

Best Ideas

Miguel Antón, Randolph B. Cohen, and Christopher Polk*

This Draft: April 2021

*Antón: Department of Finance, IESE Business School, Av. Pearson 21, 08034 Barcelona. Email: man-ton@iese.edu; Cohen: Department of Finance, Harvard Business School, Boston MA 02163. Email: rbcohen@hbs.edu. Polk (corresponding author): Department of Finance, London School of Economics, London WC2A 2AE. Email: c.polk@lse.ac.uk. We are grateful to Ken French for providing us with some of the data used in this study. A very special thanks goes to Tapio Pekkala and Bernhard Silli who were instrumental in producing an earlier version of this paper. Polk gratefully acknowledges financial support from the Paul Woolley Center at the LSE. Antón gratefully acknowledges the financial support of the Department of Economy and Knowledge of the Generalitat de Catalunya (Ref. 2014 SGR 1496), and of the Ministry of Economy, Industry and Competitiveness (Ref. ECO2015-63711-P). We thank seminar participants at Boston College, Exeter University, Harvard University, London School of Economics, University of British Columbia, University of California San Diego, Yale University, and the NBER Behavioral Finance meeting for helpful comments and ML Cooper for copy editing. All errors are our own.

Best Ideas

Abstract

We find that the stocks in which active mutual fund or hedge fund managers display the most conviction towards ex-ante, their “Best ideas,” outperform the market, as well as the other stocks in those managers’ portfolios, by approximately 2.8 to 4.5 percent per year, depending on the benchmark employed. The vast majority of the other stocks managers hold do not exhibit significant outperformance. Thus, the organization of the money management industry appears to make it optimal for managers to introduce stocks into their portfolio that are not outperformers. We argue that investors would benefit if managers held more concentrated portfolios.

JEL classification: G11, G23

I Introduction

When asked to discuss their portfolio, the typical investment manager will identify a position therein and proceed to describe the opportunity and the investment thesis with tremendous conviction and enthusiasm. Frequently the listener is overwhelmed by the persuasiveness of the presentation. This discussion leads to a natural follow-up question: how many investments make up the portfolio? Informed that the answer is, e.g., 150, the questioner will often wonder how anyone could possess such depth of knowledge and passion for so many disparate companies. Pressed to explain, investment managers have been known to sheepishly confess that their portfolio contains a few core high conviction positions—the best ideas—and then a large number of additional positions which may provide less expected abnormal return but which serve to “round out” the portfolio.

This paper attempts to identify *ex ante* which of the investments in managers’ portfolios were their best ideas and to evaluate the performance of those investments. We find that best ideas not only generate statistically and economically significant risk-adjusted returns over time but they also systematically outperform the rest of the positions in managers’ portfolios. We find this result is consistent across many specifications: different benchmarks, different risk models, and different definitions of best ideas. The level of outperformance varies depending on the specification, but for our primary tests falls in the range of 2.8 to 4.5 percent per year. This abnormal performance appears permanent, showing no evidence of subsequent reversal, even several years later.

Interestingly, cross-sectional tests indicate that active managers’ best ideas are most effective in illiquid, growth, momentum stocks, or for funds that have outperformed in the past. We also document similar results in a sample of hedge funds; those managers’ best ideas outperform in a similar fashion. Moreover, we find that the same characteristics that spread performance in best ideas in mutual funds also spread performance in hedge funds. In addition, we document striking heterogeneity in the performance of hedge fund ideas based on fund size. The best ideas of small

funds outperform those of large funds by a remarkable 15 percent per year.

We argue that these findings have powerful implications for our understanding of stock market efficiency. Previous research has generally found that money managers do not outperform benchmarks net of fees. Many years ago, Mark Rubinstein referred to this fact as the efficient-markets faction's "nuclear bomb" against the "puny rifles" of those who argue risk-adjusted returns are forecastable (but see Berk and Green (2004) for reasons to reconsider this way of thinking). Other work, e.g. Wermers (2000), has shown quite modest outperformance of around one percent per year for the stocks selected by managers (ignoring all fees and costs). In contrast, our paper makes an important contribution by presenting evidence that the typical active manager can select stocks that deliver economically large abnormal returns (relative to our risk models).

Consequently, this paper's findings are relevant for the optimal behavior of investors in managed funds. Our results suggest that while the typical manager has a small number of good investment ideas that provide positive alpha in expectation, the remaining ideas in the typical managed portfolio add little or no alpha. Managers have clear incentives to include zero-alpha positions. Without them, the portfolio would contain only a few names, leading to increased volatility, price impact, illiquidity, and regulatory/litigation risk. In particular, adding additional stocks to the portfolio can reduce volatility, increasing portfolio Sharpe ratio. Perhaps most importantly, adding names enables the manager to take in more assets, and thus draw greater management fees. The manager gains from diversifying the portfolio, but the typical investor is often worse off because they likely have only a modest fraction of their assets in any one fund.

Based on these observations, we examine optimal decentralized investment when managerial skill is consistent with our evidence on the performance of best ideas. We show that under realistic assumptions, investors can gain substantially if managers choose less-diversified portfolios that tilt more towards their best ideas.

II Related Literature

There are several reasons why examining total portfolio performance may be misleading concerning stock-picking skills. First, manager compensation is often tied to the size of the fund's holdings. As a consequence, managers may have incentive to continue investing fund capital after their supply of alpha-generating ideas has run out. This tension was studied by, among others, Berk and Green (2004). Second, the very nature of fund evaluation may cause managers to hold some or even many stocks on which they have neutral views concerning future performance. In particular, since managers may be penalized for exposing investors to idiosyncratic risk, diversification may cause managers to hold some stocks not because they increase the average risk-adjusted return on the portfolio but simply because these stocks reduce overall portfolio volatility. Third, open end mutual funds provide a liquidity service to investors. Edelen (1999) provides evidence that liquidity management is a major concern for fund managers and that performance evaluation methods should take it into account. Alexander et al. (2007) show explicitly that fund managers trade-off liquidity against valuation motives when making investment decisions. Finally, even if managers were to only hold stocks that they expect will outperform, it is likely that they believe that some of these bets are better than others.

Sharpe (1981); Elton and Gruber (2004); and van Binsbergen, Brandt, and Koijen (2008) study the way myopic decision rules for decentralized investment management can lead to suboptimal outcomes. Theoretical work by Van Nieuwerburgh and Veldkamp (2010) has highlighted the importance of specialization in managerial information acquisition. Their study shows that returns to such specialization imply that investors should not hold diversified portfolios. Our results may help to shed some light on Van Nieuwerburgh and Veldkamp's (2010) conclusions.

There are several empirical papers with findings related to ours. Evidence that managers select stocks well can be found in Wermers (2000), Chen, Jegadeesh, and Wermers (2000), Cohen, Gompers, and Vuolteenaho (2002) and Massa, Reuter, and Zitzewitz (2010). Evidence that man-

agers who focus on a limited area of expertise outperform more than the typical manager can be found in Kacperczyk, Sialm, and Zheng (2005). Cremers and Petajisto (2009) demonstrate that the share of portfolio holdings that differ from the benchmark (what they define as active share) forecasts a fund's abnormal return.

Other research suggests that extracting managers' beliefs about expected returns from portfolio holdings might be useful. For example, Shumway, Szeffler, and Yuan (2009) show that the precision of the implied beliefs from a manager's holdings concerning expected returns helps to identify successful managers. Pomorski (2009) shows that when multiple funds belonging to the same company trade the same stock in the same direction, that stock outperforms. Jiang, Verbeek, and Wang (2014) find that stocks that are overweighted by the mutual fund industry in the aggregate outperform. Wermers, Yao, and Zhao (2012) document that stocks that are held in greater weight by previously-successful managers than by less-successful managers tend to perform well.

Our paper has two primary novel contributions.

1. While we are not the first to show that managers are capable of successfully selecting outperforming stocks, we think we are the first to document that the typical manager can identify a set of stocks that outperform over and above average levels of trading costs and fees. Wermers (2000) and similar subsequent work found managers picking stocks that deliver 1-1.5% alpha, leaving net returns to investors below the index. We show that in the stocks managers identify via the weights they choose as their best ideas, outperformance is much higher, 2.8-4.5% per year, easily enough to deliver substantial alpha after fees and trading costs. Later work, some referenced in the previous paragraph, finds effects as large as ours if we select stocks by taking advantage of the information embedded in the holdings of all mutual fund managers. These "wisdom of crowds" results are in a separate category from ours, as we show that ordinary individual managers – not superstar managers, or thousands of managers working in concert, but just single, average managers working on their own – have the ability to select one or more stocks that will likely generate net alpha for investors.

2. Our paper explores the implications of our novel and surprising finding that a typical active manager's value added primarily comes from their handful of best ideas rather than from their entire 100-plus-positions portfolio. We develop a simple framework that illustrates how large the benefits to investors can be if managers concentrate their portfolios in their best ideas. Under reasonable conditions consistent with our findings, we show that investing in just the managers' best ideas will double the Sharpe ratio of an investor's portfolio relative to the case where managers maximize the Sharpe ratio of their individual funds. This effect is much larger than might be suggested by standard intuition.

III Methodology

We proceed in three stages. First, for each stock in a manager's portfolio, we attempt to estimate from the portfolio choices the manager made what CAPM alpha and information ratio the manager believed that stock would deliver. Then, we label each manager's holding with the highest estimated information ratio as their "best idea," and form each period the portfolio of all managers' best ideas. Finally, we test the performance of this "best ideas" portfolio.

To formally motivate our approach to extracting the best ideas of portfolio managers, we first consider a simple portfolio optimization problem. A manager is selecting a portfolio from a set of N risky assets. We assume their goal is to maximize portfolio Sharpe ratio. Let the vector of portfolio weights be λ_t . Then the manager must solve the problem:

$$\max_{\lambda_t} \lambda_t' (E_t \mathbf{R}_{t+1} - R_{f,t+1} \iota) - \frac{k}{2} \lambda_t' \mathbf{\Omega}_t \lambda_t$$

where ι is a vector of ones, $(E_t \mathbf{R}_{t+1} - R_{f,t+1} \iota)$ is the vector of expected excess returns on the N risky assets over the riskless interest rate, $R_{f,t+1}$, and $\mathbf{\Omega}_t$ is the return variance-covariance matrix. The well-known solution to this maximization problem is

$$\lambda_t = \frac{1}{k} \mathbf{\Omega}_t^{-1} (\mathbf{E}_t \mathbf{R}_{t+1} - R_{f,t+1} \mathbf{1})$$

One can reverse engineer this solution to find that the manager's expected excess returns are $\mu_t = (\mathbf{E}_t \mathbf{R}_{t+1} - R_{f,t+1} \mathbf{1}) = k \mathbf{\Omega}_t \lambda_t$. This formula holds for both the market (where μ_{Mt} is a vector of CAPM expected excess returns, k_M is the market's aggregate risk aversion, and λ_{Mt} are market weights) and any fund manager (where μ_{ft} is a vector of portfolio manager f 's expected excess returns, k_f is the manager's risk aversion, and λ_{ft} is that manager's portfolio weights): $\mu_{ft} = k_f \mathbf{\Omega}_t \lambda_{ft}$ and $\mu_{Mt} = k_M \mathbf{\Omega}_t \lambda_{Mt}$.

Although (as noted later) our approach does not require this assumption, for simplicity of explication, we assume that risk aversion is homogeneous, $k_f = k_M = k$, and measure the difference between the vector of the manager's subjective expected excess returns and the vector of CAPM expected excess returns: $\mu_{ft} - \mu_{Mt} = k \mathbf{\Omega}_t (\lambda_{ft} - \lambda_{Mt})$. This difference, $\mu_{ft} - \mu_{Mt}$, is just the vector of the manager's subjective CAPM alphas. Renaming the portfolio alpha vector α_{ft} and defining $\mathbf{tilt}_{ft}^{market} = \lambda_{ft} - \lambda_{Mt}$ as the vector of over/underweights of the manager's portfolio, we have:

$$\alpha_{ft} = k \mathbf{\Omega}_t \mathbf{tilt}_{ft}^{market}$$

In order to make the problem of covariance estimation more tractable we make some strong assumptions. Most important, we assume that all off-diagonal elements of the covariance matrix are due to a single factor, the market return. We can then decompose the covariance matrix into orthogonal components: $\mathbf{\Omega}_t = \mathbf{\Sigma}_t + \sigma_{Mt}^2 \mathbf{B}_t$. $\mathbf{\Sigma}_t$ is a diagonal matrix with the idiosyncratic variance (σ_{it}^2) of each stock as its diagonal element, σ_{Mt} is the volatility of the market, and $\mathbf{B}_t$ is a matrix whose (i, j) element equals $\beta_i \beta_j$ where β_k is the market beta of the k th stock.

We then make the additional simplifying assumption that all mutual funds in our sample have a fund beta, β_f , close to 1.0, the market beta, β_M . When we post-multiply the $\mathbf{B}_t$ matrix by the $\mathbf{tilt}_{ft}^{market}$ vector, we get a column vector whose k th element is $\beta_k (\beta_f - \beta_M)$; if each portfolio has unit beta these terms all equal zero and we are left with: $\alpha_{ft} = k \mathbf{\Sigma}_t \mathbf{tilt}_{ft}^{market}$. We use this expression

to estimate subjective alphas and therefore to select the ex-ante best idea(s) for manager f that has the highest information ratio. We call this approach “market alpha.”

$$\begin{aligned}\alpha_{ft}^{market} &= k \Sigma_t \mathbf{tilt}_{ft}^{market} \\ IR_{ift}^{market} &= \alpha_{ift}^{market} / \sigma_{it} \\ best_{ft}^{market\ alpha} &= \arg \max_i IR_{ift}^{market}\end{aligned}$$

Using a multi-factor model or assuming funds have betas far from 1.0 are mathematically straightforward extensions, but they clutter the explication. Moreover, given the difficulty of accurate covariance estimation, these complications may not increase the accuracy of our estimation of managers’ subjective alpha. For similar reasons, we confirm that robustness checks based on the assumption that all stocks have equal idiosyncratic risk provide qualitatively similar results.

Many investment managers are limited in the set of stocks that they consider investing in – for example, a manager’s universe may consist only of large-cap stocks or of a particular industry sector. In these cases, the relevant tilt is arguably not the overweight/underweight relative to the market, but rather relative to the value-weight portfolio of stocks in the manager’s universe. We choose as an alternative to our market alpha method, an approach we call “portfolio alpha.” Here we simply use $\mathbf{tilt}_{ft}^{portfolio}$, the difference between the manager’s weight vector, λ_{ft} , and the value-weight portfolio consisting only of stocks the manager actually holds, λ_{fVt} :

$$\begin{aligned}\mathbf{tilt}_{ft}^{portfolio} &= \lambda_{ft} - \lambda_{fVt} \\ \alpha_{ft}^{portfolio} &= k \Sigma_t \mathbf{tilt}_{ft}^{portfolio} \\ IR_{ift}^{portfolio} &= \alpha_{ift}^{portfolio} / \sigma_{it} \\ best_{ft}^{portfolio\ alpha} &= \arg \max_i IR_{ift}^{portfolio}\end{aligned}$$

Although portfolio alpha is simplistic, it has powerful intuitive appeal. In particular, suppose

that a manager approaches portfolio selection the following way: first, identify a set G (“good”) of stocks that they expect will perform well. Then, choose weights that maximize Sharpe Ratio subject to the condition that no stock has negative weight, and all stocks not in G have zero weight. If a manager uses this approach, using portfolio alpha as defined in the equations above will provide an alpha measure that recovers the ranking of the manager’s subjective alpha. Based on our conversations with money managers, we believe the heuristic described above is a good approximation of manager behavior.

In our analysis, we typically report the performance of best ideas based on both a market alpha and a portfolio alpha perspective. Our assumption of homogeneous risk aversion is without loss of generality. As our focus is only on the ordinal aspects of a comparison between λ_{ft} and either λ_{Mt} or λ_{fVt} , a manager’s best idea does not depend on whether they are more or less risk averse than the market.

In summary, we call the output of the first approach the market alpha and the second the portfolio alpha. In the market method, a stock that is held in the same proportion it has in the market is viewed as having zero expected alpha. In the portfolio method, a zero-alpha stock is one held with a weight equal to its value weight within a portfolio of the stocks the manager holds. For non-zero-alpha stocks, the magnitude of alpha depends not only on the degree of tilt towards that particular stock but also on the magnitude of the stock’s market-model idiosyncratic variance.

IV Data and Sample

Our stock return data comes from CRSP (Center for Research for Security Prices) and covers assets traded on the NYSE, AMEX and NASDAQ. We use the mutual fund holdings data from Thompson Reuters. Our sample consists of US domestic equity funds that report their holdings in the period from January 1983 to December 2018.¹ The holdings data are gathered from quarterly

¹Though the dataset begins in 1980, we start our analysis in 1983 so that the analysis discussed in Section V.C and presented in Table IV that sorts best ideas into terciles does not result in thin portfolios with only a handful of stocks. We have confirmed that our key result discussed in Section V.B and presented in Table II) that measures

filings of every U.S.- registered mutual fund with the Securities Exchange Commission. The mandatory nature of these filings implies that we can observe the holdings of the vast majority of funds that are in existence during that period. For a portfolio to be eligible for consideration, it must have total equity investments exceeding \$5 million and at least 5 recorded holdings.²

A crucial assumption of our analysis is that fund managers try to maximize the information ratio of their portfolios. Therefore, we exclude portfolios that are unlikely to be managed with this aim in mind, such as index or tax-managed funds.³ We also exclude international funds from the sample. We identify best ideas as of the true holding date of the fund manager's portfolio as we are primarily interested in whether managers have stock-picking ability, not whether outsiders can piggyback on the information content in managers' holdings data.⁴ Of particular note, we focus on the top 25% most active mutual funds, defined as those funds whose maximum position-level information ratio defined above is in the top 25% of all corresponding information ratios at that point in time. All of our results are robust to reasonable variation in this cutoff; we report findings based on different cutoffs in the Internet Appendix.

Table I Panel A provides summary statistics every five years for our sample of mutual fund portfolios over the 35-year period under consideration. The table highlights the impressive growth of the industry, the result of not only growth in the market itself but also increased demand for equity mutual fund investment. While the number of funds in our active mutual fund sample increases tenfold from 1983 to the end of 2018, assets under management increase from \$54 billion to \$3.8 trillion in the same time span. Columns four and five indicate that active mutual funds as a

the performance of all best ideas and thus does not require sorting are robust. For example, Table II Panel A reports that the six-factor alpha of the best ideas portfolio starting in 1983 is 37 basis points per month (t -stat of 3.45). The six-factor alpha of a best ideas portfolio starting in 1980 is 33 basis points per month (t -stat of 3.36).

²This minimum requirement on the amount of net assets and number of holdings is standard in the literature and imposed to filter out the most obvious errors present in the holdings data as well as incubated funds.

³We focus on active equity funds as follows. We screen and exclude funds whose names include INDEX, Index, IDX, Idx, S&P, s&p, Fixed, FIXED, TAX, tax, Tax, CONVERTIBLE, Convertible, annuity, ANNUITY, ANN, ann, VAR, Var, CONV, or Conv. Furthermore, we only include funds that have the investment objective code of 2, 3, 4, 7, or 8 and that have two-thirds of their AUM in common shares.

⁴Our work also documents that best-ideas trading strategies based only on holdings information as of the date the positions are made public also generate economically and statistically significant abnormal returns. See Figure 4 for direct evidence on this question.

whole have grown to be dominant investors in U.S. equity markets with the typical fund managing more than 2 billion and holding roughly 160 stocks. Interestingly, the average number of stocks in these mutual fund portfolios has more than doubled in our sample period. In summary, our analysis covers a substantial segment of the professional money management industry that in turn scans a substantial part of the U.S stock market for investment ideas.

Table I Panel A documents that our sample of active managers tends to hold very close to the typical NYSE stock in terms of the well-know Fama and French (2015) characteristics. These preferences are rather stable over time. Table I Panel B reports summary statistics for the stocks we identify as the top ex ante best idea while Table I Panel C reports the same statistics for the top five best ideas. Best ideas tend to be much larger stocks with relatively low book-to-market ratios, relatively high profitability, and relatively high asset growth. As a consequence, though these best ideas are a small fraction of the market in number, in total, they make up a non-trivial portion of US stock market value. The panel highlights that the typical fund's stake in their best idea is roughly 1.3% of the market of the firm in question and 10.6% of the fund's AUM.

V Empirical Results

A The Distribution of Best Ideas

In theory, the number of best ideas that exist in the industry at any point in time could be as many as the number of managers or as few as one (if each manager had the same best idea). Of course, this latter case is quite unlikely since mutual fund holdings make up a substantial proportion of the market. Therefore, massive overweighting of a stock by mutual funds would be difficult to reconcile with financial market equilibrium. In Figure 1, we show the way the distribution of best ideas across managers based on our market measure of information ratio, $\mathbf{IR}_{ft}^{market}$, has changed over time. The solid line with circles in Figure 1 indicates that best ideas generally do not overlap across managers. Over the entire sample period, approximately three-

fourths (72%) of best ideas are a best idea of only one manager at the relevant time. We note that this statistic has declined over our sample from 90% in 1983 to 65% in 2018. The solid line with squares shows that best ideas are considered as such by two managers around 14% of the time. A stock is the best idea of six or more funds only 5% of the time. Figure 1 shows that these statistics do not vary that much over the sample.

Clearly, managers' best ideas are not entirely independent. However, the best idea portfolios we identify do not consist of just a few names that are hot on Wall Street. Rather, this portfolio represents the opinions of hundreds of managers each of whom independently found at least one stock about which they appeared to have real conviction. Figure 1 shows that the percentage of best ideas that are the best idea of only one manager has decreased over the years but is always greater than 50%.

Figure 2 graphs the way key aspects of our best ideas measure—tilt, implied information ratio, and implied alpha—vary over time. The top panel depicts the evolution of the median market and portfolio tilts, the middle panel reports the corresponding value of the associated information ratio measure (the measure with which we identify the best ideas), and the bottom panel shows the corresponding implied alpha. The typical market tilt associated with best ideas averages around 7.4% while the typical portfolio tilt averages around 4.9%. The corresponding information ratio averages 1.5% for the market tilt measure and 1.1% for the portfolio tilt measure. As the graphs show, variation in these measures clearly reflects trends in idiosyncratic volatility over time.⁵ This variation is a desirable feature of our alpha measures: A 2% tilt away from the benchmark in 2000 is a stronger sign of conviction than a 2% tilt in 1997, since idiosyncratic risk has risen in between. The resulting alphas average 0.35% and 0.26%.

⁵Campbell, Lettau, Makiel, and Xu (2001) document a positive trend in idiosyncratic volatility during the 1962 to 1997 period. See Brandt, Brav, Graham, and Kumar (2010) for post-1997 evidence on this time-series variation.

B The Features and Performance of Best Ideas

We measure the performance of best ideas using two approaches. Our primary approach is to measure the out-of-sample performance of a portfolio of all active managers' best ideas.⁶ Each best idea in the portfolio is equal-weighted (if more than one manager considers a stock a best idea, we overweight accordingly). Results are qualitatively similar if we equal-weight unique names in the portfolio, if we weight by market capitalization, or if we weight by the amount of dollar invested in the best idea. The portfolio is rebalanced on the first day of every quarter to reflect new information on the stock holdings of fund managers, and its performance is tracked until the end of the quarter.

Each best ideas portfolio differs according to which of the two tilt measures we use to identify best ideas. We examine not only the excess return on the portfolio but also a matched version following Daniel, Grinblatt, Titman, and Wermers (2000). Our final approach examines "best-minus-rest" portfolios, where for every manager, we are long their best idea and short the remaining stocks in the manager's portfolio (excluding the top ten ideas and with the weights for the "rest" of the portfolio being proportional to the manager's weights). Thus, for each manager, we have a style-neutral best-idea bet, which we as before aggregate over managers by equal-weighting.⁷ We then track the monthly performance of these six portfolios (three for each tilt measure) over the following three months and rebalance thereafter.

We benchmark these six portfolios in two different ways. We choose these models to reflect industry standards in fund evaluation and to make our results comparable to the findings of previous work in the literature. We first examine the alpha with respect to the Fama French five-factor model. We also report performance results measured by a six-factor specification, which adds Carhart's (1997) momentum factor to the model. All factor return data are gathered from

⁶We require stocks to have price greater or equal than \$5 at portfolio formation.

⁷Note that our best-minus-rest approach has at least one attractive benefit: By comparing the manager's best idea to other stocks in the manager's portfolio, the best-minus-rest measure tends to cancel out most style and sector effects that might otherwise bias our performance inference.

Kenneth French's website.⁸

The first two rows of Table II Panel A benchmark the best ideas of active fund managers against the five- and six-factor models. Our estimates of five-factor alpha are 47 basis points (t -statistic of 4.09) and 43 basis points (t -statistic of 3.89) for the market and portfolio alpha measures respectively. Our six-factor estimates are 37 basis points (t -statistic of 3.34) and 35 basis points (t -statistic of 3.29) for the market and portfolio alpha measures respectively.

One concern is that the factor model may not perfectly price characteristic-sorted portfolios. The small-growth portfolio and the large-growth portfolio have three-factor alphas of -34 bps/month (t -stat -3.16) and 21 bps/month (t -stat of 3.20) in Fama and French (1993). As Daniel, Grinblatt, Titman, Wermers (1997) (DGTW) point out, this fact can distort performance evaluation. For example, the passive strategy of buying the S&P 500 growth and selling the Russell 2000 growth results in a 44 bps/month Carhart alpha (Cremers, Petajisto, and Zitzewitz; 2012). As a consequence, we also adjust the returns on the best-ideas strategy using characteristic-sorted benchmark portfolios as in DGTW. Specifically we assign each best idea to a passive portfolio according to its size, book-to-market, and momentum rank and subtract the "passive" portfolio's return, $r_{DGTW,t}$, from the best idea's return. The DGTW excess return measure is thus simply $r_{p,t} - r_{DGTW,t}$. The third and fourth rows of Table II Panel A show the mean of the benchmarked return along with results from five- and six-factor regressions. The characteristic-adjusted results remain economically and statistically significant. In particular, the six-factor alphas on the DGTW-adjusted returns are 26 basis points (t -stat of 2.91) and 29 basis points (t -stat of 3.29) respectively.

The analysis in the last two rows of Table II Panel A indicates that missing controls are probably not responsible for the alphas we measure by examining the performance of a best-minus-rest strategy. Unless the best ideas of managers systematically have a different risk or characteristic profile than the rest of the stocks in their portfolios, this strategy controls for any unknown style effects that the manager may possibly be following. For both measures, alphas

⁸http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

are comparable to our two other approaches shown in Table II. More importantly, this analysis confirms that the abnormal performance of the typical manager's overall portfolio is not spread evenly across their positions. As an additional useful benchmark, Table II Panel B measures the performance of a portfolio that holds all of the portfolios of the managers in our sample, in proportion to their investment weights. We find that the six-factor alpha of a portfolio of "All Ideas" is only 6 basis points and statistically insignificant.

Our analysis has focused on the top 25% best ideas across the universe of active managers in order to make sure we were not examining passive funds, sometimes labeled "closet indexers." In the Internet Appendix, Panels A through F of Table A.1 document that our findings concerning the performance of best ideas generally hold as we vary this threshold from 50% to the top 12.5% of active tilts.

Table III repeats the best-minus-rest analysis of Table II Panel A but replacing the best idea in the long side of the bet with the manager's top three ideas (Panel A) or top 5 ideas (Panel B). These top three/five positions are weighted within funds by the size of the manager's position in the stock and then equally-weighted across managers. We find that after generalizing what managers feel to be their top picks, the results continue to show economically and statistically significant performance. Consistent with the idea that managers' tilts reflect their views concerning stocks' prospects, the alpha of the trading strategy is smaller as we include the lower-ranked stocks. Consistent with diversification benefits, the standard error of the estimate is usually lower as more stocks are included on the long side.

We examine more precisely how views concerning alpha that are implicit in managers' portfolio weights line up with subsequent performance. We repeat the analysis replacing every manager's best idea with their second-best idea. We then move on to the third-best idea, and so on, down to the tenth-ranked idea. Figure 3 plots how the six-factor alphas evolve when one moves down the list of best ideas for the market tile measure of alpha along with two standard error bounds. The figure is striking: the point estimates decline as we move down managers' rankings. We are able to reject the hypothesis that these ideas are greater than zero for each of the top-five ideas,

identified using our market alpha measure. The estimates for ideas six through ten are less than that of any of the top-five ideas, and often substantially so.

Figure 4 plots the returns on the best ideas portfolio in event time where the post-formation long-horizon returns are measured as in Cohen, Polk, and Vuolteenaho (2009) and adjusted for risk using the six-factor model.⁹ Panel A reports the result for the full sample while Panel B reports the results for the second half of the long-horizon sample. The figure shows that the superior performance of best ideas is not transitory in nature.¹⁰ The buy-and-hold abnormal returns of the stocks in our best ideas portfolio does not revert in the months or even years after first appearing in the portfolio. If one were to buy the best-ideas portfolios of Table II and hold those bets for the next decade, the result long-horizon abnormal performance would have been slightly over 20%. This finding indicates that these stocks are dramatically underpriced relative to the six-factor model. We find similar patterns for the long-horizon performance of best ideas in the second half of the sample indicating that our finding is robust in that way.

C Where are best ideas most effective?

In this subsection, we examine the extent to which the performance of best ideas varies as a function of stock or fund characteristics. In particular, each month we sort the best ideas portfolio into tritiles based on a particular characteristic and report the differences in six-factor alphas in Table IV.¹¹ Our first variable is stock illiquidity which we measure using the bid-ask spread. We expect stronger outperformance of best ideas among illiquid stocks. Consistent with this view, we find that relatively illiquid best ideas outperform relatively liquid best ideas by 59 basis points per month with a t -statistic of 2.56.

⁹We measure abnormal return at post-formation month T using a portfolio strategy that invests in T portfolios: the portfolio based on month $t-1$ best ideas, the portfolio based on $t-2$ best ideas, ... , and the portfolio based on month $t-T$ best ideas. Furthermore, the weight on each of these different portfolios declines exponentially— ρ to the power of months from the sort minus one. The parameter ρ is the Campbell-Shiller parameter that we set to have an annual value of 0.95. See Cho and Polk (2020) for more sophisticated ways to measure price-level distortions.

¹⁰To ensure that the number of data points for each post-formation month is the same, we form our last portfolio in 2008.

¹¹Differences in five-factor alphas are very similar and available on request.

The next two variables with which we sort stocks within the best ideas portfolio are growth and momentum. In terms of the first, growth stocks are arguably often overpriced, so, if a best idea is a growth stock, then the manager must really believe in it. Consistent with this view, we find that best ideas with relatively higher market-to-book ratios outperform by 86 basis points per month with a t -statistic of 3.94. For momentum, the typical institutional investor tilts away from that characteristic, perhaps in part because of tracking error concerns.¹² We find that best ideas with relatively higher stock momentum outperform by 76 basis points (t -statistic of 3.21). Finally, we sort stocks into tritiles based on the fund manager's 36-month four-factor (Carhart, 2001) alpha. The best ideas of funds with relatively higher alpha outperform funds with relatively lower alpha by 38 basis points a month with an associated t -statistic of 2.23.

We measure the joint economic and statistical significance of these variables by combining them into a single composite sort variable. To do so, we convert each variable into percentile ranks, average these ranks for each stock into a composite variable ($COMP$), and then sort stocks based on the resulting average. We find that stocks with a high $COMP$ outperform stocks with a low $COMP$ by an impressive 91 basis points a month with a corresponding t -statistic of 3.92. To measure the importance of the average incremental effect of each of the four variables, we repeat the analysis but first orthogonalizing each characteristic to each other. The resulting orthogonal composite variable ($COMP - O$) delivers a still impressive spread of 66 basis points a month (t -statistic of 2.43) even though it purposefully discards the common component of these four variables.

D Robustness

Our findings are robust to reasonable variation in our methodology. In terms of subsample analysis, Appendix Table A.III documents that our findings do vary with the business cycle. Consistent with work by Kacperczyk, van Nieuwerburgh, and Veldkamp (2015) who argue that

¹²Gompers and Metrick (2001) document a strong negative cross-sectional relation between momentum and institutional ownership. Lou, Polk, and Skouras (2019) confirm this finding in more recent data.

mutual fund managers display stock picking skill during booms but market timing skill during recessions, best ideas are statistically significant during NBER expansions but not during NBER recessions. Appendix Table A.III confirms that variation in PE ratios also forecasts the conditional performance of best ideas.

In results not shown, we have added a demeaned trend to the portfolio attribution regressions within each of these subsamples, finding some evidence of the deterioration of the performance of best ideas over time within the expansion / market booms subsample. This deterioration is not necessarily surprising given the fact that the best mutual fund managers may have left to join the hedge fund industry. In addition, since the first draft of this paper, several investment banks have marketed investment products explicitly based on our best ideas approach.

We use forward-looking forecasts of CAPM idiosyncratic volatility when measuring the ex ante information ratio for stocks. The predictive regressions that generate those forecasts are described in Table A.V. The forecasted volatility should logically correspond to the idiosyncratic volatility the manager expects to experience over their expected holding period and should also correspond to the horizon of our trading strategies. As our holding period is three months, we forecast next quarter's idiosyncratic volatility. In particular, we estimate

$$\sigma_t = a + b * SIZE_{t-1} + c * BM_{t-1} + c * \sigma_{t-1} + d * \sigma_{t-5}$$

where t is measured in quarters. The estimated coefficients in conjunction with the time- $t - 1$ values generate our raw volatility forecast. We shrink this estimate to the cross-sectional average of σ_{t-1} using equal weights. In unreported results, our findings are robust to adding additional lags of the independent variables to the specification. Nevertheless, Appendix Table A.V shows that our results continue to hold if we simply use the most recent value of idiosyncratic volatility, σ_{t-1} , as our estimate of future idiosyncratic volatility.

Appendix Table A.VI documents that our findings are robust to using a twelve-month holding period rather than a three-month holding period. Appendix Table A.VI reports the way our key

findings change if we use other performance attribution models. One challenge here is that our methodology for identifying best ideas, which relies on the ratio of expected alpha to idiosyncratic volatility, tends to select high-idiosyncratic-volatility stocks. All else equal, such stocks are likely to be poor performers, as Ang, Hodrick, Xing, and Zhang (2006) document that stocks with relatively high idiosyncratic volatility have not only negative CAPM alpha but also underperform controlling for size, book-to-market, and momentum. However, Fama and French (2015) report that their five-factor model, unlike the three-factor version, prices the idiosyncratic volatility anomaly. Thus, our finding that best ideas tend to have higher Fama-French five-factor alphas compared to their FF three-factor alphas is unsurprising. Indeed, Jordan and Riley (2015) argue that failure to account for these factors can lead to substantial mismeasurement of mutual fund skill.

E Hedge Funds

We confirm the usefulness of our approach through an analysis based on hedge fund portfolios. These tests are independent in the sense that we obtained this data and performed the analysis long after the first draft of this paper was written, thus they constitute an out-of-sample test on our original findings. To identify hedge funds and their holdings, we use the identification approach of Agrawal, Jiang, Tang, and Yang (2012) who manually classify all 13-filing institutions based on a wide range of sources including institutional websites, SEC filings, industry publications, and broader news sources. Since 13F data is for the institution rather than the fund, we study only relatively “pure-play” hedge funds and thus exclude banks and mutual fund companies. Our final sample consists of 1,662 unique hedge funds.

Since hedge funds arguably care less about benchmarks and tracking error than mutual funds, we modify our methodology to also consider a third way of identifying best ideas that ignores these aspects,

$$\alpha_{ft}^{conviction} = \lambda_{ft}$$

$$best_{ft}^{conviction\ alpha} = \arg \max_i \alpha_{ift}^{conviction}$$

Table V documents that the equity best ideas of hedge funds strongly outperform for all of the methods we study. The first three rows of Table V Panel A benchmark the best ideas of active fund managers against the five- and six-factor models. Our estimates of five-factor alpha are 31 basis points (t -statistic of 3.21), 31 basis points (t -statistic of 3.11), and 36 basis points for the market, portfolio, and conviction measures respectively. These estimates barely change when we add momentum to the factor model: Our six-factor estimates are 28 basis points (t -statistic of 2.83), 28 basis points (t -statistic of 2.78), and 32 basis points (t -statistic of 3.17) for the market, portfolio, and conviction measures respectively.

The fourth through sixth rows of Table V Panel A show our conclusions continue to hold when using either the DGTW approach or best-minus-rest spread bets. The six-factor alphas on the DGTW-adjusted returns are 20 basis points (t -stat of 2.46), 25 basis points (t -stat of 2.92), and 20 basis points (t -stat of 2.54) respectively. We find higher estimates than these for the best-minus-rest portfolios. Six-factor alphas range from 31 basis points (t -stat of 3.11) to 25 basis points (t -stat of 2.41).

As we did for the mutual fund analysis, we report the performance of a portfolio that holds the entire portfolios of the managers in our sample, in proportion to AUM invested, as a benchmark. We find that the six-factor alpha of a portfolio of “All Ideas” is only 1 basis point and statistically insignificant.

As with our mutual fund analysis, this finding is robust as strategies which hold the top three or top five best ideas strongly outperform as well. Appendix Table A.XI documents that the six-factor alphas range from 11 basis points to 25 basis points across our six measures and two attribution models. The corresponding t -statistics are consistently above 2 for the top-three best

ideas class of strategies. For the top-five best ideas strategies, t -stats are always above 1.8 and in some cases are as high as 3.38.

Finally, in Table V Panel C, we document that the performance of best ideas varies with fund size. Best ideas strongly outperform in small hedge funds, defined as those funds with AUM in equities of 100 million or less. Specifically, the six-factor alpha based on our market-tilt-based information ratio measure is 58 basis points per month with a t -statistic of 2.40. Estimates using our other measures are broadly consistent. The point estimates of giant hedge funds, defined as funds with AUM in equities of 10 billion or more, are negative and large in absolute value, though generally not statistically significant at normal levels of significance. The difference between small and big funds best ideas six-factor alpha for our primary measure is more than 1.26% per month with an associated t -statistic of 2.58.

Appendix Table A.XII documents that the same characteristics that spread performance in best ideas in mutual funds also spread performance in hedge funds. Though the statistical significance of individual point estimates are generally weaker than in the mutual fund analysis, the six-factor alpha of a strategy based on our orthogonal composite signal, $COMP - O$ delivers a spread of 91 basis points a month (t -statistic of 3.20).

Our analysis has focused on the top 25% best ideas across the universe of active hedge fund managers. The Internet Appendix shows (Panels A through F of Table A.IX) that our findings concerning the performance of hedge fund best ideas generally hold as we vary this threshold from 50% to the top 12.5% of active tilts. Similarly, Internet Appendix Table A.X documents that the performance of the All Ideas portfolio is poor regardless of the active tilt threshold we implement.

VI Discussion and Implications

Modern portfolio theory (MPT) makes clear normative statements about optimal investing by managers on behalf of their clients. The core principle is straightforward: since linear combinations of mean-variance efficient portfolios are themselves MVE, managers should hold MVE portfolios on

behalf of their clients. But it was William Sharpe himself (1981) who noted that in the presence of constraints, delegating security selection to managers unaware of the client's other portfolio holdings is likely to lead to suboptimal outcomes. We extend these sorts of arguments to show that the problem can be quite severe when managers have investment opportunities of widely varying quality; i.e. if their "best ideas" really are significantly better than the remainder of the opportunities they perceive, as our empirical results suggest they are. Below we employ simple examples to illustrate the magnitude of this effect. In section VI.A we use a one-manager, two-idea example so concise one can do the calculations in one's head. In Section VI.B we make the situation more realistic by allowing a large number of managers, each of whom has one set of good ideas they can mix with the market in proportions of their choosing. Section VI.C briefly discusses properties of a more complete model while Section VI.D summarizes the resulting implications and Section VI.E touches upon additional considerations.

A A simple case with the market and two alpha opportunities

Suppose an endowment fund with mean-variance preferences has three possible investments: M (the global market portfolio), and X and Y (the two ideas for trades that a skilled manager possesses).

Let the riskless rate be zero and the expected returns on the assets be: $E[R_M] = 7\%$, $E[R_X] = 2\%$, and $E[R_Y] = 1\%$. Further suppose all three assets are uncorrelated and each has the same volatility. Assume the manager charges no fees. To fix ideas, imagine that the bets are purchases of catastrophe bonds: X , a bond that pays 3% in the 99% likely case that Florida hurricane losses fall below some cutoff and -100% otherwise, and Y , a similar bond that pays 2% on the 99% chance of below-threshold Japanese windstorm losses and -100% otherwise.

Unconstrained optimization delivers the result 70% in M , 20% in X , and 10% in Y , the portfolio that maximizes Sharpe ratio. The problem is separable: if we optimize the active manager's portfolio, we'll find that $2/3 X$ and $1/3 Y$ is optimal. At Stage 2, we can then optimize between

the market and the manager to get 70% and 30%, bringing us back to 70%, 20%, and 10% in M , X , and Y respectively. Everything is as expected, and the manager has not hurt their investor by maximizing Sharpe ratio in their two-asset sub-portfolio.

But, suppose the endowment decides in advance that it will not allocate more than 10% to the manager. Now in many cases, the best we can do in terms of Sharpe ratio in the absence of short-selling is if the manager puts 100% in the better bet X and zero in Y . In fact, if the managers can sell short, Sharpe ratio may often be further increased if they short Y to fund greater investment in X . Once we put in place the extremely realistic constraint that an endowment fund will cap the allocation to any given manager, then the manager is hurting the endowment's expected utility if they select the Sharpe-ratio-maximizing (SRM) portfolio of their ideas rather than concentrating on their very best idea(s).

Figure 5 shows the Sharpe ratios obtained at different allocations to the ideas X and Y . Each line on the graph shows the results for a different constraint on the total fraction of assets that are managed (i.e. invested in either X or Y). We indicate on each line with a star the amount in the best idea consistent with a myopic allocation by the active manager that only maximized their own-portfolio Sharpe ratio. If 30% of the investor's portfolio is allocated to the active manager, the global optimum is also the unconstrained choice (i.e., the highest Sharpe ratio occurs at 20%, which is $2/3$ of 30%). When we constrain the managed assets to either 10% or 20% of the portfolio, the maximum Sharpe ratio is reduced of course. Less obviously, when managed assets are constrained, the fraction of managed assets that should be held in the best idea X grows from $2/3$ in the unconstrained case to $4/5$ if managed holdings are capped at 20% of the portfolio (since the maximum Sharpe ratio is obtained at a 16% investment in the best idea). Finally, in the case where the fixed allocation to the active manager is only 10%, the optimal investment in the best idea becomes $11/10$, implying a short position of $-1/10$ in the second-best idea.

In summary, Figure 5 demonstrates that constraining the allocation to a manager should simultaneously also incentivize the client to push the fund manager to allocate more to best ideas. Otherwise, if managers act myopically by maximizing only their sub-portfolio's Sharpe ratio, the

overall Sharpe ratio may be reduced. In the example above, the magnitude of the reduction in Sharpe ratio is modest. In order for the true impact of the effect to be appreciated, one needs to consider the more realistic situation where the investor allocates to multiple managers, which we do next.

B A more realistic setting

Suppose that the assumptions underlying the CAPM hold, except that each manager has identified a single unit-beta investment opportunity X that has positive CAPM alpha. We assume that there are N managers, each of whom has one best idea so that each manager's portfolio consists of a combination of the best idea and the market portfolio. Note that the best idea could be thought of as an immutable basket of the manager's good ideas. For simplicity, we assume that each manager's idea has the same expected return, volatility, and beta and that the unsystematic components of managers' best ideas are uncorrelated. In Figure 6, we display the Sharpe ratios for such portfolios based on the following set of assumptions. Suppose that each investment X has 4% annual alpha and that the market premium is 6%; let the market's annual volatility be 15% and X 's be 40% (with the assumption of unit beta, every X must have a correlation of .375 with M , where M again represents the market portfolio). We continue to assume that the risk-free rate is zero.

The optimal risky portfolio for an investor to hold will be a mix of the X s and M , with each X having equal weight. The weights that are optimal are the weights that maximize the resulting portfolio's Sharpe ratio. If each individual manager maximizes their Sharpe ratio, the result will be that each manager will have 89% in the market and 11% in their best idea. And if the investor has access to only a single manager, the manager's choice will be the optimal choice for the investor as well. But as Figure 6 shows, the conclusion changes dramatically as the number of managers grows. For example, if the investor is allocating among five equally-skilled managers, the resulting portfolio will be optimized if each manager allocates approximately 47% to their best idea. If the

investor has access to fifty equally-skilled managers, the optimum is found when managers put 468% in their best idea (and -368% in the market).

The top line in Figure 6 shows the Sharpe ratio that would result if managers followed this optimal policy. The lowest line shows the Sharpe ratio the investor will obtain if each manager instead mean-variance optimizes their own portfolio. The middle line gives the resulting Sharpe ratios if managers choose the portfolio that is best for the investor but with the constraint that they cannot sell the market short.

Differences in Sharpe ratios are substantial. For fifty managers, manager-level optimization leads to a Sharpe ratio of 0.4 while the optimum optimum delivers 0.8, and the best-case scenario with short selling constraints provides 0.6. Moreover, optimal weights in the manager's best ideas are dramatically larger than what results from myopically maximizing manager-level Sharpe ratio.

In general, it seems likely that borrowing, lending, shorting, and maximum-investment constraints will create a situation where the investor's optimum requires the manager to choose a weight in X far greater than the SRM weight. Such a scenario would appear to be the case in typical real-world situations. A manager has a small number of good investment ideas. Modern portfolio theory says that any portfolio of stocks that maximizes CAPM information ratio is equally good for investors. But in truth, if the manager offers a portfolio with small weights in the good ideas and a very large weight in the market [or a near-market portfolio of zero- (or near-zero-) alpha stocks], the results for investors will be entirely unsatisfactory. The small allocation that investors make to any given manager, combined with the small weight such a manager places in the good ideas, mean that the manager adds very little value.

Suppose managers have optimized their Sharpe ratios and the investor wishes to obtain the constrained optimum. In a world where shorting the market was costless and common, an investor could take 100 dollars of capital and, instead of giving two dollars to each manager, could short the market to the tune of 800 dollars, giving 18 to each manager. Then, if each manager maximized Sharpe ratio and put 11% into their best idea, the investor would have about two dollars in each best idea and would approximately match the allocation the investor would have had if they had

given two dollars to each manager and each manager had put 100% of this capital in their best idea. In reality, it would be shocking to see an endowment fund pursue such an extreme market-shortening strategy. But this example shows that investors can alleviate the costs of constraints on allocations to managers if they can short and lever at the portfolio level.

So: MPT says all X - M combinations are equally good because investors can go long or short the market to return to the optimum. A natural choice for managers would be the SRM portfolio (11/89 in our example). But we see that the more realistic constrained case suggests that managers can serve their clients better by putting a much greater weight in X than the SRM weight – e.g. 100% instead of 11%. And yet as we see in Figure 2, over weights of best ideas by actual managers are smaller than 11%. Indeed overweights of that magnitude are rare. Of course, the 11% figure came from our simple example; perhaps managers view their best ideas as having far less than 4% alpha. But this seems unlikely, since we find actual outperformance of this order of magnitude despite our very poor proxy for best ideas (plus which, in the real world, managers are probably *overly* optimistic about the expectations for their best ideas, not *insufficiently* optimistic).

C Multiple managers, each with multiple positive-alpha ideas

Of course, in the real world, investors face a problem that combines the complexities of each of the two examples above – multiple managers, each with multiple ideas. Modelling this general case is beyond the scope of this paper. But we observe that if a manager wishes to serve a client optimally by dividing their allocation among their positive-alpha ideas, they will need to know what other managers are in the investor’s portfolio and what the quality of *those* managers’ best ideas is. Moreover, that answer will differ for each of the manager’s existing investors, not to mention potential investors the manager would like to court. In the next section, we discuss reasons managers might choose to employ mean-variance optimization of their holdings even when this is suboptimal for the underlying investors. It’s worth considering that one additional possible reason is that managers, faced with a problem where they lack sufficient information – and where

even if they possessed that information there would still be no way to please everyone – may simply throw up their hands in frustration and revert to simple rules of thumb.

D Why don't managers concentrate more in their best ideas?

Other conditions may differ from our simple example, but it appears probable that what we are observing is a decision by managers to diversify as much or more than the SRM portfolio despite the argument above that their clients would be best served by them diversifying far less than SRM. We identify four reasons managers may overdiversify.

1. **Regulatory/legal.** A number of regulations make it impossible or at least risky for many investment funds to be highly concentrated. Specific regulations bar overconcentration; additionally vague standards such as the “Prudent man” rule make it more attractive for funds to be better diversified from a regulatory perspective. Managers may well feel that a concentrated portfolio that performs poorly is likely to lead to investor litigation against the manager. Anecdotally, discussions with institutional fund-pickers make their preference for individual funds with low idiosyncratic risk clear. Some attribute the effect to a lack of understanding of portfolio theory by the selectors. Others argue that the selector's superior (whether inside or outside the organization) will tend to zero in on the worst performing funds, regardless of portfolio performance. Whatever the cause, we have little doubt that most managers feel pressure to be diversified.
2. **Price impact, liquidity and asset-gathering.** Berk and Green (2004) outline a model in which managers attempt to maximize profits by maximizing assets under management. In their model, as in ours, managers mix their positive-alpha ideas with a weighting in the market portfolio. The motivation in their model for the market weight is that investing in an individual stock will affect the stock's price, each purchase pushing it toward fair value.¹³

¹³See empirical work by Berk and van Binsbergen (2015), Pastor, Stambaugh, and Taylor (2015), and Zhu (2018) on empirical findings related to scale and skill in the fund management industry and Pastor, Stambaugh,

Thus, there is a maximum number of dollars of alpha that the manager can extract from a given idea. In the Berk and Green model managers collect fees as a fixed percentage of assets under management, and investors react to performance, so that in equilibrium each manager will raise assets until the fees are equal to the alpha that can be extracted from their good ideas. This choice leaves the investors with zero after-fee alpha. Clearly in the world of Berk and Green, and in the real world of mutual funds, managers with one great idea would be foolish to invest their entire fund in that idea, for this would make it impossible for them to capture a very high fraction of the idea's alpha in their fees. In other words, while investors benefit from concentration as noted above, managers under most commonly-used fee structures are better off with a more diversified portfolio. The distribution of bargaining power between managers and investors may therefore be a key determinant of diversification levels in funds.

3. **Manager risk aversion.** While the investor is diversified beyond the manager's portfolio, the manager himself is not. The portfolio's performance is likely the central determinant of the manager's wealth, and as such we should expect them to be risk averse over fund performance. A heavy bet on one or a small number of positions can, in the presence of bad luck, cause the manager to lose their business or their job (and perhaps much of their savings as well, if they are heavily invested in their own fund, as is common practice). If manager talent were fully observable this would not be the case – for a skilled manager, the poor performance would be correctly attributed to luck, and no penalty would be exacted. But when ability is being estimated by investors based on performance, risk-averse managers will have incentive to overdiversify.

4. **Investor irrationality.** There is ample reason to believe that many investors – even sophisticated institutional investors – do not fully appreciate portfolio theory and therefore

and Taylor (2017) for theoretical and empirical findings related to fund turnover in the presence of time-varying profit opportunities.

tend to judge individual investments on their expected Sharpe ratio rather than on what those investments are expected to contribute to the Sharpe ratio of their portfolio.¹⁴ For example, Morningstar's well-known star rating system is based on a risk-return trade-off that is highly correlated with Sharpe ratio. It is very difficult for a highly concentrated fund to get a top rating even if average returns are very high, as the star methodology heavily penalizes idiosyncratic risk. Since a large majority of all flows to mutual funds are to four- and five-star funds, concentrated funds would appear to be at a significant disadvantage in fund-raising.¹⁵ Other evidence of this bias includes the prominence of fund-level Sharpe ratios in the marketing materials of funds, as well as maximum drawdown and other idiosyncratic measures.

Both theory and evidence suggest that investors would benefit from managers holding more concentrated portfolios.¹⁶ As institutional investors such as endowment funds are certainly well-diversified, a move by fund managers to hold more concentrated funds would likely have little effect on the efficiency of (end) institutional investors' portfolios. In contrast, offering more concentrated funds to retail investors might possibly result in less-diversified portfolios for those investors, given their lower investment scale and sophistication. However, given the well-known MBA textbook result that most of the gains from diversification are achieved with a small number of individual stocks, such lack-of-diversification costs seem limited. In particular, Fisher and Lorie (1970) show that 80% of the gains from diversification are achieved by holding only eight stocks and 95% of those gains are achieved with just 32 stocks.

With that logic in hand, Curcuru, Heaton, Lucas, and Moore (2010) document that stockholding households with undiversified portfolios are actually a relatively small percentage of the data, less than 14% in the 2001 Survey of Consumer Finances.¹⁷ Indeed, Benartzi and Thaler (2001)

¹⁴This behavior is consistent with the general notion of "narrow framing" proposed by Kahneman and Lovallo (1993), Rabin and Thaler (2001), and Barberis, Huang, and Thaler (2006).

¹⁵Del Guercio and Tkac (2008) show that Morningstar star rating is the strongest variable predicting mutual fund flows out of those they consider, subsuming alpha in their analysis.

¹⁶See work by Van Nieuwerburgh and Veldkamp (2010).

¹⁷On a related note, Polkovnichenko (2005) emphasizes the diversification benefits of the coexistence of both

show that investors tend to follow the “1/n” strategy and divide contributions across all of the funds they are offered in their defined contribution plan rather than just hold one fund. Thus, the move by fund managers to hold more concentrated portfolios would also tend to have little impact on the ultimate efficiency of most retail portfolios.

More importantly, for those retail investors aiming to prioritize diversification over alpha in their allocation to mutual funds, a key consequence of choosing a diversified active mutual fund is that active funds charge “alpha fees.” Thus, rather than simply choosing an expensive diversified active fund, a cheaper solution for a diversification-driven retail investor would be to allocate a portion of their portfolio to one or more of the proposed concentrated best-ideas versions of managers’ funds and then allocate the remaining percentage to a low-cost index fund. Of course, while we argue that the first aspect of that strategy would be a beneficial change to the investment opportunity set, the second aspect of that strategy has been easily available since the 1976 introduction of Vanguard’s first index fund.

Our view is that we fail to see managers focusing on their best ideas for a number of reasons. Most of these relate to benefits to the manager of holding a diversified portfolio. But if those were the only causes, we would be hearing an outcry from investors about overdiversification by managers, while in fact such complaints are rare. Thus, we speculate that investor irrationality (or at least bounded rationality) in the form of manager-level analytics and heuristics that are not truly appropriate in a portfolio context, play a major role in causing overdiversification.

E Additional considerations

While we argue that ideas other than the “best ideas” enter the manager’s portfolio primarily to “round it out” and thereby reduce volatility, this view is of course an oversimplification. After all, if diversification were their only purpose, we might see managers simply put a fraction of their portfolio into their best ideas and the remainder into an index ETF or similar. However, such a

individual stocks and mutual funds within stockholding households’ portfolios.

strategy might be too flagrant and anger investors who, after all, see themselves as paying fees on/for the entire portfolio. And it might not even be optimal; a set of best ideas concentrated in, say, healthcare and tech might be best diversified by a selection of industrials, energy firms, and the like. Beyond these issues, though, we fully appreciate that managers do their best to add value with every position in their portfolio. No doubt even the least-high-conviction stock in the fund was chosen with some degree of care. Our point is merely that the market is very tough to beat, while diversifying is relatively easy. Stocks selected in the hope of simultaneously adding alpha and reducing portfolio volatility, and selected with sincere effort but, time constraints on human endeavor being what they are, perhaps only a modicum thereof, will likely be far more successful at diversifying than at alpha-adding.

If best ideas perform so much better than others, a natural response by fund management firms might be to hire a large number of managers to deliver a greater number of high-alpha stocks for the portfolio. Why don't we see this? One answer is that we do. There is considerable variation among money managers in how many staff are employed per holding, or per dollar invested. Additionally there is variation in the quality of staff applied to the management of each dollar or asset. Some evidence suggests that applying more resources to the asset-management challenge does bear fruit. For example, Chaudhuri, Ivkovic, Pollet, and Trzcinka (2020) show that funds where PhDs play a key role outperform otherwise similar firms. This variation brings us to deep issues in the industrial organization of the money-management business. Financial professionals are expensive, talented ones even more so. Firms have to weigh the trade-off between paying for more talent and accepting a higher probability of poor performance. Because there is so much luck in which managers perform well in any given year, it is plausible that many firms will choose the lower-cost, lower-expected-performance strategy.

Additionally, there are psychological factors at play in this human resources challenge: skilled stock-pickers want to be in control, to have the portfolio manager title and the accompanying prestige. Consequently, it may be challenging to keep, say, several dozen top managers working together to generate several dozen best ideas for a single fund, even if the fund is so large it can

compensate them all quite well for their efforts – they cannot all be the PM! There have been a few famous cases of this sort of structure working well, at least for a period of time, e.g. in the glory years of Julian Robertson’s Tiger Management. According to firm lore, he had a group of skilled investors funneling their best ideas up to Robertson, who then selected the best of the best. It may be that holding such talent in subordinate positions requires a manager of legendary status. And even so it did not last forever, most of the key Tiger staff went on to manage their own funds.

Finally, we argue that with asset managers paid on assets under management, it is understandable that they would work to maximize the AUM that they can manage well enough to keep investors in the fund. This point leads to a natural question: Is paying a percentage of AUM the optimal compensation scheme for asset managers? Of course, we do see other contracts; indeed, fees based on both a fraction of assets and a fraction of profits (“2 and 20”) is one of the defining features of the hedge fund industry. And it has been often argued by hedge fund advocates that this structure provides incentives to hedge funds to focus less on asset-gathering and more on performance than traditional contracts tied only to AUM.

At the same time, it is common for hedge fund investors to claim that asset-gathering is still a central goal of hedge fund managers. A full exploration of optimal contract design in asset management is beyond the scope of this paper. Our view is that the fee structure is not the primary factor leading managers to overdiversify. Rather, as we discuss in section VI.D, not only do managers have strong incentives to favor diversified portfolios, investors may inappropriately view diversified portfolios as in their favor as well. Consequently, there may be little pressure to create new contracts that would push managers toward concentration. And of course floating above all this theorizing is the simple fact that in the United States there are severe regulatory restrictions on the types of contracts mutual funds can offer their customers. Almost all funds use a flat percentage of AUM because this is an entirely legal way to charge fees, and while there is some flexibility for mutual funds to be paid more directly for performance, these exceptions are rather narrow and difficult to implement, especially given that such approaches are viewed as

non-standard and therefore suspicious by many investors.

VII Conclusion

How efficient are stock prices? This question is perhaps the central issue in the study of investing. Many have interpreted the fact that skilled professionals fail to beat the market by a significant amount as very strong evidence for the efficiency of the stock market.

This paper asks a related simple question. What if each mutual fund manager had only to pick a few stocks, their best ideas? Could they outperform under those circumstances? We document strong evidence that they could, as the best ideas of active managers generate up to an order of magnitude more alpha than their portfolio as a whole, depending on the performance benchmark and the type of stocks and funds in question.

We argue that these results present powerful evidence that the typical mutual fund manager can indeed pick stocks. The poor overall performance of mutual fund managers in the past is not due to a lack of stock-picking ability, but rather to institutional factors that encourage them to overdiversify, i.e. pick stocks beyond their best alpha-generating ideas. We point out that these factors may include not only the desire to have a very large fund and therefore collect more fees [as detailed in Berk and Green (2004)] but also the desire by both managers and investors to minimize a fund's idiosyncratic volatility: Though of course managers are risk averse, it seems investors may judge funds irrationally by measures such as Sharpe ratio or Morningstar rating. Both of these measures penalize idiosyncratic volatility, a penalty whose benefits in a portfolio context are extremely questionable.

References

- Alexander, Gordon, Gjergji Cici, and Scott Gibson, 2007, Does Motivation Matter When Assessing Trade Performance? An Analysis of Mutual Funds, *Review of Financial Studies* 12 125–150.
- Ang, Andrew, Robert Hodrick, Yuhang Xing and Xiaoyan Zhang, 2006, The Cross-Section of Volatility and Expected Returns, *Journal of Finance* 51 259–299.
- Barberis, Nicholas, Ming Huang, and Richard H. Thaler, 2006, Individual Preferences, Monetary Gambles, and Stock Market Participation: A Case for Narrow Framing, *American Economic Review* 96 1069–1090.
- Benartzi, Shlomo and Richard H. Thaler, 2001, Naive Diversification Strategies in Defined Contribution Saving Plans, *American Economic Review* 91 79–98.
- Berk, Jonathan B. and Richard Green, 2004, Mutual fund flows and performance in rational markets, *Journal of Political Economy* 112 1269–1295.
- Berk, Jonathan and Jules van Binsbergen, 2015, Measuring Skill in the Mutual Fund Industry, *Journal of Financial Economics* 118 1–20.
- van Binsbergen, Jules, Michael W. Brandt, and Ralph S. Koijen, Optimal Decentralized Investment Management, *Journal of Finance* 63 1849–1895.
- Brandt, Michael W., Alon Brav, John R. Graham, and Alok Kumar, 2010, The Idiosyncratic Volatility Puzzle: Time Trend or Speculative Episode?, *Review of Financial Studies* 23 863–899.
- Campbell, John Y., Martin Lettau, Burton Makiel, and Yexiao Xu, 2001, Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk, *Journal of Finance* 56 1–43.
- Carhart, Mark, 1997, On Persistence in Mutual Fund Performance, *Journal of Finance* 52 57–82.
- Chaudhuri, Ranadeb, Zoran Ivkovic, Joshua Pollet, Charles Trzcinka, 2020, A Tangled Tale of Training and Talent: PhDs in Institutional Asset Management, *Management Science* 66 5623–5647.
- Chen, Hsiu-Lang, Narasimhan Jegadeesh, and Russ Wermers, 2000, The Value of Active Mutual Fund Management, *Journal of Financial and Quantitative Analysis* 35 343–368.
- Chen, Joseph, Harrison Hong, Ming Huang, and Jeffrey D. Kubik, 2004, Does Fund Size Erode Mutual Fund Performance? The Role of Liquidity and Organization, *American Economic Review* 94 1276–1302.
- Chevalier, Judith and Glenn Ellison, 1997, Risk Taking by Mutual Funds as a Response to Incentives, *Journal of Political Economy* 105 1167–1200.

- Cho, Thummim and Christopher Polk, 2021, Putting the Price in Asset Pricing, London School of Economics working paper.
- Cohen Randolph B., Paul A. Gompers, and Tuomo Vuolteenaho, 2002, Who underreacts to cash-flow news? Evidence from trading between individuals and institutions, *Journal of Financial Economics* 66 409–462.
- Cohen Randolph B., Christopher Polk, and Tuomo Vuolteenaho, 2009, The Price is (Almost) Right, *Journal of Finance* 64 2739–2782.
- Cremers, Martijn and Antti Petajisto, 2009, How Active is Your Fund Manager? A New Measure That Predicts Fund Performance, *Review of Financial Studies* 22 3329–3365.
- Cremers, Martijn, Antti Petajisto, and Eric Zitzewitz, 2012, Should Benchmark Indices Have Alpha? Revisiting Performance Evaluation, *Critical Finance Review* 2 1–48.
- Curcuro, Stephanie, John Heaton, Deborah Lucas, and Damien Moore, 2010, Heterogeneity and portfolio choice: Theory and evidence, *The Handbook of Financial Econometrics*, Vol 1, North Holland.
- Daniel, Kent, Mark Grinblatt, Sheridan Titman, and Russ Wermers, 1997, Measuring Mutual Fund Performance With Characteristic-Based Benchmarks, *Journal of Finance* 52 1035–1058.
- Del Guercio, Diane and Paula A. Tkac (2008), The determinants of the flow of funds of managed portfolios: Mutual funds versus pension funds, *Journal of Financial and Quantitative Analysis* 43 907–936.
- Edelen, R., 2002, Investor flows and the assessed performance of open-end fund managers, *Journal of Financial Economics* 53 439–466.
- Elton, Edwin J., and Martin J. Gruber, 2004, Optimum centralized portfolio construction with decentralized portfolio management, *Journal of Financial and Quantitative Analysis* 39 481–494.
- Fama, Eugene F. and Kenneth R. French, 1992. The cross-section of expected stock returns. *Journal of Finance* 47 427–465.
- Fama, Eugene F. and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33 3–56.
- Fama, Eugene F. and Kenneth R. French, 2015, A five-factor asset pricing model, *Journal of Financial Economics* 116 1–22.
- Fisher, Lawrence and James H. Lorie, 1970, Some Studies of Variability of Returns on Investments in Common Stocks, *Journal of Business* 43 99–134.
- Gompers, Paul and Andrew Metrick, 2001, Institutional Investors and Asset Prices, *Quarterly Journal of Economics* 116 229–259.

- Grossman, Sanford J. and Joseph E. Stiglitz, 1980, On the Impossibility of Informationally Efficient Markets, *American Economic Review* 70 393–408.
- Jensen, Michael C., 1968, The Performance of Mutual Funds in the Period 1945-1964, *Journal of Finance* 23 389–416.
- Jiang, Hao, Marno Verbeek, and Yu Wang, 2014, Information Content When Mutual Funds Deviate from Benchmarks, *Management Science* 60 1861–2109.
- Jordan, Bradford, and Timothy Riley, 2015, Volatility and mutual fund skill, *Journal of Financial Economics* 118 289–298.
- Kacperczyk, Marcin, Clemens Sialm, and Lu Zheng, 2005, On the Industry Concentration of Actively Managed Equity Mutual Funds, *Journal of Finance* 60 1983–2021.
- Kahneman, Daniel and Amos Tversky, 1979, Prospect Theory: Analysis of Decision Under Uncertainty, *Econometrica* 47 263–291.
- Kahneman, Daniel and Daniel Lovallo, 1993, Timid Choices and Bold Forecasts: A Cognitive Perspective on Risk Taking, *Management Science* 39 17–31.
- Lou, Dong, Christopher Polk, and Spyros Skouras, 2019, A Tug of War: Overnight vs. Intraday Expected Returns, *Journal of Financial Economics* 134 192–213.
- MacKinlay, A. Craig and Lubos Pastor, 2000, Asset Pricing Models: Implications for Expected Returns and Portfolio Selection, *Review of Financial Studies* 13 883–916.
- MacKinlay, A. Craig, 1995, Multifactor models do not explain deviations from the CAPM, *Journal of Financial Economics* 38 3–28.
- Markowitz, Harry, 1952, Portfolio Selection, *Journal of Finance* 7 77–91.
- Massa, Massimo, Jonathan Reuter, and Eric Zitzewitz, 2010, When should firms share credit with employees? Evidence from anonymously managed mutual funds, *Journal of Financial Economics* 95 400–424.
- van Nieuwerburgh, Stijn and Laura Veldkamp, 2010, Information Acquisition and Under-Diversification, *Review of Economic Studies* 77 779–805.
- Pastor, Lubos and Robert Stambaugh, 2003, Liquidity Risk and Expected Stock Returns, *Journal of Political Economy* 111 642–685.
- Pastor, Lubos, Robert Stambaugh, and Lucian Taylor, 2015, Scale and Skill in Active Management, *Journal of Financial Economics* 116 23–45.
- Pastor, Lubos, Robert Stambaugh, and Lucian Taylor, 2017, Do Funds Make More When They Trade More?, *Journal of Finance* 72 1483–1528.
- Polkovnichenko, Valery, 2005, Household portfolio diversification: A case for rank dependent preferences, *Review of Financial Studies* 18, 1467–1502.

- Pomorski, Lukasz, 2009, Acting on the Most Valuable Information: ‘Best Idea’ Trades of Mutual Fund Managers, University of Toronto working paper.
- Rabin, Matthew and Thaler, Richard H., 2001, Anomalies: Risk Aversion, *Journal of Economic Perspectives* 15 219–232.
- Rubinstein, Mark, 2001, Rational Markets? Yes or No: The Affirmative Case. *Financial Analysts Journal* 57 15–29.
- Sadka, Ronnie, 2006, Momentum and Post-earnings-announcement drift anomalies: The Role of Liquidity Risk, *Journal of Financial Economics* 80 309–349.
- Sharpe, William F., 1966, Mutual Fund Performance, *Journal of Business* 39 119–138.
- Sharpe, William F., 1981, Decentralized Investment Management, *Journal of Finance* 36 217–234.
- Shumway, Tyler, Maciej Szeffler, and Kathy Yuan, 2009, The Information Content of Revealed Beliefs in Portfolio Holdings, University of Michigan working paper.
- Sirri, Erik R. and Peter Tufano, 1998, Costly search and mutual fund flows, *Journal of Finance* 53 1589–1622.
- Wermers, Russ, 2000, Mutual fund performance: An empirical decomposition into stock picking talent, style, transaction costs, and expenses. *Journal of Finance* 55 1655–1695.
- Wermers, Russ, Tong Yao, and Jane Zhao, 2012, Forecasting stock returns through an efficient aggregation of mutual fund holdings, *Review of Financial Studies* 25 3490–3529.
- Zhu, Min, 2018, Informative Fund Size, Managerial Skill, and Investor Rationality, *Journal of Financial Economics* 130 114–134.

Table I: Summary Statistics

This table reports year-end summary statistics sampled every five years from 1983 to 2018 for all mutual fund portfolios detailed on Thompson that contain at least five stocks. Panel A contains information about all positions. Column 1 reports the number of all mutual funds while columns 2 to 8 focus on characteristics of only active equity mutual funds, those funds whose maximum position-level information ratio is in the top 25% of all corresponding information ratios at the time. Columns 5-8 report the median decile portfolio for the Fama and French (2014) characteristics for all stocks in those portfolios. Panel B details summary statistics for the best idea of the 25% most active mutual funds. Owner fraction is the average share of ownership held by a manager's best idea. AUM fraction is the weight that managers allocate to their best idea, averaged across funds. Panel C shows the same statistics as in Panel B, but for the top five best ideas within these portfolios.

PANEL A: All Positions								
Year	Number of Funds	Number of Active Funds	Avg Fund Size (\$b)	Stocks Per Fund	Size Decile	BM Decile	OP Decile	INV Decile
1983	471	270	0.2	65	3.9	4.7	5.5	5.9
1985	707	335	0.2	65	3.8	4.7	5.6	6.2
1990	1170	456	0.2	67	4.4	4.8	5.7	6.2
1995	2687	908	0.5	92	3.6	5.2	5.0	6.1
2000	8544	1819	0.8	115	4.3	4.9	4.7	6.0
2005	9365	1655	0.7	133	3.7	5.1	4.8	5.9
2010	7454	1593	1.0	148	4.0	5.3	5.0	5.9
2015	10597	1763	1.8	155	4.4	5.1	4.9	6.0
2018	9747	1658	2.3	165	4.5	5.0	4.9	6.1

PANEL B: Top Best Idea Among Top 25% Active Tilters							
	Unique	Owner	AUM	Size	BM	OP	INV
Year	Best Ideas	Fraction	Fraction	Decile	Decile	Decile	Decile
1983	60	2.6%	10%	7.0	3.5	6.4	7.1
1985	67	2.3%	11%	7.0	3.0	6.2	6.7
1990	98	1.8%	12%	7.0	2.8	6.6	7.4
1995	155	1.2%	12%	7.2	2.7	6.3	6.8
2000	223	0.8%	11%	8.1	2.3	5.3	7.1
2005	183	0.7%	9%	7.5	3.4	5.4	6.2
2010	170	0.7%	11%	8.0	3.8	6.1	6.5
2015	155	0.9%	10%	7.8	3.4	5.6	6.3
2018	168	0.6%	10%	8.0	3.8	6.1	6.7

PANEL C: Top Five Best Ideas Among Top 25% Active Tilters							
	Unique	Owner	AUM	Size	BM	OP	INV
Year	Best Ideas	Fraction	Fraction	Decile	Decile	Decile	Decile
1983	264	1.3%	23%	6.8	3.7	6.2	6.6
1985	292	1.2%	26%	7.0	3.5	6.3	6.4
1990	382	1.1%	28%	6.7	3.4	6.3	6.8
1995	566	1.0%	25%	6.8	3.6	5.9	6.4
2000	682	1.1%	26%	7.4	2.9	5.4	6.6
2005	645	0.8%	24%	6.8	3.8	5.4	6.3
2010	586	0.9%	27%	7.2	4.2	5.8	6.4
2015	605	1.1%	26%	7.2	3.8	5.9	6.1
2018	606	0.8%	28%	7.5	4.0	5.9	6.4

Table II: Performance of Best Ideas

This table reports coefficients from monthly six-factor regressions, $r_{p,t}^{excess} = \alpha_6 + bRMRF_t + sSMB_t + hHML_t + rRWM_t + cCMA_t + mUMD_t + \varepsilon_{p,t}$, where $r_{p,t}^{excess}$ is either $r_{p,t} - r_{f,t}$ the equal-weight excess return on the portfolio of the stocks that include the best idea of each active manager; $r_{p,t} - r_{DGTW,t}$ the equal-weight DGTW-adjusted return on the best idea portfolio; or $spread_{p,t}$, the return on an equal-weight long-short portfolio that is long a dollar in each manager's best idea and short a dollar in each manager's investment-weight portfolio of the rest of their ideas, where the rest is defined as positions outside each manager's top ten ideas. The best idea is determined within each fund as the stock with the maximum value of one of two possible information ratio estimates: 1) $IR_{ift}^{market} = \sigma_{it}(\lambda_{ift} - \lambda_{iMt})$ or 2) $IR_{ift}^{portfolio} = \sigma_{it}(\lambda_{ift} - \lambda_{iVt})$, where λ_{ift} is manager f 's portfolio weight in stock i , λ_{iMt} is the weight of stock i in the market portfolio, λ_{iVt} is the value-weight position of stock i in manager f 's portfolio, and σ_{it} is the most-recent forecast of a stock's CAPM-idiosyncratic volatility. Explanatory variables are from Ken French's website. We also report intercept estimates, α_5 , when UMD is excluded from the regression. As a benchmark, Panel B shows the results when we instead form investment-weight portfolios based on all of the ideas of managers instead of just their best ideas. In both panels, we restrict the analysis to those managers whose maximum position-level information ratio is in the top 25% of all corresponding information ratios at the time. We report t -statistics in parentheses and denote regression intercepts (Mean, α_5 , and α_6) that are statistically significant at the 5% level in bold font. The sample period for the dependent variables is January 1983-March 2019.

PANEL A: Performance of Best Ideas										
	Mean	α_5	α_6	RMRF	SMB	HML	RMW	CMA	UMD	R ²
$r_{1,t} - r_{f,t}$	0.97% (3.19)	0.47% (4.09)	0.37% (3.45)	1.12 (41.87)	0.15 (4.02)	-0.26 (-5.10)	-0.35 (-6.99)	-0.23 (-3.11)	0.19 (8.11)	89%
$r_{2,t} - r_{f,t}$	1.01% (3.23)	0.43% (3.89)	0.35% (3.29)	1.13 (42.65)	0.43 (11.18)	-0.19 (-3.82)	-0.31 (-6.33)	-0.19 (-2.68)	0.16 (6.95)	90%
$r_{1,t} - r_{DGTW,t}$	0.21% (2.05)	0.33% (3.56)	0.26% (2.91)	0.07 (3.31)	0.01 (0.36)	-0.13 (-3.08)	-0.24 (-6.06)	-0.13 (-2.13)	0.15 (7.44)	38%
$r_{2,t} - r_{DGTW,t}$	0.25% (2.43)	0.35% (3.88)	0.29% (3.29)	0.08 (3.52)	0.07 (2.32)	-0.13 (-3.10)	-0.24 (-5.84)	-0.12 (-1.95)	0.13 (6.72)	39%
$spread_{1,t}$	0.21% (1.79)	0.39% (3.52)	0.26% (2.71)	0.05 (2.10)	-0.05 (-1.34)	-0.11 (-2.35)	-0.32 (-7.31)	-0.17 (-2.54)	0.25 (11.54)	43%
$spread_{2,t}$	0.24% (1.95)	0.34% (3.39)	0.23% (2.60)	0.07 (3.13)	0.25 (7.55)	-0.07 (-1.56)	-0.30 (-7.21)	-0.14 (-2.23)	0.21 (10.34)	52%
PANEL B: Benchmark - Performance of All Ideas										
	Mean	α_5	α_6	RMRF	SMB	HML	RMW	CMA	UMD	R ²
$r_{all,t} - r_{f,t}$	0.65% (3.09)	0.04% (1.12)	0.06% (1.55)	0.93 (102.42)	0.15 (11.26)	-0.11 (-6.40)	-0.02 (-1.13)	-0.06 (-2.23)	-0.03 (-3.48)	97%

Table III: Performance of Top Three and Top Five Best Ideas

This table reports coefficients from monthly six-factor regressions, $r_{p,t}^{excess} = a_6 + bRMRF_t + sSMB_t + hHML_t + rRWM_t + cCMA_t + mUMD_t + \varepsilon_{p,t}$, where $r_{p,t}^{excess}$ is either $r_{p,t} - r_{f,t}$ the equal-weight excess return on the portfolio of the stocks that include the top-three (Panel A) or top-five (Panel B) best ideas of each active manager; $r_{p,t} - r_{DGTW,t}$ the equal-weight DGTW-adjusted return on that best ideas portfolio; or $spread_{p,t}$, the return on an equal-weight long-short portfolio that is long a dollar in a portfolio of the top-three (Panel A) or top-five (Panel B) best ideas and short a dollar in each manager's investment-weight portfolio of the rest of their ideas, where the rest is defined as positions outside each manager's top ten ideas. Best ideas are determined within each fund as the stock with the highest values of one of two possible information ratio estimates: 1) $IR_{ift}^{market} = \sigma_{it}(\lambda_{ift} - \lambda_{iMt})$ or 2) $IR_{ift}^{portfolio} = \sigma_{it}(\lambda_{ift} - \lambda_{iVt})$, where λ_{ift} is manager f 's portfolio weight in stock i , λ_{iMt} is the weight of stock i in the market portfolio, λ_{iVt} is the value-weight position of stock i in manager f 's portfolio, and σ_{it} is the most-recent forecast of a stock's CAPM-idiosyncratic volatility. Explanatory variables are from Ken French's website. We also report intercept estimates, α_5 , when UMD is excluded from the regression. We restrict the analysis to those managers whose maximum alpha is in the top 25% of all maximum alphas at the time. We report t -statistics in parentheses and denote regression intercepts (Mean, α_5 , and α_6) that are statistically significant at the 5% level in bold font. The sample period for the dependent variables is January 1983-March 2019.

PANEL A: Best Three Ideas										
	Mean	α_5	α_6	RMRF	SMB	HML	RMW	CMA	UMD	R ²
$r_{1,t} - r_{f,t}$	0.89% (3.30)	0.33% (4.18)	0.26% (3.45)	1.08 (58.28)	0.16 (5.84)	-0.22 (-6.17)	-0.26 (-7.42)	-0.20 (-3.96)	0.14 (8.46)	93%
$r_{2,t} - r_{f,t}$	0.92% (3.33)	0.29% (3.65)	0.23% (3.03)	1.10 (57.79)	0.39 (14.29)	-0.16 (-4.57)	-0.20 (-5.64)	-0.14 (-2.72)	0.10 (6.04)	93%
$r_{1,t} - r_{DGTW,t}$	0.14% (1.92)	0.22% (3.65)	0.17% (2.91)	0.05 (3.62)	0.03 (1.46)	-0.10 (-3.62)	-0.18 (-6.49)	-0.13 (-3.18)	0.10 (7.89)	43%
$r_{2,t} - r_{DGTW,t}$	0.17% (2.36)	0.23% (3.77)	0.18% (3.13)	0.06 (4.16)	0.07 (3.29)	-0.11 (-3.95)	-0.15 (-5.51)	-0.09 (-2.24)	0.08 (6.41)	42%
$spread_{1,t}$	0.12% (1.43)	0.28% (3.47)	0.16% (2.43)	0.02 (1.37)	-0.04 (-1.92)	-0.08 (-2.46)	-0.26 (-8.54)	-0.16 (-3.63)	0.22 (14.91)	51%
$spread_{2,t}$	0.15% (1.72)	0.22% (3.21)	0.13% (2.23)	0.05 (3.17)	0.24 (11.11)	-0.04 (-1.55)	-0.21 (-7.38)	-0.10 (-2.42)	0.16 (11.78)	58%
PANEL B: Best Five Ideas										
	Mean	α_5	α_6	RMRF	SMB	HML	RMW	CMA	UMD	R ²
$r_{1,t} - r_{f,t}$	0.86% (3.31)	0.28% (4.16)	0.22% (3.49)	1.07 (66.74)	0.16 (6.77)	-0.21 (-6.90)	-0.20 (-6.68)	-0.17 (-3.85)	0.10 (7.12)	95%
$r_{2,t} - r_{f,t}$	0.90% (3.38)	0.25% (3.67)	0.21% (3.16)	1.09 (64.31)	0.37 (15.31)	-0.15 (-4.66)	-0.16 (-4.96)	-0.12 (-2.52)	0.07 (4.74)	94%
$r_{1,t} - r_{DGTW,t}$	0.12% (1.92)	0.18% (3.61)	0.14% (2.93)	0.05 (3.96)	0.03 (1.77)	-0.10 (-4.45)	-0.13 (-5.78)	-0.11 (-3.18)	0.08 (6.98)	44%
$r_{2,t} - r_{DGTW,t}$	0.16% (2.60)	0.20% (3.94)	0.17% (3.37)	0.06 (4.58)	0.06 (3.50)	-0.10 (-4.12)	-0.12 (-5.14)	-0.07 (-2.09)	0.06 (5.51)	41%
$spread_{1,t}$	0.10% (1.38)	0.24% (3.52)	0.14% (2.48)	0.01 (0.80)	-0.06 (-2.81)	-0.08 (-2.94)	-0.22 (-8.35)	-0.14 (-3.67)	0.19 (15.46)	52%
$spread_{2,t}$	0.14% (1.94)	0.19% (3.29)	0.12% (2.34)	0.04 (3.30)	0.24 (13.12)	-0.03 (-1.14)	-0.17 (-6.81)	-0.07 (-1.96)	0.13 (11.33)	59%

Table IV: Performance of Best Ideas Based on Stock and Fund Characteristics

This table reports the intercept from monthly six-factor regressions, $r_{p,t}^{excess} = a_6 + bRMRF_t + sSMB_t + hHML_t + rRWM_t + cCMA_t + mUMD_t + \varepsilon_{p,t}$, where $r_{p,t}^{excess}$ is either $r_{p,t} - r_{f,t}$ the equal-weight excess return on the portfolio of the stocks that represent the best idea of each active manager or the equal-weight DGTW-adjusted return on that best ideas portfolio. The best idea, p , is determined within each fund as the stock with the maximum value of the following information ratio estimate: $IR = \sigma_{it}(\lambda_{ift} - \lambda_{iMt})$, where λ_{ift} is manager f's portfolio weight in stock i, λ_{iMt} is the weight of stock i in the market portfolio, and σ_{it} is the most-recent forecast of a stock's CAPM-idiosyncratic volatility. Explanatory variables are from Ken French's website. We restrict the analysis to those managers whose maximum position-level information ratio is in the top 25% of all corresponding information ratios at the time. We report decompositions of the resulting a_6 estimates based on the following stock and fund characteristics – illiquidity (ILIQ) is the bid-ask spread; growth (GROWTH) is the market-to-book ratio; momentum (MOM) is the past 11-month return skipping the most recent month; alpha (ALPHA) is the fund's Carhart alpha, using rolling windows of 36 months, composite (COMP) is a combination of the percentile rankings of the previous characteristics, and orthogonal composite (COMP-O) is the combination of the percentile rankings of those characteristics, after first being orthogonalized to each other. We report t -statistics in parentheses and denote regression intercepts (a_6) that are statistically significant at the 5% level in bold font. Differences across the High and Low portfolios are shaded. The sample period for the dependent variables is January 1983-March 2019.

Performance of Best Ideas Conditional on Ex Ante Stock/Fund Characteristics						
	ILIQ	GROWTH	MOM	ALPHA	COMP	COMP-O
$r_{LOW,t} - r_{f,t}$	0.02% (0.24)	-0.04% (-0.40)	-0.04% (-0.41)	0.13% (0.90)	-0.05% (-0.66)	-0.08% (-0.68)
$r_{MED,t} - r_{f,t}$	0.30% (2.20)	0.13% (1.13)	0.27% (2.37)	0.37% (2.42)	0.12% (1.04)	0.23% (1.52)
$r_{HIGH,t} - r_{f,t}$	0.61% (2.89)	0.82% (4.05)	0.72% (3.37)	0.50% (3.17)	0.86% (3.90)	0.82% (4.07)
$r_{HIGH} - r_{LOW}$	0.59% (2.56)	0.86% (3.94)	0.76% (3.21)	0.38% (2.23)	0.91% (3.92)	0.90% (3.93)
$r_{LOW,t} - r_{DGTW,t}$	0.00% (0.05)	-0.05% (-0.72)	-0.06% (-0.73)	0.05% (0.39)	-0.05% (-0.78)	-0.03% (-0.23)
$r_{MED,t} - r_{DGTW,t}$	0.16% (1.39)	0.09% (0.86)	0.19% (2.00)	0.21% (1.57)	0.02% (0.26)	0.10% (0.68)
$r_{HIGH,t} - r_{DGTW,t}$	0.48% (2.64)	0.59% (3.48)	0.52% (2.82)	0.37% (2.65)	0.65% (3.50)	0.55% (2.83)
$r_{HIGHDGTW} - r_{LOWDGTW}$	0.47% (2.50)	0.64% (3.53)	0.58% (2.87)	0.32% (2.13)	0.70% (3.58)	0.58% (2.62)

Table V: Performance of Best Ideas in Hedge Funds

This table reports coefficients from monthly regressions of $r_{p,t}^{excess} = \alpha_6 + bRMRF_t + sSMB_t + hHML_t + rRWM_t + cCMA_t + mUMD_t + \varepsilon_{p,t}$, where $r_{p,t}^{excess}$ is either $r_{p,t} - r_{f,t}$ the equal-weight excess return on the portfolio of the stocks that include the best idea of each active manager; $r_{p,t} - r_{DGTW,t}$ the equal-weight DGTW-adjusted return on the best ideas portfolio; or $spread_{p,t}$, the return on an equal-weight long-short portfolio, long a dollar in each manager's best idea and short a dollar in each manager's investment-weight portfolio of the rest of their ideas, where the rest is defined as positions outside each manager's top ten ideas. The best idea is determined within each fund as the stock with the maximum value of one of three possible measures: 1) $IR_{ift}^{market} = \sigma_{it}(\lambda_{ift} - \lambda_{iMt})$, 2) $IR_{ift}^{portfolio} = \sigma_{it}(\lambda_{ift} - \lambda_{iVt})$, or 3) $\alpha_{ift}^{conviction} = \lambda_{ift}$ where λ_{ift} is manager f 's portfolio weight in stock i , λ_{iMt} is the weight of stock i in the market portfolio, λ_{iVt} is the value weight of stock i in manager f 's portfolio, and σ_{it} is the most-recent estimate of a stock's CAPM-idiosyncratic volatility. Explanatory variables are from Ken French's website. We also report intercept estimates, α_5 , when UMD is excluded from the regression. As a benchmark, Panel B shows the results when we instead form portfolios based on all of the ideas of managers instead of just their best ideas. Panel C decomposes the results in Panel A as a function of fund size. In all three panels, we restrict the analysis to those managers whose maximum position-level information ratio is in the top 25% of all corresponding information ratios at the time. We report t -statistics in parentheses and denote regression intercepts (Mean, α_5 , and α_6) that are statistically significant at the 5% level in bold font. The sample period for the dependent variables is January 1983-March 2019.

PANEL A: Performance of Best Ideas										
	Mean	α_5	α_6	RMRF	SMB	HML	RMW	CMA	UMD	R ²
$r_{1,t} - r_{f,t}$	0.96% (3.77)	0.31% (3.21)	0.28% (2.83)	1.02 (41.89)	0.34 (9.84)	0.08 (1.71)	-0.19 (-4.21)	-0.11 (-1.65)	0.07 (3.30)	87%
$r_{2,t} - r_{f,t}$	0.96% (3.71)	0.31% (3.11)	0.28% (2.78)	1.01 (40.24)	0.47 (12.93)	0.04 (0.84)	-0.22 (-4.53)	-0.02 (-0.27)	0.06 (2.81)	87%
$r_{3,t} - r_{f,t}$	0.97% (3.91)	0.36% (3.51)	0.32% (3.17)	0.98 (38.17)	0.35 (9.52)	0.07 (1.37)	-0.16 (-3.40)	-0.15 (-2.16)	0.07 (2.85)	85%
$r_{1,t} - r_{DGTW,t}$	0.15% (1.89)	0.22% (2.78)	0.20% (2.46)	0.01 (0.57)	0.02 (0.77)	0.00 (0.03)	-0.17 (-4.53)	-0.07 (-1.27)	0.05 (2.63)	11%
$r_{2,t} - r_{DGTW,t}$	0.20% (2.30)	0.28% (3.23)	0.25% (2.92)	0.00 (-0.01)	0.05 (1.71)	-0.04 (-0.96)	-0.21 (-5.04)	0.01 (0.21)	0.05 (2.59)	13%
$r_{3,t} - r_{DGTW,t}$	0.17% (2.28)	0.22% (2.82)	0.20% (2.54)	-0.01 (-0.69)	0.04 (1.42)	0.01 (0.32)	-0.07 (-1.81)	-0.08 (-1.50)	0.04 (2.28)	4%
$spread_{1,t}$	0.27% (2.65)	0.38% (3.69)	0.31% (3.11)	0.02 (0.64)	-0.01 (-0.21)	0.04 (0.76)	-0.23 (-4.82)	-0.16 (-2.36)	0.13 (5.75)	15%
$spread_{2,t}$	0.24% (2.32)	0.31% (3.06)	0.26% (2.58)	0.00 (0.18)	0.14 (4.03)	-0.01 (-0.30)	-0.20 (-4.16)	-0.03 (-0.39)	0.10 (4.40)	17%
$spread_{3,t}$	0.20% (2.01)	0.32% (2.97)	0.25% (2.41)	-0.03 (-1.02)	-0.03 (-0.92)	0.02 (0.36)	-0.13 (-2.63)	-0.17 (-2.40)	0.13 (5.37)	10%

PANEL B: Benchmark - Performance of All Ideas										
	Mean	α_5	α_6	RMRF	SMB	HML	RMW	CMA	UMD	R ²
$r_{all,t} - r_{f,t}$	0.55% (3.25)	-0.01% (-0.22)	0.00% (-0.03)	0.77 (81.09)	0.19 (14.14)	0.02 (1.26)	0.02 (0.89)	0.02 (0.90)	-0.01 (-1.59)	96%

Table V: Performance of Best Ideas in Hedge Funds (continued)

PANEL C: Performance of Best Ideas for Different Fund Size Groups								
	SMALL (F<100m)		MED (100m<F<10bn)		BIG (F>10bn)		SMALL - BIG	
	a5	a6	a5	a6	a5	a6	a5	a6
$r_{1,t} - r_{f,t}$	0.57% (2.38)	0.58% (2.40)	0.13% (1.32)	0.12% (1.22)	-0.67% (-1.65)	-0.68% (-1.66)	1.24% (2.56)	1.26% (2.58)
$r_{2,t} - r_{f,t}$	0.61% (1.84)	0.61% (1.82)	0.27% (2.30)	0.27% (2.31)	-0.90% (-1.92)	-0.94% (-1.99)	1.51% (2.61)	1.54% (2.65)
$r_{3,t} - r_{f,t}$	0.70% (2.35)	0.71% (2.40)	0.22% (1.72)	0.18% (1.44)	-0.63% (-1.78)	-0.60% (-1.69)	1.33% (2.98)	1.32% (2.94)




Figure 1: Overlap in Best Ideas.

This figure displays the popularity of the stocks that we select as managers' best ideas from 1983-2018. Popularity is defined as the percentage of managers at any point in time that consider a particular stock their best idea. Best ideas are determined within each fund as the stock with the largest ex ante information ratio, $IR_{ift}^{market} = \sigma_{it}(\lambda_{ift} - \lambda_{iMt})$, where λ_{ift} is manager f 's portfolio weight in stock i , λ_{iMt} is the weight of stock i in the market portfolio, and σ_{it} is the most-recent estimate of a stock's CAPM-idiosyncratic volatility.




Figure 2: Measures to identify Best Ideas over time.

This figure graphs the value of the two measures we use to identify the best idea of a portfolio for the median manager over the time period in question as well as their key components. Best ideas are determined within each fund as the stock with the maximum value of one of two possible information ratio measures: 1) $IR_{ift}^{market} = \sigma_{it} \text{tilt}_{ift}^{market}$ or 2) $IR_{ift}^{portfolio} = \sigma_{it} \text{tilt}_{ift}^{portfolio}$, where $\text{tilt}_{ift}^{market} = (\lambda_{ift} - \lambda_{iMt})$ and $\text{tilt}_{ift}^{portfolio} = (\lambda_{ift} - \lambda_{iVt})$, λ_{ift} is manager f 's portfolio weight in stock i , λ_{iMt} is the weight of stock i in the market portfolio, λ_{iVt} is the value weight of stock i in manager f 's portfolio, and σ_{it} is the most-recent estimate of a stock's CAPM-idiosyncratic volatility. In the first panel, we graph the tilt component of these two measures, in the second panel we plot the information ratio measures, and in the third panel, we plot the implied alpha, where the information ratios are multiplied by σ_{it} .




Figure 3: Performance of Best Ideas.

This figure graphs the monthly six-factor alpha of portfolios based on managers' best idea, second-best idea, down to their tenth-best idea. The rank of a manager's ideas are determined by our information ratio measure $IR_{ijt}^{market} = \sigma_{it} tilt_{ijt}^{market}$.




Figure 4: Buy-and-hold abnormal returns of Best Ideas.

This figure graphs the risk-adjusted cumulative buy-and-hold abnormal returns of the best ideas portfolio as identified by our market alpha measures. All cumulative abnormal returns are adjusted using the Fama-French-Carhart six factor model $r_{p,t}^{excess} = a_6 + bRMRF_t + sSMB_t + hHML_t + rRWM_t + cCMA_t + mUMD_t + \varepsilon_{p,t}$, where $r_{p,t}^{excess}$ is $r_{p,t} - r_{f,t}$ the equal-weight excess return on the portfolio of the stocks that represent the best idea of each active manager. Best ideas are determined within each fund as the stock with the maximum value of the following information ratio measure, $IR_{ift}^{market} = \sigma_{it}(\lambda_{ift} - \lambda_{imt})$ where λ_{ift} is manager f 's portfolio weight in stock i , λ_{imt} is the weight of stock i in the market portfolio, and σ_{it} is the most-recent estimate of a stock's CAPM-idiosyncratic volatility. We restrict the analysis to those managers whose maximum position-level information ratio is in the top 25% of all maximum measure at that time. The sample period for the dependent variables is 1983 to 2019 in Panel A and 1996-2019 in Panel B.




Figure 5: Optimal Active Policies Under a Constant Active Allocation Rule.

This figure shows the Sharpe ratios obtained from different allocations to the market (M), a best idea (X), and a second-best idea (Y). In particular, we consider the Sharpe ratio of portfolios where an investor allocates a fixed percentage to an active manager choosing a portfolio of X and Y and puts the remaining capital in M. The riskless rate is zero and the expected returns on the three assets in question are: $E[R_M] = 7\%$, $E[R_X] = 2\%$, and $E[R_Y] = 1\%$. All three assets are uncorrelated and each has the same volatility. Each line on the graph shows the results for a different constraint on the total fraction of assets that are managed by the active manager (i.e. invested in either X or Y). The star on each line represents the myopic allocation by the active manager that maximizes simply his or her own-portfolio Sharpe ratio.




Figure 6: Sharpe Ratios of Equal-weight Portfolios of Active Managers.

This figure shows the Sharpe ratio of equal-weight portfolios of N active managers, each of whom has one best idea, X , that he or she combines with the market portfolio. We assume each investment X has 4% annual alpha, unit beta, the market premium is 6%, the market's annual volatility is 15%, X 's volatility is 40%, the risk-free rate is zero, and that the unsystematic components of each manager's best idea are uncorrelated. The top line shows the Sharpe ratio that would result if managers followed the optimal policy. The lowest line shows the Sharpe ratio the investor will obtain if each manager instead mean-variance optimizes his own portfolio. The middle line gives the resulting Sharpe ratios if managers choose the portfolio that is best for the investor but with the constraint that they cannot sell the market short.