and Schadt, 1996; Avramov and Wermers, 2006; and Banegas, et al., 2013).

Table 6 reports the outcome of this analysis. We find that many of the state variables are significantly correlated with the width of the top FCS portfolio. Notably, the proportion of funds included in the top FCS tends to be higher in environments with a high short interest rate, a steep yield curve (high term spread), and a small default yield.³⁹

Turning to the funds in the bottom FCS, we find somewhat different results: the proportion of funds whose performance is judged to be inferior tends to be higher in economic states with high interest rates, a high default spread, and a high dividend yield. A higher default spread is often associated with economic recessions, and, so, this suggests a very different pattern than that uncovered for the FCS comprising funds with superior performance.

We conclude, from this evidence, that economic states with higher expected future growth (higher term spread) and lower risk premia (smaller default yield) are associated with a broader set of funds being identified as superior. Such states are more likely to occur in economic expansions, and, so, our finding is consistent with that of Kacperczyk et al. (2014) that some funds can pick stocks during economic upturns. Conversely, inferior performance tends to become more widespread during recessions.

³⁹A similar set of results apply if we focus on the number of funds (rather than the proportion of funds), although, now, the T-bill rate is insignificant.

6. Style and Industry Exposures

To gain insights into the investment strategies underlying the abnormal returns generated by the FCS portfolios, we next analyze the evolution in the FCS-chosen funds' style exposures. To this end, we exploit that we have access to quarterly holdings for all funds included in our analysis. The holdings data include information on funds' exposure to size, book-to-market, and momentum risk factors, along with their industry concentrations. For each FCS portfolio, we use these data to construct equal-weighted cross sectional averages of the characteristics and industry concentrations for the underlying funds. Because the time series of holdings data can be quite volatile, we consider rolling one-year averages of the portfolio characteristics.

6.1. Style Exposures

Table 7 presents attributes of the funds in the top and bottom FCS portfolios and, for comparison purposes, the average computed across all funds. The average attributes are calculated by first taking cross-sectional averages of the relevant set of funds at a given point in time, then averaging this cross-sectional mean over time. The attributes include the characteristic selectivity measure computed over the past one-month (CS-1m) and three-month (CS-3m) horizons, total net assets measured in millions of dollars (TNA), portfolio-holdings-based style attributes for size (Size), book-to-market (BTM), and momentum (MOM) style factors, the gap between actual returns and the hypothetical return based on reported portfolio holdings (return gap), fund flows, net cash inflows, expense ratio, and turnover.

We observe substantial differences in the attributes of superior and inferior funds, and between either of these types of funds and the average fund. Specifically, superior funds have substantially higher CS measures than average funds, which, in turn, have higher CS measures than inferior funds. This is to be expected, since the CS measure is used to estimate fund performance and, thus, determines which funds get allocated to the superior and inferior portfolios.

Interestingly, the portfolio of superior funds consists of funds whose average size (\$3.3bn) is considerably larger than that of both the average fund (\$1.7bn) and inferior funds (\$0.5bn), indicating that the average superior fund performs well enough to more than offset any diseconomies-of-scale. Turning to the style factors, the portfolio of superior funds holds firms with largely similar size, book-to-market ratios, and momentum, compared to the average fund, while the portfolio of inferior funds exhibits some differentiation on the size dimension. Both superior and inferior funds tend to exhibit a similar return gap, which indicates that the return gap is not a useful proxy for the skills captured by the FCS technique.

Further, inflows are much larger for superior funds, as expected (due to their higher past returns). Superior funds also have a much lower expense ratio and somewhat lower turnover than

inferior funds. This is consistent with the findings of Carhart (1997) that higher expense ratios and higher turnover tend to reduce net return performance.

The above-mentioned style exposure results, which were averaged over time, mask important differences during particular time periods. Accordingly, Figure 11 plots the time-series of style exposures for funds in the top (left column) and bottom FCS (right column), along with style exposures for the average fund. Comparing the plots for the top FCS portfolio to that of the average fund, superior funds are more volatile in all three style dimensions. This result indicates that the FCS rotates among funds residing in different styles, depending on macroeconomic conditions.

Funds in the superior FCS portfolio tend to have a heavier allocation to small- and mid-cap stocks than the average fund—the most notable exception occurring during the global financial crisis (2008-2009) where the superior funds overweighted large-cap stocks. The top FCS portfolio tilts towards growth stocks from 1994 to 2001 (consistent with growth stocks having the most uncertainty in valuations during this period), but overweights value stocks from 2001-2007. As shown in the bottom panel of Figure 11, the top FCS funds have a notably higher exposure to the momentum factor than the average fund only for brief spells during 1998, 2009 and 2012, suggesting that this portfolio is not predominantly capturing rewards for momentum risk.

A very different picture emerges for the inferior FCS funds. These funds overweight large cap stocks between 1993 and 1997, and again after 2009, overweighting small stocks during the interim period from 1997-2008. The inferior FCS funds overweight value stocks for most of the sample, while their exposure to the momentum risk factor only deviates substantially from that of the average fund during 1994-1997.

Figure 12 shows histograms depicting the performance of the top FCS funds (top row) for different realizations of the size, book-to-market, and momentum risk factors sorted into quintiles. For example, the left-most blocks (within each style/FCS segment, e.g., “Size Superior FCS”) show the average risk-adjusted performance of the FCS funds during months with the lowest quintile of realized returns of the different risk factors, while the right-most blocks show the average performance during months in which the realized value of the factors were in the top quintile.

The top left panel reveals a systematic and monotonically increasing relation between the performance of the superior FCS funds and the size factor: the better large-cap stocks perform relative to small stocks, the better funds in the superior FCS perform during a particular month. The relation between the book-to-market factor and the performance of the superior FCS funds shows that superior funds perform far better (with an alpha close to 100 bps/month) during months where this factor was at its lowest, i.e., during months where growth stocks outperform value stocks. Finally, as shown in the right panel, superior FCS funds tend to perform better when the momentum risk factor is very low.

The bottom row of panels in Figure 12 presents a similar set of plots for the inferior FCS funds. There is evidence that the inferior FCS funds perform worse in periods where large stocks

outperform (left panel), when value stocks outperform growth stocks by a large margin (middle panel), and during times with a low realized value of the momentum factor (bottom panel).

Overall, these results indicate that the alphas of superior and inferior funds have a common component that varies with the performance of the market capitalization and value/growth dimensions in the marketplace and, to a lesser extent, with the outcome of the momentum risk factor.

6.2. *Industry Concentration*

Table 8 reports time-series average industry concentrations for the top FCS funds, bottom FCS funds, and for all funds (equal-weighted averages) based on an industry classification that uses 48 industries (obtained from Ken French's data site). Industry concentrations are also reported separately for high-volatility and low-volatility states. A result that stands out is that funds in the top FCS have much greater average concentrations in communications, business services, computers, and electronic equipment, than funds in the bottom FCS, indicating the importance of manager skill in choosing among technology firms during the 1985 to 2012 period. Moreover, these differences are particularly high in the high-volatility state for the communications, business services, and electronic equipment industries.

One way to summarize the extent to which the superior and inferior FCS portfolios take industry bets is by calculating “active share” (Cremers and Petajisto, 2009) at the industry level, i.e., one-half times the absolute deviation of each industry weight, measured relative to the average mutual fund’s industry weighting, summed across industries, then averaged across time. For superior FCS funds, this industry active share measure is 27.7 overall versus an active share of 10.8 for inferior FCS funds. Industry active share is twice as high in recessions (45.1 and 21.7 for superior and inferior funds, respectively) compared to expansions (25.7 and 10.1), suggesting that both superior and inferior funds deviate substantially more from the average fund’s industry weights during recessions. Moreover, for the superior FCS funds, industry active share is notably higher in high-volatility states (32.5) than in low-volatility states (25.9), whereas it varies less across these states for the inferior FCS funds.

To get a sense of the dynamic industry tilting strategies implied by the portfolios of superior and inferior funds, Figure 13 plots one-year rolling averages of industry concentrations for the top FCS funds and the average fund. We show results for the eight industries with the largest (absolute) differences in industry concentration (relative to the average). The top funds greatly overweight computer and electronic equipment stocks after 2005, and also overweight business services and machinery during shorter spells. Conversely, the top funds underweight retail stocks and, in particular, banking stocks throughout most of the sample.

Very different industry concentrations emerge for the inferior FCS funds (Figure 14). These funds overweight apparel, medical equipment, computers, and wholesale industries during the early

1990s, while conversely banking, business services and pharmaceutical products are underweighted during this period. Towards the end of the sample, banking, computers, and communication shares get overweighted by the inferior FCS funds.

Overall, this analysis shows that FCS-identified superior funds, while tilting more toward technology-intensive companies, are focused on different sectors of the market during different years. This indicates that the turnover of funds in the top FCS portfolio, over time, that was documented earlier in this paper is at least partially due to the time-varying skills of managers who invest in a particular sector of the market. Moreover, it also suggests that it is difficult for a manager to maintain his performance over long periods of time.

7. Conclusion

This paper presents a new approach to selecting funds with superior performance. Our approach goes well beyond earlier studies, in that we predict the *set* of funds most likely to outperform in the future. This goal is very different from predicting whether an individual fund might outperform in the future. By conducting a large set of pair-wise performance comparisons across a large set of mutual funds, the approach iteratively eliminates funds with inferior performance. In line with the finding in recent studies that only a relatively narrow—and declining—set of funds is capable of outperforming on a risk-adjusted basis, our results suggest that it is important to choose a stringent procedure that is capable of eliminating funds with inferior performance.

Several insights emerge from our analysis. First, we find that the set of funds identified ex ante to have superior performance subsequently goes on to generate risk-adjusted returns that are substantially higher than the returns obtained by alpha-ranked funds in the top 5-10%, or by funds with significantly positive alphas. Clearly there is substantial heterogeneity in performance even among the funds with the highest alpha estimates.

Second, because of the considerable sampling error surrounding estimates of individual funds' alphas, our results show that it is important to combine funds' return records with information obtained from their holdings to obtain sufficiently accurate performance estimates making it possible to discriminate superior from inferior funds.

Third, the proportion of funds—as well as the identify of the individual funds—deemed to be superior varies considerably over time and generally is far smaller than the 5-10% of funds in existence often assumed in studies that use portfolio decile sorts to gauge top funds' performance. Consistent with predictions from the theoretical analysis in Kacperczyk et al. (2016), we find that the cross-sectional dispersion in fund performance increases in states with high market volatility and that the performance of funds identified as being superior also rises when aggregate (or idiosyncratic) volatility is high. This suggests that the payoffs to mutual funds from allocating more attention to the stock market is higher during such periods.

Fourth, superior funds achieved their high returns by substantially deviating from the average fund's industry concentration and loadings on systematic risk factors. For example, the superior funds overweighted value stocks from 2003 to 2009, switched between over- and underweighting small cap stocks, but generally took only small bets on momentum risk. These funds also undertook sizeable bets on individual sectors, overweighting computer and electronic equipment stocks but underweighted retail and banking stocks.

Internet Appendix A: Estimation and prediction of fund performance

This appendix explains how we estimate the unknown parameters of the latent state models and generate predictions of the conditional alpha.

For each fund, i , we observe a sample of excess returns, R_{it} . We cast the return model into state space form as follows:

$$\begin{aligned} R_{it} &= a_i F_{it-1} + b_i F_{it-1}^2 - k_i + (\bar{\beta}'_i + \mathbf{c}'_i F_{it-1}) \mathbf{z}_t + \varepsilon_{it} \\ &= G_{it}(F_{it-1}) + \varepsilon_{it}, \\ F_{it} &= \nu_i F_{it-1} + \eta_{it}. \end{aligned}$$

We focus on models where R_{it} is a linear function of the signal. Define $\hat{F}_{it|t-1}$ and $P_{it|t-1}$ as the conditional mean and variance of the i th fund's signal, given information at time $t-1$. The extended Kalman filter relies on a linear approximation of $G_{it}(F_{it-1})$ around $\hat{F}_{it-1|t-2}$,

$$G_{it}(F_{it-1}) \approx G_{it}(\hat{F}_{it-1|t-2}) + G_{F,it}(F_{it-1} - \hat{F}_{it-1|t-2}),$$

where

$$G_{F,it} = \left. \frac{\partial G_{it}(F_{it-1})}{\partial F_{it-1}} \right|_{F_{it-1} = \hat{F}_{it-1|t-2}}.$$

Given starting values for $\hat{F}_{i1|0}$ and $P_{i1|0}$, the following recursions constitute the extended Kalman filter:

$$\begin{aligned} v_{it} &= R_{it} - G_{it}(\hat{F}_{it-1|t-2}), \\ \hat{F}_{i,t-1|t-1} &= \hat{F}_{it-1|t-2} + P_{it-1|t-2} G'_{F,it-1} K_{it-1}^{-1} v_{it-1}, \\ \hat{F}_{it|t-1} &= \nu_i \hat{F}_{i,t-1|t-1}, \\ K_{it} &= G_{F,it} P_{it-1|t-2} G'_{F,it} + \varepsilon_{it}^2, \\ P_{i,t-1|t-1} &= P_{it-1|t-2} - P_{it-1|t-2} G'_{F,it-1} K_{it-1}^{-1} G_{F,it-1} P_{it-1|t-2}, \\ P_{it|t-1} &= \nu_i^2 P_{i,t-1|t-1} + \eta_{it-1}^2. \end{aligned}$$

Using information up to time $t-1$, we can estimate the parameters of the latent state models presented in Section 3. Let $\hat{\theta}_{i,t-1}$ denote the parameter estimates based on time $t-1$ information. We use the Kalman filter to forecast the signal one step ahead:

$$\hat{F}_{i,t|t-1} = \hat{F}_{i,t-1}(\hat{\theta}_{i,t-1}).$$

For each fund, $i = 1, \dots, N_t$, we also predict the alpha one step ahead

$$\hat{\alpha}_{i,t|t-1} = \hat{\alpha}_{i,t|t-1}(\hat{F}_{i,t|t-1}).$$

The forecast of alpha at time t , given information at time $t - 1$, is therefore a function of the forecast of the signal and the parameter estimates available at time $t - 1$.

Internet Appendix B: Local estimates of fund performance

This appendix shows how we generate non-parametric estimates of local (in time) return performance.

To track variation in the performance of portfolios over time, we use a flexible, nonparametric approach that allows for time-varying alpha performance, as well as time-varying factor exposures through the following smooth time-varying parameter model:

$$R_{pt} = \alpha_p(t/T) + \beta'_p(t/T)\mathbf{z}_t + \varepsilon_{pt}, \quad t = 1, \dots, T. \quad (22)$$

Here $\alpha_p(\cdot)$ and $\beta_p(\cdot)$ are unknown smooth functions that are allowed to depend on the sample ‘‘fraction’’, t/T , and, thus, can vary over time.⁴⁰ To see how we can nonparametrically estimate the parameters, define the vector of regressors $\mathbf{X}_t = (1 \ \mathbf{z}_t)$ and parameters $\boldsymbol{\theta}_p(t/T) = (\alpha_p(t/T) \ \beta'_p(t/T))'$, and rewrite equation (22) as

$$R_{pt} = \mathbf{X}'_t \boldsymbol{\theta}_p(t/T) + \varepsilon_{pt}. \quad (23)$$

The parameters $\boldsymbol{\theta}_p(t/T)$ can be estimated by means of a two-step procedure that first considers the OLS estimator $\hat{\gamma}_p = (\hat{\gamma}_{p0}, \hat{\gamma}_{p1})'$ of the transformed model

$$k_{st}^{1/2} R_{ps} = k_{st}^{1/2} \mathbf{X}'_s \gamma_{p0} + k_{st}^{1/2} \left(\frac{s-t}{T} \right) \mathbf{X}'_s \gamma_{p1} + \varepsilon_{ps}, \quad s = 1, \dots, T, \quad (24)$$

where $k_{st} = k(\frac{s-t}{Th})$ is a kernel function, and h is the bandwidth.⁴¹ In the second step, we construct an estimator of $\boldsymbol{\alpha}_{pt} = \boldsymbol{\alpha}_p(t/T)$ as

$$\hat{\boldsymbol{\alpha}}_{pt} = (\mathbf{e} \otimes \mathbf{I}) \hat{\gamma}_p, \quad (25)$$

where $\mathbf{e} = (1, \mathbf{0})$, \mathbf{I} is an identity matrix, and \otimes denotes the Kronecker product; see Cai (2007) and Chen and Hong (2012) for further discussion of this approach, and its ability to capture time variation in parameter estimates.

⁴⁰Technically, the $\alpha_p(\cdot)$ and $\beta_p(\cdot)$ functions allow for a finite number of discontinuities.

⁴¹Following Chen and Hong (2012), we set $h = T^{-1/5}/\sqrt{12}$.

The estimate $\hat{\alpha}_{p,t}$ from (25) portrays the evolution in the performance of the different portfolios in a way that does not “average out” potentially interesting time-variation in α_p . This can be contrasted with the conventional full-sample approach or even a rolling-window procedure which does not take into account how much the performance varies over time.

References

- Avramov, D., Wermers, R., 2006. Investing in mutual funds when returns are predictable. *Journal of Financial Economics* 81, 339–377.
- Banegas, A., Gillen, B., Timmermann, A., Wermers, R., 2013. The cross-section of conditional mutual fund performance in European stock markets. *Journal of Financial Economics* 108, 269–726.
- Barras, L., O. Scaillet, and R. Wermers, 2010, False Discoveries in Mutual Fund Performance: Measuring Luck in Estimated Alphas. *Journal of Finance* 65, 179-216.
- Berk, J.B. and R. C. Green, 2004, Mutual fund flows and performance in rational markets. *Journal of Political Economy* 112, 1269-1295.
- Blake, D., A.G. Rossi, A. Timmermann, I. Tonks, and R. Wermers, 2013, Decentralized Investment Management: Evidence from the Pension Fund Industry. *Journal of Finance* 68, 1133-1178.
- Bollen, N., Busse, J., 2004. Short-term persistence in mutual fund performance. *Review of Financial Studies* 18, 569–597.
- Cai, Z., 2007. Trending time-varying coefficient time series models with serially correlated errors. *Journal of Econometrics* 136, 163–188.
- Carhart, M., 1997. On persistence in mutual fund performance. *Journal of Finance* 52, 569–597.
- Chen, B., Hong, Y., 2012. Testing for smooth structural changes in time series models via nonparametric regression. *Econometrica* 80, 1157–1183.
- Christopherson, J., Ferson, W., Glassman, D., 1998. Conditioning manager alphas on economic information: Another look at the persistence of performance. *Review of Financial Studies* 11, 111–142.
- Cremers, M. and A. Petajisto, 2009, How Active is Your Fund Manager? A New Measure That Predicts Performance. *Review of Financial Studies* 22(8), 3329-3365.
- Daniel, K., Grinblatt, M., Titman, S., Wermers, R., 1997. Measuring mutual fund performance with characteristic-based benchmarks. *Journal of Finance* 52, 1035–1058.
- Doornik, J., 2006. Object-Oriented Matrix Programming Using Ox, 3rd Edition. London: Timberlake Consultants Press and Oxford: www.doornik.com.
- Ferson, W.E., and Schadt, R.W., 1996, Measuring fund strategy and performance in changing economic conditions. *Journal of Finance* 51, 425-461.

- Foerster, S., J.T. Linnainmaa, B.T. Melzer, and A. Previtero, 2015, Retail Financial Advice: Does One Size Fit All? Forthcoming in *Journal of Finance*.
- Garleanu, N., and L.H. Pedersen, 2016, Efficiently Inefficient Markets for Assets and Asset Management. Manuscript, UC Berkeley and Copenhagen Business School.
- Glode, V., Hollifield, B., Kacperczyk, M., Kogan, S., 2011. Time-varying predictability in mutual fund returns. Unpublished Manuscript, Wharton School.
- Hansen, P., Lunde, A., 2010. Mulcom 2.0, econometric toolkit for multiple comparisons.
- Hansen, P., Lunde, A., Nason, J., 2011. The model confidence set. *Econometrica* 2 (79), 453–497.
- Harvey, C.R., and Y. Liu, 2016, Rethinking Performance Evaluation. Unpublished manuscript, Duke University.
- Kacperczyk, M., C. Sialm, and L. Zheng, 2005, On the Industry Concentration of Actively Managed Equity Mutual Funds. *Journal of Finance* 60, 1983-2012.
- Kacperczyk, M., C. Sialm, and L. Zheng, 2008, Unobserved Actions of Mutual Funds. *Review of Financial Studies* 21, 2379-2416.
- Kacperczyk, M., van Nieuwerburgh, S., Veldkamp, L., 2014. Time-varying fund manager skill. *Journal of Finance* 69, 1455-1484.
- Kacperczyk, M., S van Nieuwerburgh, and L. Veldkamp, 2016, A Rational Theory of Mutual Funds' Attention Allocation. *Econometrica* 84(2), 571-626.
- Kinnel, R., 2010. How expense ratios and star ratings predict success. *Morningstar Advisor* 10, August.
- Kosowski, R., 2011. Do mutual funds perform when it matters most to investors. *Quarterly Journal of Finance* 1, 607–664.
- Liu, J., Timmermann, A., 2013. Optimal convergence trade strategies. *Review of Financial Studies* 26 (4), 1048–1086.
- Mamaysky, H., Spiegel, M., Zhang, H., 2007. Improved forecasting of mutual fund alphas and betas. *Review of Finance* 11, 359–400.
- Mamaysky, H., Spiegel, M., Zhang, H., 2008. Estimating the dynamics of mutual fund alphas and betas. *Review of Financial Studies* 21 (1), 233–264.
- Morey, M., 2005. The kiss of death: A 5-star Morningstar rating?, *Journal of Investment Management* 3 (2).

- Pastor, L., R.F. Stambaugh, and L.A. Taylor, 2015, Scale and Skill in Active Management. *Journal of Financial Economics* 116, 23-45.
- Patton, A., Timmermann, A., 2010. Monotonicity in asset returns: New tests with applications to the term structure, the CAPM and portfolio sorts. *Journal of Financial Economics* 98, 605–625.
- Romano, J., Wolf, M., 2005. Stepwise multiple testing as formalized data snooping. *Econometrica* 73 (4), 1237–1282.
- Wei, K., R. Wermers, and T. Yao, 2015, Uncommon Value: The Investment Performance of Contrarian Funds. *Management Science*, October, 2394–2414.
- Welch, I., and A. Goyal, 2008. A comprehensive look at the empirical performance of equity premium prediction. *The Review of Financial Studies* 21 (4), 1455–1508.
- Wermers, R., 2000, Mutual Fund Performance: An Empirical Decomposition into Stock-Picking Talent, Style, Transactions Costs, and Expenses. *Journal of Finance* 55, 1655-1695.
- White, H., 2000. A reality check for data snooping. *Econometrica* 68 (5), 1097–1126.

Table 1: Cross-section of alpha estimates

	Benchmark	Latent Skill	Latent Skill Holdings
Mean	-0.057	-0.057	-0.066
Min	-1.114	-1.168	-1.610
Max	0.942	0.813	1.713
5%	-0.361	-0.370	-0.376
10%	-0.272	-0.279	-0.273
25%	-0.166	-0.162	-0.157
50%	-0.058	-0.051	-0.062
75%	0.053	0.056	0.034
90%	0.169	0.167	0.136
95%	0.255	0.244	0.211

This table shows the cross-sectional distribution of four-factor alpha estimates (in percent per month) obtained from monthly returns data on U.S. equity mutual funds over the period 1980:06-2012:12. Mutual fund returns are net of transaction costs and fees and are calculated in excess of a one-month T-bill rate. The four risk factors are excess returns on a market portfolio, small-minus-big market cap and high-minus-low book-to-market value Fama-French factors and a momentum risk factor. The three columns report results for different performance models. The benchmark model assumes that alpha is constant and estimated by OLS using a 60-month rolling window. The latent skill (LS) model allows for a time-varying alpha and uses returns data to compute the expected value of each fund's alpha. The latent skill-holdings (LSH) model also allows for a time-varying alpha but combines returns and holdings data to estimate each fund's alpha. The extended Kalman filter is used to extract the latent signal underlying the models with time-varying alphas. The reported estimates of the cross-sectional distribution is based on the individual funds' average alpha estimates.

Table 2: Risk-adjusted performance for alpha-ranked decile portfolios

Portfolio	Benchmark		Latent Skill		Latent Skill Holdings	
P1A	-0.22	(-2.66)	-0.25	(-3.65)	-0.18	(-2.51)
P1B	-0.16	(-2.53)	-0.17	(-2.66)	-0.14	(-2.24)
P1	-0.19	(-2.74)	-0.21	(-3.32)	-0.16	(-2.51)
P2	-0.14	(-2.46)	-0.15	(-2.57)	-0.14	(-2.51)
P3	-0.08	(-1.60)	-0.10	(-2.06)	-0.10	(-2.15)
P4	-0.09	(-1.76)	-0.08	(-1.75)	-0.09	(-1.86)
P5	-0.07	(-1.44)	-0.08	(-1.74)	-0.09	(-2.25)
P6	-0.07	(-1.45)	-0.08	(-1.76)	-0.06	(-1.43)
P7	-0.03	(-0.65)	-0.05	(-1.15)	-0.05	(-1.23)
P8	-0.05	(-1.08)	-0.03	(-0.61)	-0.02	(-0.45)
P9	-0.02	(-0.42)	0.04	(0.75)	0.00	(0.04)
P10	0.06	(0.83)	0.05	(0.70)	0.02	(0.31)
P10A	0.00	(0.07)	0.05	(0.75)	-0.00	(-0.04)
P10B	0.11	(1.26)	0.05	(0.57)	0.05	(0.59)
P10-P1	0.25	(3.46)	0.26	(4.94)	0.18	(4.00)
MR p-val	0.07		0.00		0.00	

This table reports four-factor alphas (in percent per month) for a set of alpha-ranked decile portfolios, with t-statistics in brackets. Each month we rank the set of mutual funds in existence according to their expected alphas and allocate them to equal-weighted decile portfolios listed in increasing order from funds with the lowest alphas ($P1$) to funds with the highest alphas ($P10$). The alpha ranking is repeated each month during the sample 1985:07-2012:12 and so produces a time series of portfolio returns from which the reported alpha estimates are computed. The columns report results for three different performance models. The benchmark model assumes that alpha is constant and estimated by OLS using a 60-month rolling window. The latent skill model allows for a time-varying alpha and uses returns data to compute the expected value of next month's alpha. The latent skill-holdings model also allows for a time-varying alpha but combines returns and holdings data to estimate each fund's alpha. $P1A$ and $P1B$ are constructed by splitting the funds in the bottom decile portfolio into two new portfolios of equal size. $P10A$ and $P10B$ are constructed in the same way. $P10 - P1$ is a portfolio constructed from the difference between the top and bottom decile portfolios. The table reports alpha estimates in percentage terms followed by t-statistics in brackets. If a particular performance model provides an accurate ranking of the funds we would expect the alpha estimates to increase monotonically from $P1$ to $P10$. We test for this monotonically increasing pattern against a flat or a decreasing relation using the monotonic relation (MR) test proposed by Patton and Timmermann (2010). Low p-values are evidence against a flat or decreasing pattern and so are suggestive of a precise ranking of alphas.

Table 3: Risk-adjusted performance for FCS portfolios of superior funds

Panel A: Top 15 percent of funds						
	Benchmark		Latent Skill		Latent Skill Holdings	
	Equal	Value	Equal	Value	Equal	Value
FCS-tight	0.30 (1.24)	0.36 (1.41)	0.25 (1.94)	0.25 (1.81)	0.54 (2.13)	0.52 (2.08)
FCS-medium	0.27 (1.59)	0.23 (1.33)	0.15 (1.53)	0.14 (1.24)	0.18 (1.35)	0.18 (1.24)
FCS-wide	0.20 (1.51)	0.19 (1.38)	0.10 (1.05)	0.02 (0.16)	0.06 (0.51)	-0.02 (-0.14)
P10	0.06 (0.83)	0.07 (0.89)	0.05 (0.70)	0.04 (0.61)	0.02 (0.31)	0.07 (1.06)
Panel B: Top 5 percent of funds						
	Benchmark		Latent Skill		Latent Skill Holdings	
	Equal	Value	Equal	Value	Equal	Value
FCS-tight	0.36 (1.34)	0.39 (1.41)	0.10 (0.57)	0.10 (0.50)	0.48 (1.77)	0.47 (1.74)
FCS-medium	0.45 (2.28)	0.41 (2.02)	0.14 (1.06)	0.18 (0.99)	0.36 (1.92)	0.43 (2.04)
FCS-wide	0.36 (2.03)	0.28 (1.56)	0.08 (0.63)	0.06 (0.40)	0.24 (1.61)	0.21 (1.40)
P10	0.06 (0.83)	0.07 (0.89)	0.05 (0.70)	0.04 (0.61)	0.02 (0.31)	0.07 (1.06)
Panel C: Top 25 percent of funds						
	Benchmark		Latent Skill		Latent Skill Holdings	
	Equal	Value	Equal	Value	Equal	Value
FCS-tight	0.30 (1.33)	0.35 (1.50)	0.17 (1.60)	0.16 (1.44)	0.38 (1.94)	0.38 (1.91)
FCS-medium	0.19 (1.24)	0.20 (1.30)	0.12 (1.48)	0.11 (1.20)	0.16 (1.18)	0.17 (1.20)
FCS-wide	0.14 (1.20)	0.14 (1.16)	0.08 (1.07)	0.04 (0.51)	0.05 (0.48)	-0.00 (-0.04)
P10	0.06 (0.83)	0.07 (0.89)	0.05 (0.70)	0.04 (0.61)	0.02 (0.31)	0.07 (1.06)
Panel D: Romano and Wolf						
	Benchmark		Latent Skill		Latent Skill Holdings	
	Equal	Value	Equal	Value	Equal	Value
FCS-tight	0.20 (0.95)	0.28 (1.30)	0.08 (1.05)	0.11 (1.19)	0.48 (2.02)	0.53 (2.19)
FCS-medium	0.14 (1.04)	0.18 (1.30)	0.10 (1.51)	0.10 (1.36)	0.14 (1.01)	0.13 (0.85)
FCS-wide	0.08 (0.61)	0.07 (0.59)	0.06 (0.94)	0.09 (1.29)	0.13 (1.02)	0.11 (0.85)
P10	0.06 (0.83)	0.07 (0.89)	0.05 (0.70)	0.04 (0.61)	0.02 (0.31)	0.07 (1.06)

This table reports four-factor alphas for a set of portfolios formed by equal- or value-weighting stocks identified as being superior by the fund confidence set (FCS) approach. Alpha estimates are in percent per month with t-statistics shown in brackets. Each month we apply the FCS approach to all funds in existence so as to identify the set of superior funds. We then form a portfolio of these funds and record its return during the following month. This procedure is repeated each month during the sample 1985:07-2012:12 and so produces a time series of portfolio returns from which the reported four-factor alpha estimates are computed. In each panel the top three rows report results for three different values of the tightness parameter (λ) used to determine how strict to be when eliminating funds from the FCS. The top row uses a tight choice ($\lambda = 0.90$), resulting in a narrower set of funds being included, while the second and third rows give rise to medium ($\lambda = 0.50$) and wide ($\lambda = 0.10$) fund confidence sets. For comparison we also show results for the alpha-ranked top decile portfolio ($P10$). The FCS approach requires us to choose an initial set of candidate funds from which the superior funds are identified. Panel A uses the funds in the top 15 percent of the distribution of alpha estimates, while Panels B and C use the top 5 and top 25 percent of funds ranked by their expected alphas. Panel D uses the set of funds with significantly positive alpha, identified using the approach of Romano and Wolf (2005). The columns report results for three different performance models. The benchmark model assumes that alpha is constant and estimated by OLS using a 60-month rolling window. The latent skill (LS) model allows for a time-varying alpha and uses returns data to compute the expected value of next month's alpha. The latent skill-holdings (LSH) model also allows for a time-varying alpha but combines returns and holdings data to estimate each fund's alpha.

Table 4: Risk-adjusted performance for FCS portfolios of inferior funds

Panel A: Bottom 15 percent of funds						
	Benchmark		Latent Skill		Latent Skill Holdings	
	Equal	Value	Equal	Value	Equal	Value
FCS-tight	-0.43 (-3.04)	-0.41 (-2.78)	-0.36 (-3.23)	-0.27 (-2.49)	-0.32 (-2.48)	-0.26 (-1.74)
FCS-medium	-0.31 (-2.58)	-0.19 (-1.55)	-0.27 (-3.45)	0.01 (0.11)	-0.22 (-2.99)	-0.19 (-1.75)
FCS-wide	-0.21 (-3.09)	-0.14 (-1.78)	-0.25 (-3.32)	-0.01 (-0.12)	-0.19 (-2.99)	-0.15 (-1.54)
P1	-0.19 (-2.73)	-0.18 (-2.09)	-0.21 (-3.32)	-0.12 (-1.66)	-0.16 (-2.51)	-0.16 (-1.87)
Panel B: Bottom 5 percent of funds						
	Benchmark		Latent Skill		Latent Skill Holdings	
	Equal	Value	Equal	Value	Equal	Value
FCS-tight	-0.57 (-3.85)	-0.57 (-3.92)	-0.56 (-3.41)	-0.44 (-2.50)	-0.45 (-2.99)	-0.49 (-2.75)
FCS-medium	-0.49 (-3.82)	-0.45 (-3.47)	-0.41 (-4.14)	-0.14 (-1.23)	-0.37 (-3.97)	-0.30 (-2.28)
FCS-wide	-0.40 (-4.04)	-0.39 (-4.37)	-0.36 (-3.78)	-0.00 (-0.04)	-0.29 (-3.92)	-0.29 (-2.56)
P1	-0.19 (-2.73)	-0.18 (-2.09)	-0.21 (-3.32)	-0.12 (-1.66)	-0.16 (-2.51)	-0.16 (-1.87)
Panel C: Bottom 25 percent of funds						
	Benchmark		Latent Skill		Latent Skill Holdings	
	Equal	Value	Equal	Value	Equal	Value
FCS-tight	-0.40 (-2.82)	-0.35 (-2.31)	-0.21 (-2.12)	-0.17 (-1.58)	-0.26 (-2.68)	-0.16 (-1.44)
FCS-medium	-0.35 (-3.39)	-0.23 (-2.14)	-0.24 (-3.37)	-0.06 (-0.75)	-0.21 (-3.07)	-0.19 (-1.99)
FCS-wide	-0.18 (-2.46)	-0.10 (-1.33)	-0.23 (-3.44)	-0.04 (-0.59)	-0.19 (-3.08)	-0.13 (-1.60)
P1	-0.19 (-2.73)	-0.18 (-2.09)	-0.21 (-3.32)	-0.12 (-1.66)	-0.16 (-2.51)	-0.16 (-1.87)
Panel D: Romano and Wolf						
	Benchmark		Latent Skill		Latent Skill Holdings	
	Equal	Value	Equal	Value	Equal	Value
FCS-tight	-0.45 (-3.36)	-0.40 (-2.81)	-0.20 (-2.77)	-0.11 (-1.26)	-0.23 (-2.31)	-0.17 (-1.55)
FCS-medium	-0.25 (-3.35)	-0.11 (-1.35)	-0.20 (-3.01)	-0.11 (-1.54)	-0.15 (-2.20)	-0.13 (-1.52)
FCS-wide	-0.14 (-2.28)	-0.02 (-0.34)	-0.17 (-3.22)	-0.05 (-0.97)	-0.20 (-3.59)	-0.16 (-2.04)
P1	-0.19 (-2.73)	-0.18 (-2.09)	-0.21 (-3.32)	-0.12 (-1.66)	-0.16 (-2.51)	-0.16 (-1.87)

This table reports four-factor alphas for a set of portfolios formed by equal- or value-weighting stocks identified as being inferior by the fund confidence set (FCS) approach. Alpha estimates are reported in percent per month, with t-statistics reported in brackets. Each month we apply the FCS approach to all funds in existence so as to identify the set of inferior funds. We then form a portfolio of these funds and record its return during the following month. This procedure is repeated each month during the sample 1985:07-2012:12 and so produces a time series of portfolio returns from which the reported four-factor alpha estimates are computed. In each panel the top three rows report results for three different values of the tightness parameter (λ) used to determine how strict to be when eliminating funds from the FCS. The top row uses a tight choice ($\lambda = 0.90$), resulting in a narrower set of funds being included, while the second and third rows give rise to medium ($\lambda = 0.50$) and wide ($\lambda = 0.10$) fund confidence sets. For comparison we also show results for the alpha-ranked bottom decile portfolio ($P1$). The FCS approach requires us to choose an initial set of candidate funds from which the inferior funds are identified. Panel A uses the funds in the top 15 percent of the distribution of alpha estimates, while Panels B and C use the top 5 and top 25 percent of funds ranked by their expected alphas. Panel D uses the set of funds with significantly positive alpha, identified using the approach of Romano and Wolf (2005). The columns report results for three different performance models. The benchmark model assumes that alpha is constant and estimated by OLS using a 60-month rolling window. The latent skill (LS) model allows for a time-varying alpha and uses returns data to compute the expected value of next month's alpha. The latent skill-holdings (LSH) model also allows for a time-varying alpha but combines returns and holdings data to estimate each fund's alpha.

Table 5: Dispersion and risk-adjusted performance sorted by volatility

Panel A: Dispersion in Industry Concentration and CS Measure

	Industry concentration		CS measure	
	Aggregate risk		Aggregate risk	
	All funds	All funds	All funds	All funds
Low risk	0.464 [0.000]		2.291 [0.000]	
Med. risk	0.680 [0.000]		3.399 [0.000]	
High risk	1.023 [0.000]		6.093 [0.000]	
High-Low	0.559 [0.023]		3.802 [0.000]	

Panel B1: Dispersion of performance for superior funds (12 months)

	Aggregate risk		Idiosyncratic risk	
	FCS	P10	FCS	P10
Low risk	0.010 [0.893]	0.005 [0.834]	-0.048 [0.526]	-0.064 [0.013]
Med. risk	0.154 [0.223]	-0.094 [0.004]	0.064 [0.460]	-0.060 [0.041]
High risk	1.034 [0.000]	-0.016 [0.700]	1.183 [0.000]	0.019 [0.671]
High-Low	1.024 [0.000]	-0.021 [0.661]	1.232 [0.000]	0.083 [0.104]

Panel B2: Dispersion of performance for superior funds (36 months)

	Aggregate risk		Idiosyncratic risk	
	FCS	P10	FCS	P10
Low risk	0.045 [0.218]	-0.017 [0.169]	0.072 [0.036]	0.010 [0.381]
Med. risk	0.473 [0.000]	-0.017 [0.394]	0.095 [0.078]	-0.107 [0.000]
High risk	1.115 [0.000]	0.061 [0.006]	1.466 [0.000]	0.124 [0.000]
High-Low	1.070 [0.000]	0.078 [0.002]	1.394 [0.000]	0.114 [0.000]

Panel A of this table reports dispersion in industry concentrations and CS measure across all funds sorted into groups according to volatility. We define the aggregate risk in period t as $|\beta_t \sigma_{mt}|$, where β_t is the average market beta across funds and σ_{mt} is the realized volatility of the market based on daily returns. The idiosyncratic risk is calculated in a similar fashion. The risk measures are based on a rolling window with a length of 12 monthly observations. Dispersion in industry concentration is calculated as $\sum_{c=1}^{48} (w_{i,t}^c - w_{avg,t}^c)^2$, where $w_{i,t}^c$ is fund i 's weight in industry c at time t , and $w_{avg,t}^c$ is the weight of the average fund in industry c at time t . The dispersion in the CS measure is calculated as the average squared deviation from a fund's CS measure to the average CS measure across funds. Panel B1 reports average alpha performance of portfolios of superior funds identified by the fund confidence set (FCS) and the decile approach (P10) sorted according to the rolling 12-month volatility estimate. Average alpha estimates are reported in percent per month. Each month we apply the FCS approach to all funds in existence to identify the set of superior or inferior funds. We then form a portfolio of these funds and record its return during the following month. This procedure is repeated each month during the sample 1985:07-2012:12 and so produces a time series of portfolio returns from which four-factor alpha estimates are computed based on the same 12-month rolling window used to calculate the volatility measures. Panel B2 repeats this analysis for a 36-month rolling window. The FCS portfolios are equally weighted, selected from the top 15 percent of alpha-ranked funds and for $\lambda = 0.90$.

Table 6: Width of fund confidence set and macroeconomic state variables

State variable	Superior Funds	Inferior Funds
Constant	-0.00 (-0.28)	0.12 (2.23)
tbl	0.45 (2.65)	0.76 (2.70)
tms	1.23 (4.05)	0.30 (0.87)
dfy	-0.88 (-2.38)	1.33 (1.65)
dy	0.00 (0.67)	0.04 (2.53)

This table presents least squares estimates (with t-statistics in parentheses) from regressing the percentage of funds identified by the tight FCS approach as being superior (left panel) or inferior (right panel). The FCS uses as the initial set of candidate funds those funds identified to have significantly positive alpha estimates by the Romano-Wolf (2005) approach and assumes that $\lambda = 0.90$. We use the latent skill-holdings model with a time-varying alpha to compute each fund's alpha estimate and thus use both returns and holdings data. The macroeconomic state variables are the one-month T-bill rate (tbl), the term spread (tms), the default yield spread (dfy), and the dividend yield (dy).

Table 7: Attributes of funds included in the Fund Confidence Set

	Overall			High vol.			Low vol.		
	Superior	Inferior	Average	Superior	Inferior	Average	Superior	Inferior	Average
CS-1m($\times 100$)	0.36	-0.03	0.03	0.54	0.02	0.02	0.27	0.09	0.08
CS-3m($\times 100$)	0.42	-0.15	0.04	0.70	-0.14	0.04	0.27	-0.04	0.05
TNA	3381.16	508.36	1765.90	2540.63	762.60	1997.49	3680.65	277.20	1514.06
Size	4.20	3.88	4.14	4.38	4.01	4.20	4.06	3.66	4.07
BTM	2.84	2.97	2.89	2.70	2.89	2.86	2.92	3.03	2.91
MOM	3.17	3.19	3.14	3.02	3.20	3.10	3.25	3.12	3.13
RetGap12($\times 100$)	-0.06	-0.05	-0.04	-0.10	0.05	0.04	-0.02	-0.19	-0.08
Flow($\times 100$)	2.72	-0.85	1.40	1.92	-2.65	-0.45	3.11	0.32	2.57
ExpRatio($\times 100$)	1.13	1.65	1.22	1.09	1.49	1.20	1.21	1.78	1.25
Turnover	1.06	1.39	0.83	1.51	1.38	0.86	0.74	1.21	0.80

This table reports several attributes of the funds selected by the FCS approach. The table reports attributes for the FCS portfolio of superior funds, the FCS portfolio of inferior funds and the average across funds. The FCS portfolios are based on a tight parameterization ($\lambda = 0.9$) and the top 15% of alpha-ranked funds, using the latent skill holdings model to estimate time-varying alphas. Attributes are calculated by first taking cross-sectional averages of attributes of the funds in the portfolio at a given point in time, then taking time series averages. The attributes include the portfolio-weighted characteristic selectivity measure, one month (CS-1m) and three months (CS-3m) in percent, total net assets of funds in millions of dollars (TNA), portfolio holdings-based style attributes in the size (Size), book-to-market (BTM), and momentum (MOM) dimensions, lagged net return, compounded over the 12 months prior to each portfolio formation date (RetGap12), monthly percentage net cash inflows (Flow), computed as TNA minus one-quarter-lagged TNA, divided by three, fund expense ratio (ExpRatio), and percent monthly turnover (Turnover). These attributes are presented for all periods, for periods with high aggregate volatility only and for periods with low aggregate volatility only. Aggregate volatility is calculated as in Table 5.

Table 8: Industry concentrations of funds included in the fund confidence set

	Overall			High vol.			Low vol.		
	Superior	Inferior	Average	Superior	Inferior	Average	Superior	Inferior	Average
Agriculture	0.09	0.20	0.21	0.05	0.30	0.26	0.10	0.08	0.18
Food Products	0.84	1.36	1.56	0.37	1.25	1.38	0.87	1.22	1.63
Candy & Soda	0.16	0.27	0.11	0.25	0.27	0.13	0.12	0.20	0.09
Beer & Liquor	0.82	1.12	0.95	0.43	0.87	0.96	1.37	1.45	0.96
Tobacco Products	0.25	0.22	0.34	0.24	0.21	0.38	0.23	0.21	0.31
Recreation	0.24	0.33	0.34	0.19	0.25	0.25	0.39	0.40	0.41
Entertainment	1.56	0.88	0.99	1.82	0.72	0.72	1.78	1.12	1.21
Printing & Publishing	1.67	0.74	1.00	1.22	0.65	0.79	1.61	0.87	1.14
Consumer Goods	1.17	2.76	1.98	1.01	2.46	1.87	1.56	2.91	2.07
Apparel	0.61	1.88	0.85	0.28	0.98	0.81	0.96	3.38	0.91
Healthcare	0.90	1.26	1.31	0.66	1.11	1.10	1.02	1.51	1.45
Medical Equipment	1.19	3.12	2.13	0.70	1.96	2.14	1.11	3.58	2.10
Pharmaceutical Products	3.96	6.46	7.01	2.30	7.07	7.63	5.00	5.58	6.24
Chemicals	1.18	1.82	2.36	0.59	1.33	1.91	1.61	2.29	2.73
Rubber and Plastic Products	0.21	0.32	0.33	0.07	0.28	0.27	0.14	0.32	0.37
Textiles	0.09	0.23	0.25	0.04	0.13	0.14	0.15	0.31	0.36
Construction Materials	0.63	1.20	0.97	0.64	0.87	0.83	0.53	1.52	1.10
Construction	0.22	0.99	0.74	0.15	0.90	0.69	0.28	1.24	0.83
Steel Works Etc	0.58	1.67	1.15	0.29	1.02	0.90	0.73	2.05	1.30
Fabricated Products	0.03	0.10	0.10	0.03	0.11	0.07	0.03	0.12	0.13
Machinery	3.53	3.44	2.89	2.32	2.91	2.92	1.74	3.97	2.94
Electrical Equipment	0.87	1.28	1.50	0.34	1.58	1.42	1.59	1.26	1.52
Automobiles and Trucks	0.97	1.39	1.32	0.99	1.11	1.04	1.09	1.73	1.58
Aircraft	0.44	1.23	1.13	0.28	1.09	1.17	0.57	1.52	1.11
Shipbuilding and Railroad Equipment	0.03	0.10	0.17	0.02	0.12	0.21	0.05	0.08	0.16
Defense	0.12	0.26	0.28	0.06	0.34	0.28	0.20	0.28	0.29
Precious Metals	0.25	0.06	1.15	0.08	0.08	0.82	0.36	0.05	1.41
Non-Metallic and Industrial Metal Mining	0.26	0.30	0.65	0.19	0.28	0.64	0.34	0.37	0.71
Coal	0.17	0.13	0.15	0.06	0.13	0.20	0.39	0.21	0.16
Petroleum and Natural Gas	3.56	4.71	5.81	3.80	4.67	6.04	4.06	4.76	5.64
Utilities	1.04	2.68	4.12	1.06	2.25	3.98	1.10	2.88	4.13
Communication	10.93	4.49	4.79	16.58	4.29	5.08	9.79	4.29	4.40
Personal Services	0.62	0.80	0.65	0.32	1.06	0.66	1.05	0.54	0.64
Business Services	13.66	7.88	8.34	18.97	10.07	9.79	11.39	6.36	7.42
Computers	13.12	7.45	4.82	10.66	6.83	5.11	14.58	7.45	4.43
Electronic Equipment	11.11	6.61	5.35	12.69	6.99	5.91	8.09	6.35	5.01
Measuring and Control Equipment	1.45	1.30	1.30	0.97	1.38	1.41	0.99	1.19	1.24
Business Supplies	0.99	1.35	1.61	0.66	1.12	1.33	1.15	1.46	1.73
Shipping Containers	0.13	0.30	0.27	0.13	0.34	0.28	0.13	0.33	0.26
Transportation	1.79	1.65	2.20	2.03	1.76	1.97	1.81	1.73	2.52
Wholesale	1.19	2.94	1.91	0.74	2.65	2.00	1.72	3.61	1.85
Retail	5.67	6.12	6.47	6.78	7.09	6.52	6.05	5.61	6.46
Restaurants and Hotels and Motels	1.17	1.70	1.22	1.54	1.06	1.27	1.08	2.57	1.25
Banking	4.23	7.33	8.47	2.86	9.47	8.29	4.24	4.68	8.60
Insurance	4.14	4.04	5.43	2.65	4.48	5.03	4.72	3.72	5.73
Real Estate	0.16	0.09	0.25	0.29	0.14	0.28	0.07	0.10	0.23
Trading	1.17	2.60	2.37	0.62	3.42	2.46	1.64	1.53	2.41
Almost Nothing	0.82	0.82	0.72	0.97	0.55	0.67	0.40	1.02	0.64

This table reports industry concentrations in percent for the FCS portfolio of superior funds, the FCS portfolio of inferior funds, and the average across funds. The FCS portfolios are based on a tight parameterization ($\lambda = 0.9$) and the top 15% of alpha-ranked funds, using the latent skill holdings model to estimate time-varying alphas. Average industry concentrations are calculated by first taking cross-sectional averages of the industry concentrations of the funds in the portfolio at a given point in time, then taking time series averages. These industry concentrations are presented for all periods, for periods with high aggregate volatility and for periods with low aggregate volatility only. Aggregate volatility is calculated as in Table 5.

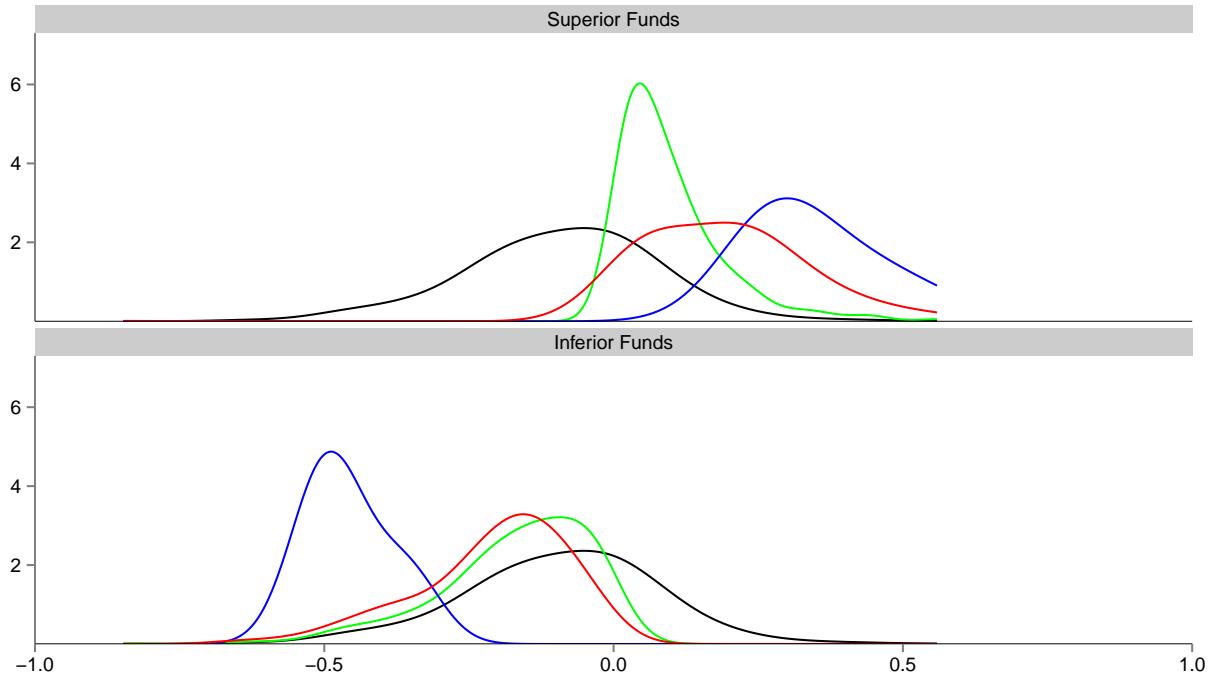


Figure 1: Time variation in alpha estimates for top-ranked funds. The figure presents distributions of predictive alphas for funds included in various portfolios at a single month in our sample (July 2006). The top panel presents distributions of predictive alphas for the top-ranked funds, while the bottom panel shows results for the bottom-ranked funds. In both cases the black line represents a portfolio consisting of all funds in existence in July 2006. The green line shows the distribution for funds with positive (top panel) or negative (bottom panel) predictive alphas. The red line shows the distribution for funds with significantly positive (top panel) or negative (bottom panel) predictive alphas, identified using the Romano Wolf (2005) approach. The blue line shows the distribution of funds included in the superior (top) or inferior (bottom) FCS, in both cases assuming a narrow set ($\lambda = 0.90$) and using the time-varying latent skill and holdings performance model.

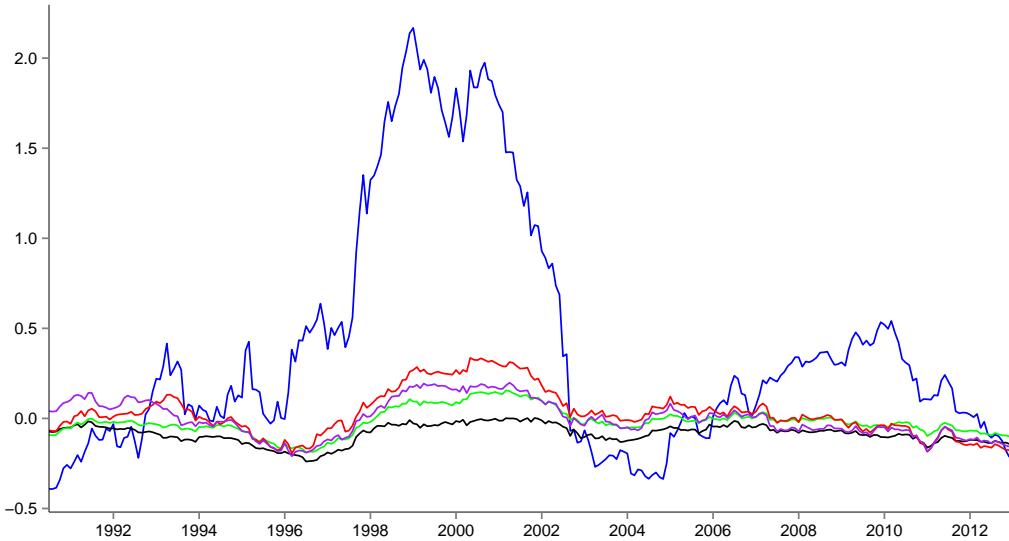


Figure 2: Time variation in alpha estimates for superior funds. This figure shows nonparametric estimates of local time-variation in four-factor alpha estimates (in percent per month) for a variety of portfolios consisting of top-ranked funds. The blue line tracks the performance of the tight FCS portfolio ($\lambda = 0.90$) of superior funds based on the latent skill holdings model. The green line tracks the performance of a portfolio consisting of all funds with positive predictive alphas while the red line tracks the performance of a portfolio consisting of the top 15% of funds with positive alpha estimates. The purple line tracks the performance of a portfolio consisting of the top decile of alpha-ranked funds ($P10$). For comparison the black line tracks the variation in the alpha of the average fund.

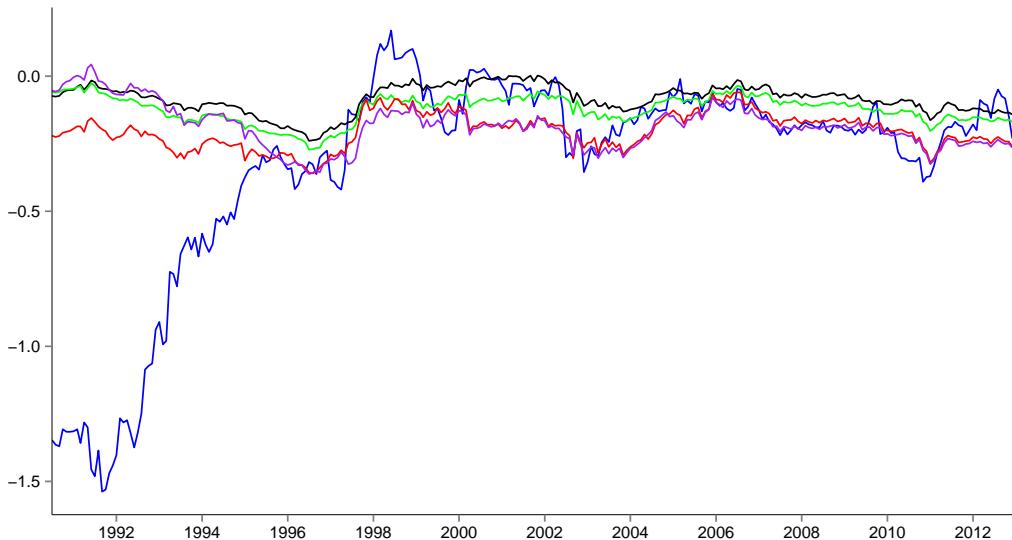


Figure 3: Time variation in alpha estimates for inferior funds. This figure shows nonparametric estimates of local time-variation in four-factor alpha estimates (in percent per month) for a variety of portfolios consisting of bottom-ranked funds. The blue line tracks the performance of the tight FCS portfolio ($\lambda = 0.90$) of inferior funds based on the latent skill holdings model. The green line tracks the performance of a portfolio consisting of all funds with negative predictive alphas while the red line tracks the performance of a portfolio consisting of the bottom 15% of funds with negative alpha estimates. The purple line tracks the performance of a portfolio consisting of the bottom decile of alpha-ranked funds ($P1$). For comparison the black line tracks the variation in the alpha of the average fund.

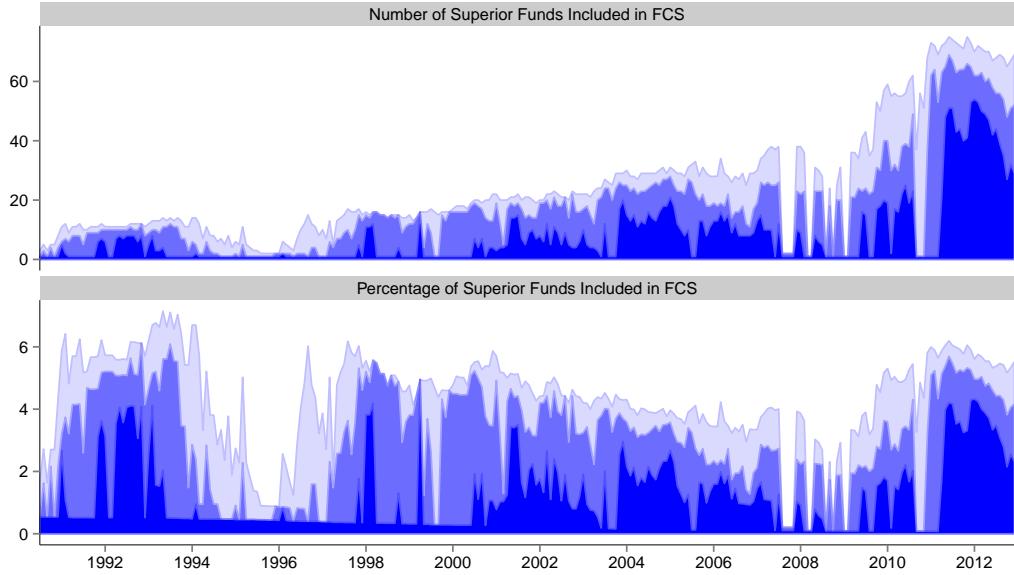


Figure 4: Evolution in the set of funds identified as being superior by the FCS approach. The top panel shows the evolution in the number of funds included in the tight, medium and wide fund confidence sets of superior funds. The bottom panel shows the corresponding evolution in the percentage of funds identified as being superior by the FCS approach. Funds are selected from the top 15 percent of alpha-ranked funds using the latent skill-holdings model with a time-varying alpha to compute each fund's alpha estimate. Dark blue areas correspond to $\lambda = 0.90$ (tight set), while lighter areas correspond to $\lambda = 0.50$ (medium set) and $\lambda = 0.10$ (wide set).

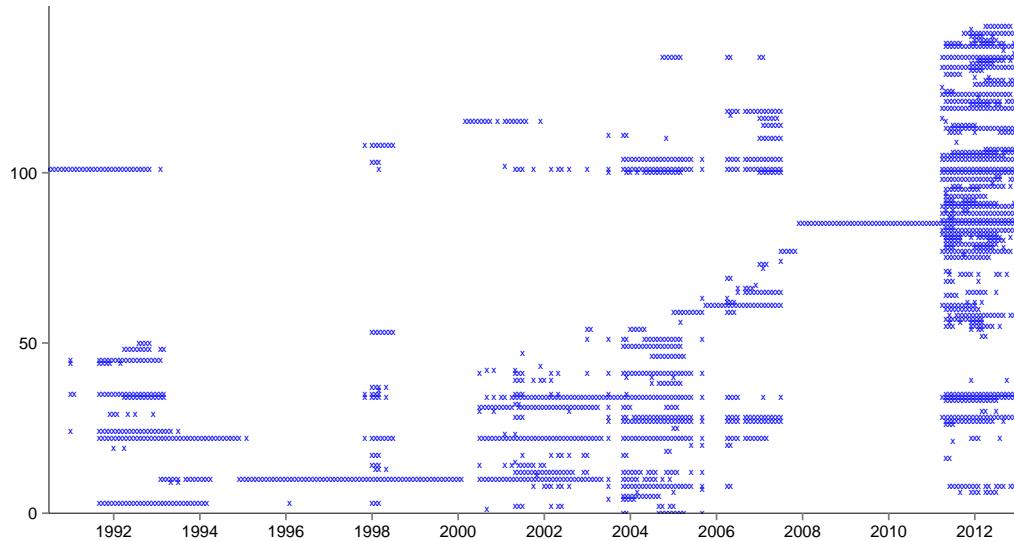


Figure 5: Identification of individual funds with superior skill by the FCS approach. The plot illustrates the evolution in the composition of individual funds with superior skills as identified by the FCS approach. Each fund that is included in the FCS at least once during the sample is assigned a unique number on the y-axis and a cross shows when this fund is included. The FCS is based on the latent skill holdings model with time-varying alpha, considers the top 15 percent of alpha-ranked funds as candidate funds and assumes $\lambda = 0.90$.

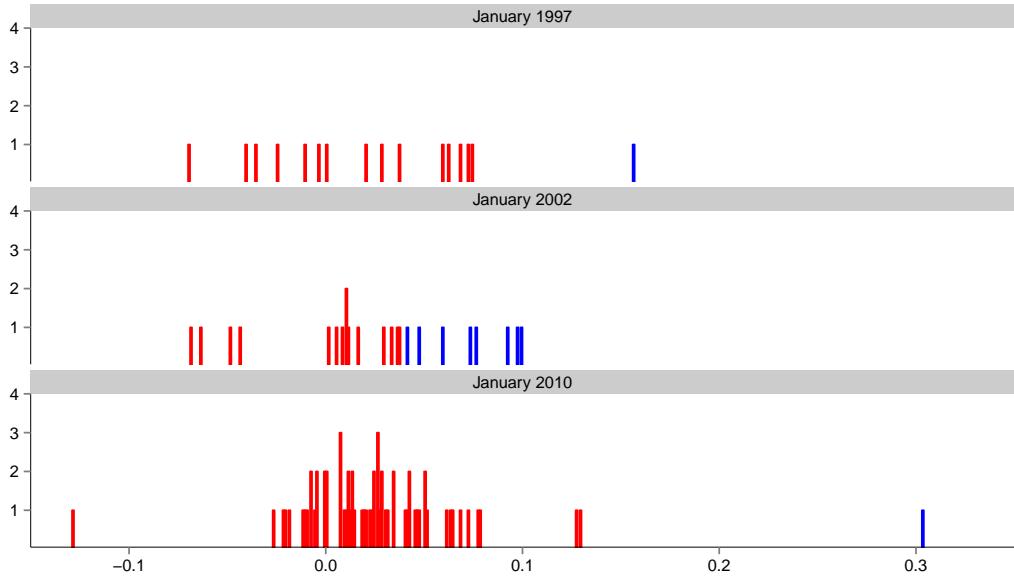


Figure 6: Histograms of the average predictive alpha measure for three different periods in the sample. Red bars indicate funds (chosen from the top 15% of alpha-ranked funds in existence) which are not included in the FCS assuming that $\lambda = 0.90$. Blue bars indicate funds which are included in the FCS.

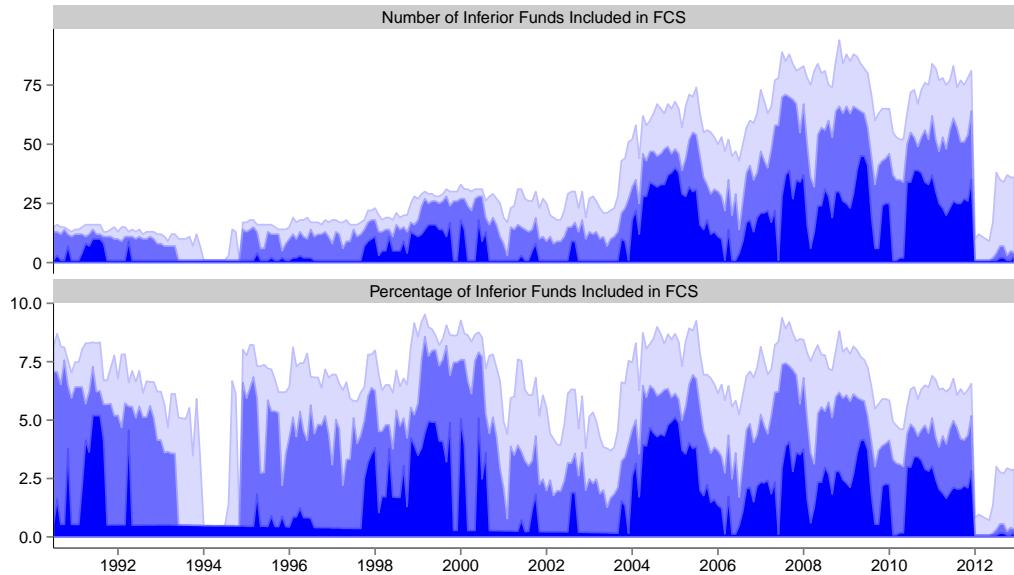


Figure 7: Evolution in the set of funds identified as being inferior by the FCS approach. The top panel shows the evolution in the number of funds included in the tight, medium and wide fund confidence sets of inferior funds. The bottom panel shows the corresponding evolution in the percentage of funds identified as being inferior by the FCS approach. Funds are selected from the bottom 15 percent of alpha-ranked funds using the latent skill-holdings model with a time-varying alpha to compute each fund's alpha estimate. Dark blue areas correspond to $\lambda = 0.90$ (tight set), while lighter areas correspond to $\lambda = 0.50$ (medium set) and $\lambda = 0.10$ (wide set).

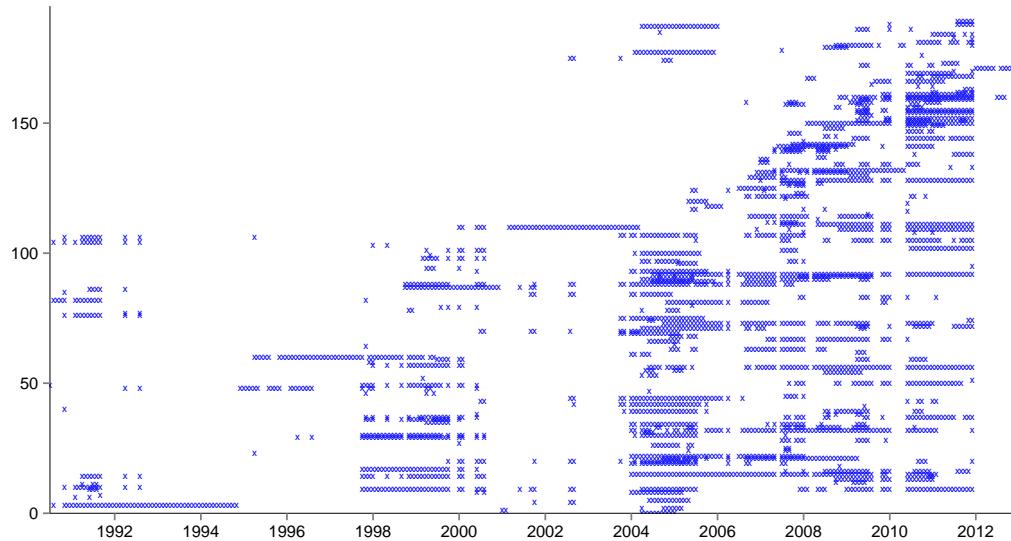


Figure 8: Identification of individual funds with inferior skill by the FCS approach. The plot illustrates the evolution in the composition of individual funds with inferior skills as identified by the FCS approach. Each fund that is included in the FCS at least once during the sample is assigned a unique number on the y-axis and a cross shows when this fund is included. The FCS is based on the latent skill holdings model with time-varying alpha, considers the bottom 15 percent of alpha-ranked funds as candidate funds and assumes $\lambda = 0.90$.

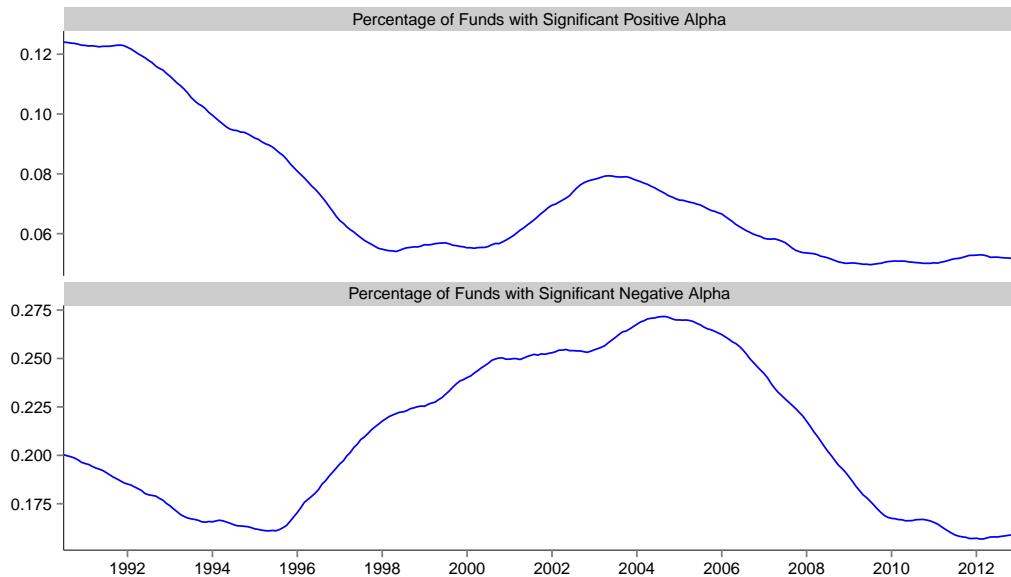


Figure 9: Kernel estimate of the evolution in the proportion of funds with significantly positive (top panel) or negative (bottom panel) alpha estimates, using the Romano-Wolf (2005) methodology. Each month we determine the set of funds with significantly positive (or negative) alphas by applying the Romano-Wolf (2005) step-wise bootstrap procedure to the alpha estimates obtained from the latent skill model that combines data on fund returns and quarterly holdings.

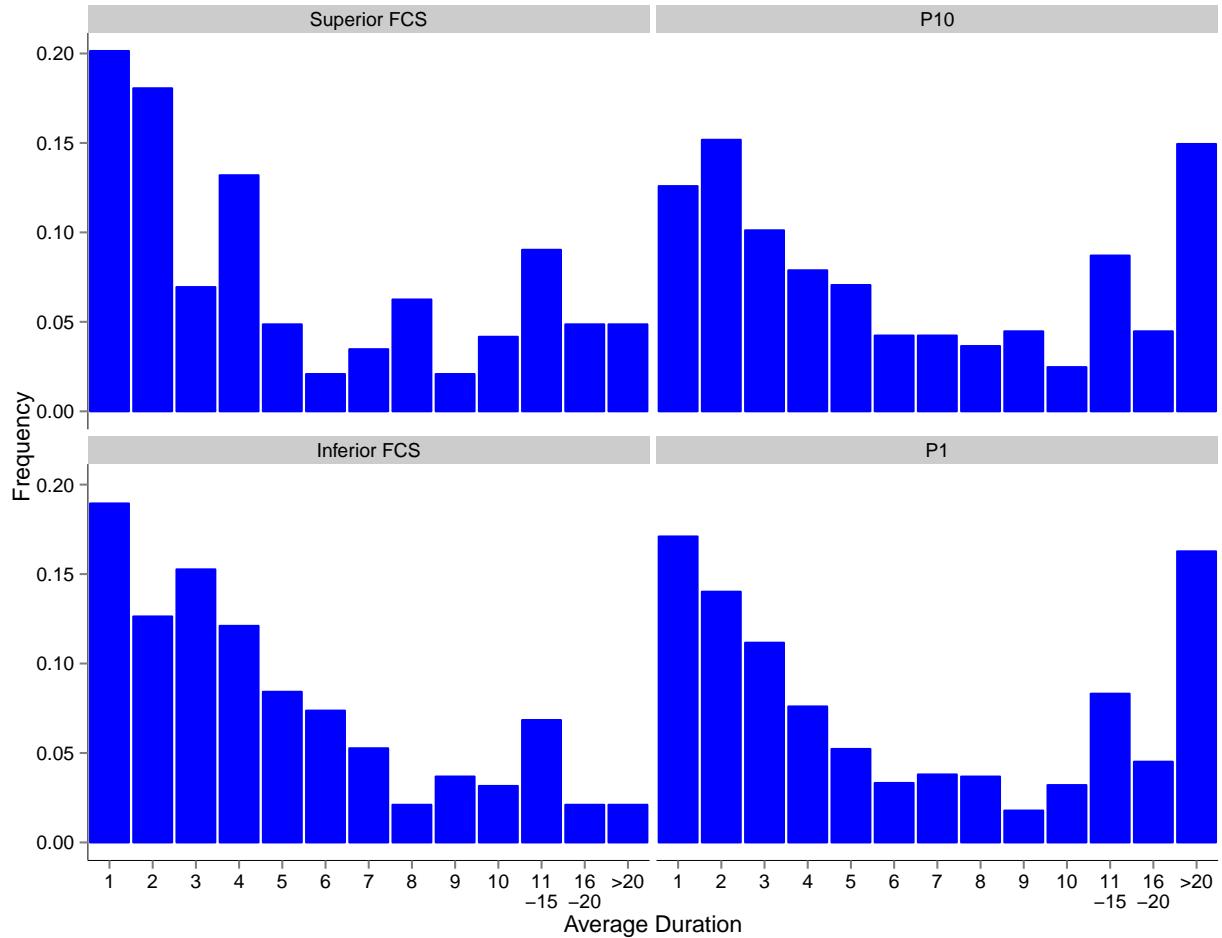


Figure 10: Average duration for funds getting included in different portfolios. The histograms show the distribution of the number of consecutive months where individual funds get included in four different portfolios, namely the FCS portfolio of superior (top left corner) and inferior funds (bottom left), the top (top right) and bottom decile (bottom right) portfolios. Only funds that get included in the portfolio at least once are represented in the plot.

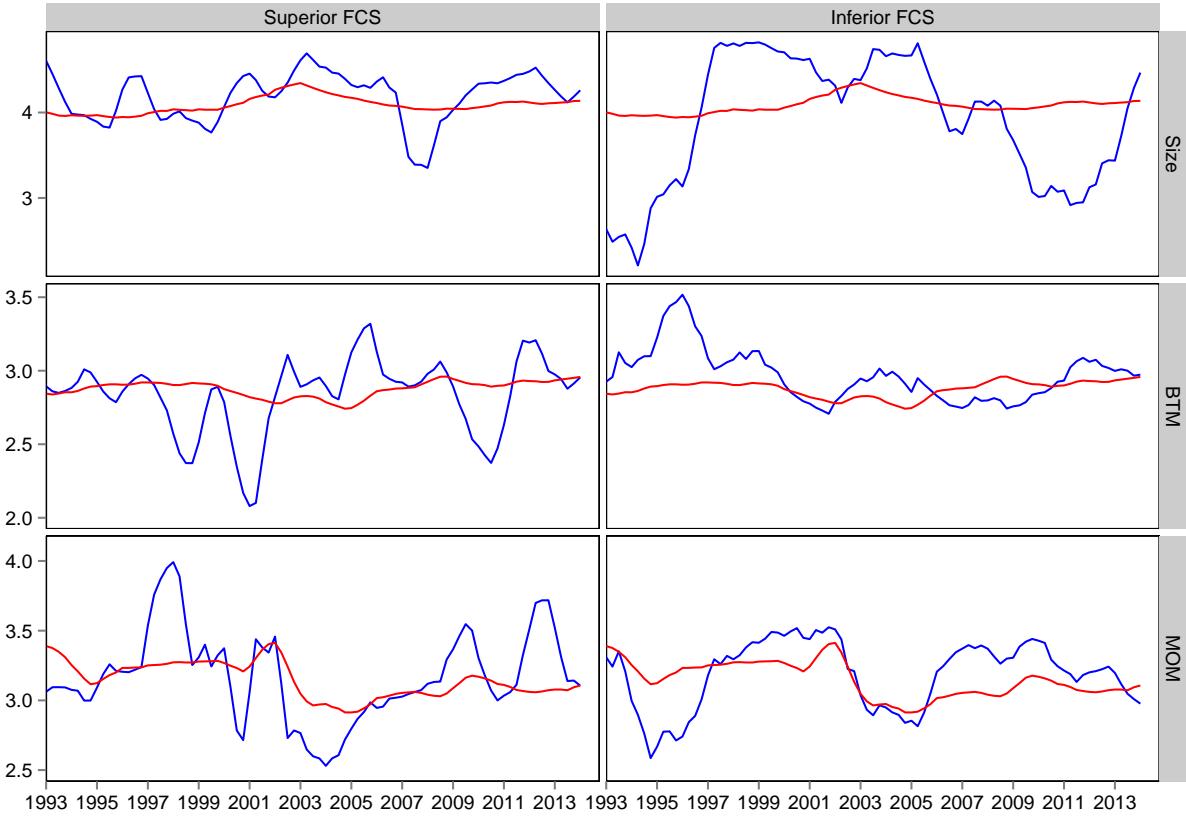


Figure 11: Evolution in loadings on risk factors. Blue lines show equal-weighted loadings for funds in the FCS portfolios. Red lines show loadings for the average fund for comparison. The left column shows one-year rolling averages of exposures to the size, book-to-market and momentum risk factors for the FCS portfolio consisting of superior funds. The right column shows one-year rolling averages of exposures to the same risk factors for the FCS portfolio consisting of inferior funds. The plots assume $\lambda = 0.90$ and are based on the time-varying alpha model that combines data on fund returns and holdings to estimate alphas.

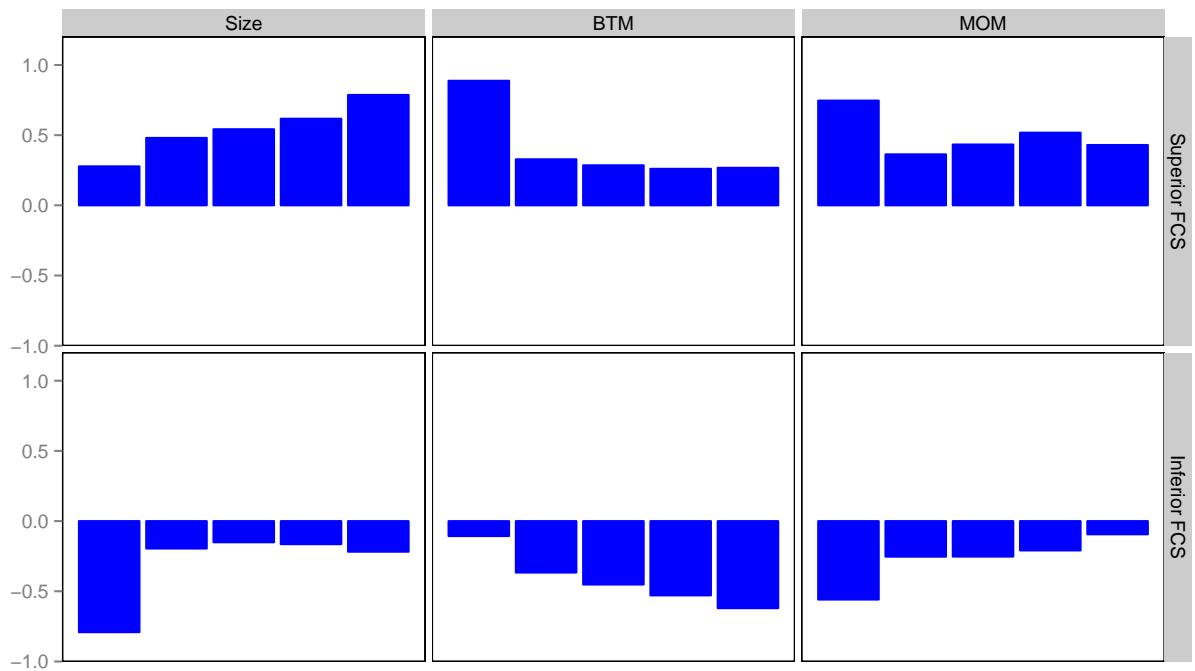


Figure 12: Performance of the superior and inferior FCS portfolios for different realizations of the risk factors. The top row shows histograms of the four-factor alphas for the FCS portfolio of superior funds for different values of the size factor (left panel), book-to-market factor (middle panel) and momentum factor (right panel). The bottom row presents four-factor alpha estimates for the FCS portfolio of inferior funds. The plots assume $\lambda = 0.90$ and are based on the time-varying alpha model that combines data on fund returns and holdings to estimate alphas.

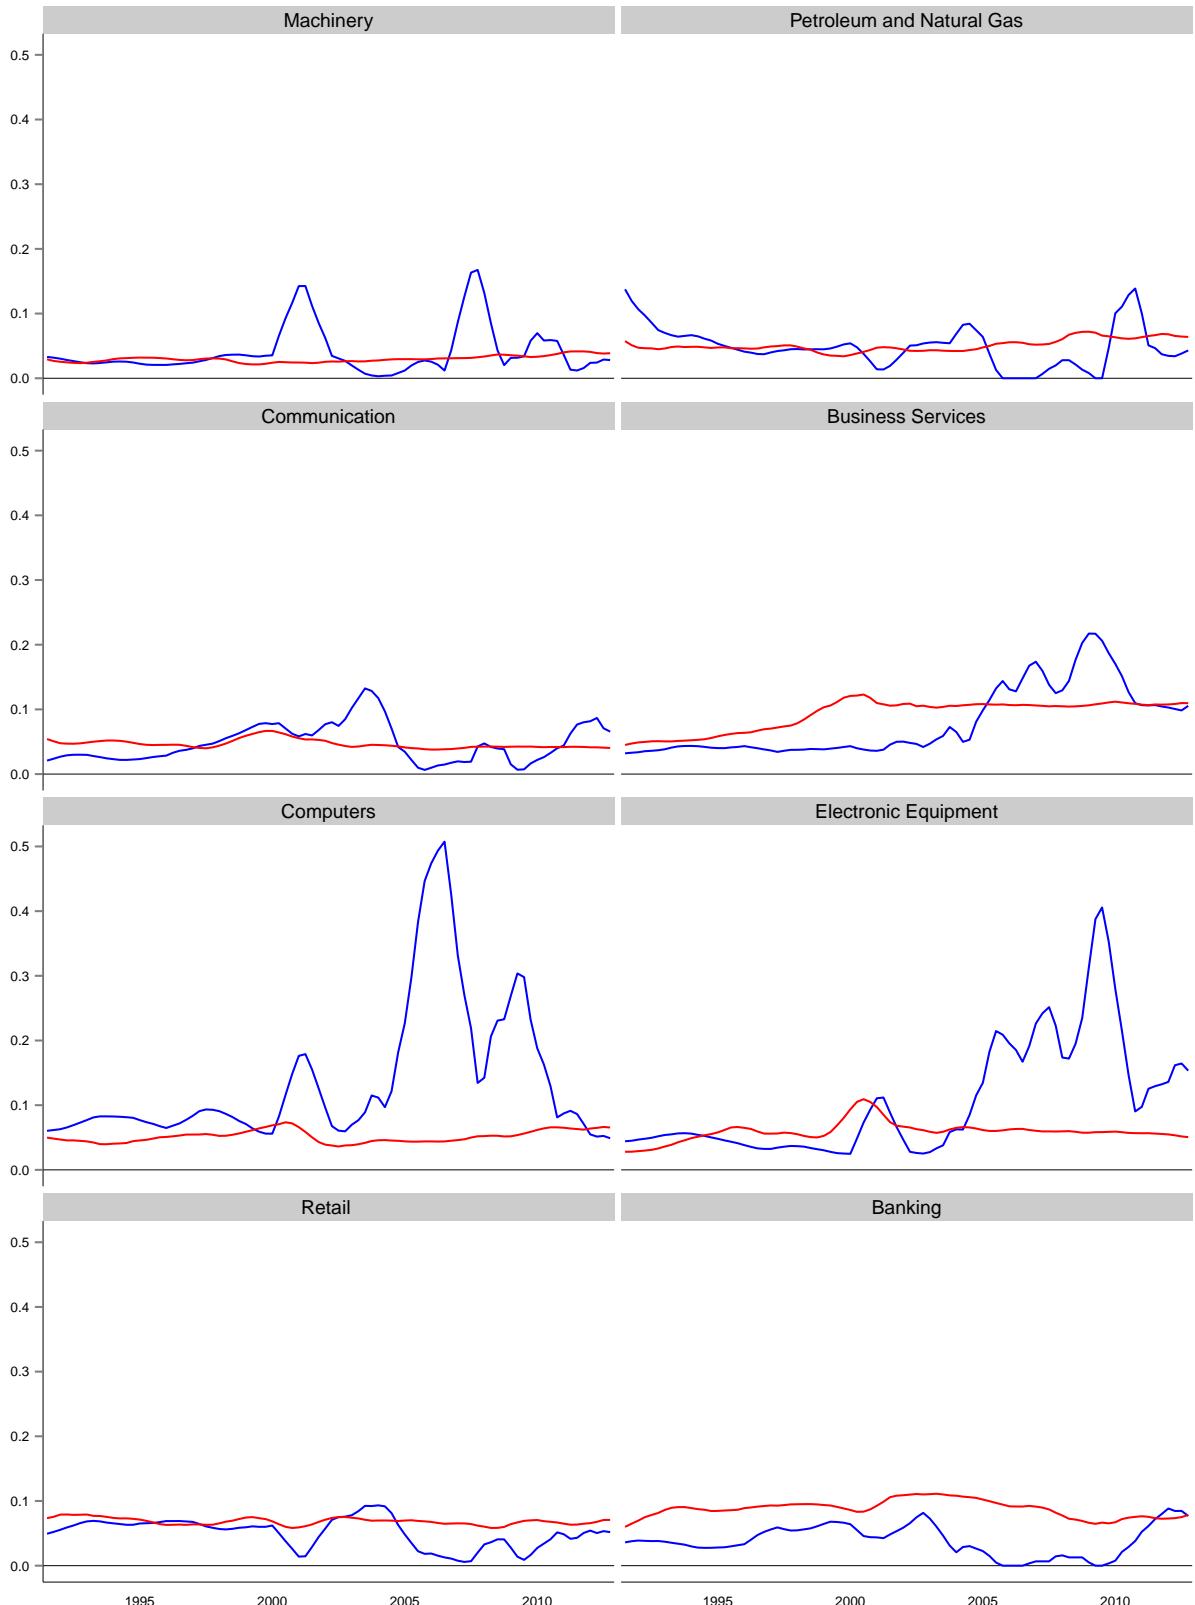


Figure 13: Industry concentrations for the FCS portfolio of superior funds. The figure presents one-year rolling averages of industry concentrations for the FCS portfolio composed of funds identified to have superior performance (blue line). We have chosen the eight industries with the highest maximum industry concentration. For comparison the red lines show the average concentration in the same industries computed across all funds in our sample. The plots assume $\lambda = 0.90$ and are based on the time-varying alpha model that combines data on fund returns and holdings to estimate alphas.

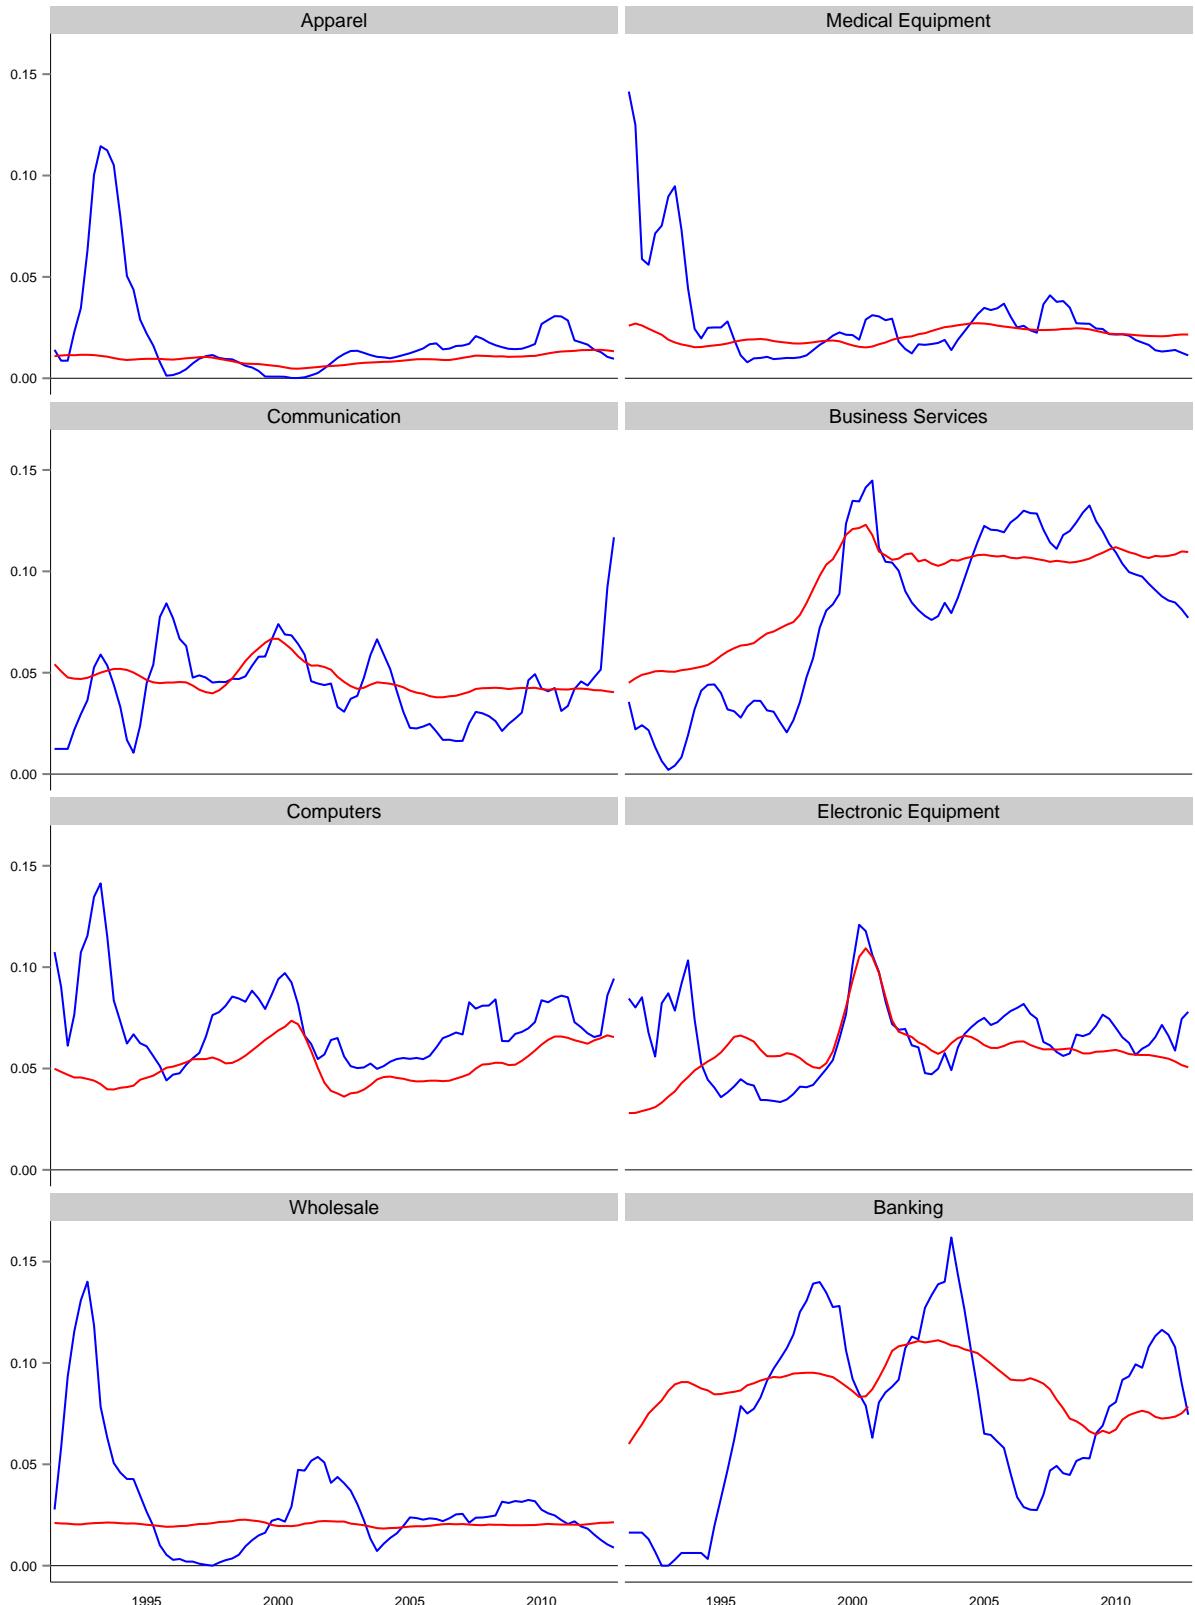


Figure 14: Industry concentrations for the FCS portfolio of inferior funds. The figure presents one-year rolling averages of industry concentrations for the FCS portfolio composed of funds identified to have inferior performance (blue line). We have chosen the 8 industries with the highest maximum industry concentration. For comparison the red lines show the average concentration in the same industries computed across all funds in our sample. The plots assume $\lambda = 0.90$ and are based on the time-varying alpha model that combines data on fund returns and holdings to estimate alphas.