The Information Content of Revealed Beliefs in Portfolio Holdings

Tyler Shumway, Maciej B. Szefler, and Kathy Yuan*†

Current Version: January 2009

Abstract

In this paper, we elicit heterogeneous fund manager beliefs on expected stock returns from funds' portfolio holdings at each quarter-end. Revealed beliefs are extracted by assuming that each fund manager aims to outperform a certain benchmark portfolio by choosing an optimal risk-return tradeoff. We then construct a measure of fund managers' forecasting ability—the *belief accuracy index* (BAI)—by correlating a manager's revealed beliefs on stock returns with the subsequently realized returns. We measure the *differences* in beliefs between funds with high BAI and all other funds, the *belief difference index* (BDI). Sorting stocks based on BDI, we find that the annualized return difference between the top and bottom decile is about two to six percent.

Journal of Economic Literature Classification Codes: G12, E4, C7.

Keywords: Performance Metric, Revealed Beliefs, Belief Accuracy Index (BAI), Belief Difference Index (BDI), Expected Return, and Information Content.

---

*© 2008 by Tyler Shumway (shumway@umich.edu), Maciej B. Szefler (mszefler@umich.edu), and Kathy Yuan (kyuan@umich.edu and K.Yuan@lse.ac.uk), Department of Finance, Stephen M. Ross School of Business, University of Michigan (Shumway, Szefler and Yuan) and London School of Economics (Yuan).

†We would like to thank Clemens Sialm and Lu Zheng for their detailed help. We would also like to thank the seminar participants at the Duke-UNC asset pricing conference and Imperial College, Josh Coval, Mark Grinblatt, Lubos Pastor, Noah Stoffman, and Günter Strobl for their insightful comments. The usual disclaimer applies.

1 Introduction

Portfolio theory is a cornerstone of modern finance. Pioneered by the work of Markowitz (1952), it has led to a vast amount of research exploring optimal portfolios under various constraints (e.g., short-sales), frictions (e.g., heterogeneous information) and computational limits (e.g., estimation of the variance-covariance matrix). This has resulted in a set of recipies for converting portfolio theory's raw ingredients—beliefs on the structure of stock returns—into its finished product, the optimal portfolio. In contrast, little attention has been devoted to the dual problem: extracting beliefs about the structure of expected stock returns from observed portfolio holdings. In this paper, we focus on the information embedded in the cross-sectional portfolio holdings of mutual fund managers, particularly the revealed heterogeneous fund manager beliefs about expected stock returns. We examine whether these revealed beliefs contain information about the skills of mutual fund managers and/or how they are embedded into the prices of common stocks.

To elicit fund managers' beliefs we make three assumptions. First, we assume that mutual fund managers possess heterogeneous beliefs. Second, we assume that each fund manager has a benchmark index. He wishes to outperform his benchmark with the minimum amount of risk subject to a performance target. This objective function is the same as that discussed in Roll (1992) and is commonly observed in practice. Last, we assume that the variance-covariance matrix of asset returns is common knowledge among investors. This assumption is motivated by the empirical finding that estimating the second moment of a return-generating processes from historical data is considerably easier than estimating the first moment. The widely-implemented Black-Litterman model (1992) also adopts this assumption.

In our model, a mutual fund manager's portfolio holdings are the outcome of an optimization based on his beliefs about stock returns (which are specific to him) and about the variance-covariance structure of these returns (which is common across investors). Therefore, his beliefs about stock returns can be easily backed out if the variance-covariance structure is known. Empirically, we estimate the variance-covariance matrix based on historical return data, which are observable to all investors. In estimating the variance-covariance matrix, we use a multi-factor model. We motivate this by noting that multi-factor models are commonly used in the money management industry.

After backing out these revealed beliefs, we construct a measure of fund managers' forecasting ability—the belief accuracy index (BAI)—by correlating each manager's revealed beliefs about stock returns with the subsequently realized returns. We then measure the *differences* in beliefs between the top 30 percent of fund managers ranked by BAI and all the remaining fund managers, the belief difference index (BDI). We construct BDI in this way because we believe that only the best fund managers have forecasting ability that might allow them to outperform the market. We conjecture that ex post returns will be more consistent with beliefs of the best managers than with the beliefs of all other managers. Hence the differences in beliefs between these two groups of managers reveal information not embedded in the stock price: A large positive BDI statistic indicates that the positive information is not embedded into the stock price while a large negative BDI statistic suggests that the negative information is not embedded into the stock price. We sort stocks into deciles according to BDI and examine the subsequent three-month performance across the decile portfolios. The results show that, on average, stocks with higher BDI statistics outperform stocks with lower BDI statistics, indicating that revealed beliefs contain valuable information about future stock returns. We find the annualized performance spread between the top and bottom decile funds is about two to six percent, which is significant, both economically and statistically. These performance differences are not explained by variations in risk or style factors.

It is important to know whether there is information in fund holdings, in part because this information allows us to make some inferences about the degree to which the equity market is informationally efficient. One of the most frequently cited arguments for efficiency is the apparent lack of ability of mutual fund managers. However, Berk and Green (2004) show that managerial ability is consistent with a lack of performance persistence in equilibrium. Therefore, assessing managerial ability requires more powerful techniques than those which simply analyze historical fund returns. Our technique shows that many managers are able to forecast returns.

Recently, there have been various attempts to investigate the information revealed by portfolio holdings in the context of performance evaluation. Grinblatt and Titman (1989); Daniel, Grinblatt, Titman, and Wermers (1997); Graham and Harvey (1996); Wermers (2000); Ferson and Khang (2002); Cohen, Coval, and Pastor (2005); Kacperczyk, Sialm, and Zheng (2005); Kacperczyk, Sialm, and Zheng (forthcoming); and Cremers and Petajisto (2006) have made contributions along this line. There have also been attempts to look beyond historical price return data to future returns. Lo and Wang (2000; 2001) find that turnover satisfies an approximately linear k-factor structure and Goetzmann and Massa

(2006) identify factors through a sample of net flows to nearly 1000 U.S. mutual funds over a year and a half period. Factors embedded in flow and turn-over data are shown to have valuable information for pricing stocks. Wermers, Yao, and Zhao (2007) find that stocks held by top ranked funds (according to measures such as Cohen, Coval, and Pastor (2005)) outperform the rest on average, indicating the investment value of mutual funds. Our paper differs from the existing literature by exploring the information embedded in the cross-sectional portfolio holding for fund managers' beliefs and how these revealed beliefs can be used for predicting future stock returns.

The remainder of this paper is organized as follows. In Section 2, we present our methodology for extracting beliefs from portfolio holdings. Section 3 provides the definitions of BAI and BDI. Section 4 describes the data used and the empirical implementation of the model. In Section 5, we use BAI and managers' heterogenous beliefs on expected returns to construct BDI and evaluate whether BDI has valuable information for predicting future stock returns. We conclude in Section 6.

2 Eliciting Fund Managers' Heterogeneous Beliefs

In this section we present a simple portfolio optimization model to highlight the theoretical foundation for eliciting portfolio managers' heterogeneous beliefs. The objective of this section is to demonstrate how one can back out heterogeneous beliefs about future excess returns from observed portfolio holdings. To do so requires an assumption about the behavior of fund managers. We assume that a fund manager is evaluated relative to some passive benchmark portfolio. Compared with the standard textbook portfolio problem where the investor seeks to minimize return volatility for a given level of expected return, the fund manager in this setup seeks to minimize tracking error volatility for a given level of return *in excess of the benchmark return*. In other words, a fund manager is indifferent to the whims of his benchmark, as long as he can outperform it. As Roll (1992) points out, managers who implement this optimization program do not hold mean-variance efficient portfolios. However, it is clear that tracking error criteria are widely used in practice. Given this, it is reasonable to assume that managers implement such an optimization program. In what follows, we first detail the return-generating process for risky and risk-free assets and the information structure among the fund managers. We then solve the fund manager's portfolio optimization problem. Then, we show how a fund manager's beliefs about stock returns can be identified up to a constant given his (optimal) portfolio holdings and benchmark.

To develop the model, we first focus on the managers' optimization problem. In this problem the investment opportunity set consist of a risk-free asset with a constant return, r_f, and n risky assets where the ith asset's *excess return* over the risk-free rate r_f is denoted as $\tilde{r}_i$ and the excess returns of the n risky assets can be written as $\tilde{\mathbf{r}} = [\tilde{r}_1, ..., \tilde{r}_n]'$, a $n \times 1$ vector . The n risky assets have the following full rank variance-covariance matrix:

$$\mathbf{\Sigma} = \begin{bmatrix} \sigma^2_{1,1} & \cdots & \sigma^2_{1,n} \\ \vdots & \ddots & \vdots \\ \sigma^2_{n,1} & \cdots & \sigma^2_{n,n} \end{bmatrix}, \tag{1}$$

We assume that there are m mutual fund managers in this economy. Fund managers possess a common knowledge of $\mathbf{\Sigma}$, but are heterogeneously informed about the risky assets' excess returns. We use μ_{mi} to denote manager m's belief of asset i's expected excess return. Manager m's belief on the n assets' expected excess returns can be written as $\mu_m = [\mu_{m1}, \ldots, \mu_{mn}]'$, and the beliefs of all the managers can be written as $\mu = [\mu_1, \ldots, \mu_m]'$, which is an $m \times n$ matrix.[1]

Let w_{m0} denote the percentage of wealth (or portfolio weight) invested by manager m in the risk-free asset and let $\mathbf{w}_m = [w_{m1}, \ldots, w_{mn}]'$ denote the vector of portfolio weights in each of the n risky assets. The portfolio weights satisfy the following budget constraint:

$$w_{m0} + \mathbf{1}'\mathbf{w}_m = 1. \tag{2}$$

Each manager has a benchmark portfolio (against which he will be judged), and we denote by $\mathbf{q}_m = [q_{m1}, \ldots, q_{mn}]'$ the vector of manager m's benchmark portfolio weights in each of the n risky assets. We assume that the benchmark consists only of risky assets; consequently, benchmark weights satisfy the following constraint:

$$\mathbf{1}'\mathbf{q}_m = 1. \tag{3}$$

Fund managers choose portfolio weights to maximize the expected return over the benchmark (i.e., active return) and minimize the tracking error volatility (i.e., active risk). We denote manager m's active return by $\tilde{z}_m$, where:

$$\begin{aligned} \tilde{z}_m &= (\mathbf{w}_m - \mathbf{q}_m)'(\tilde{\mathbf{r}} + \mathbf{1}r_0) + w_{m0}r_0 && (4) \\ &= (\mathbf{w}_m - \mathbf{q}_m)\tilde{\mathbf{r}}. && (5) \end{aligned}$$

[1]Unless otherwise noted, we use boldface letters to denote vectors or matrices.

Then, his expected active return is $E[\tilde{z}_m] = (\mathbf{w}_m - \mathbf{q}_m)'\mu_m$ and his active risk is $Var[\tilde{z}_m] = (\mathbf{w}_m - \mathbf{q}_m)'\mathbf{\Sigma}(\mathbf{w}_m - \mathbf{q}_m)$.

In other words, fund manager m, conditional on his beliefs, μ_m, his benchmark, $\mathbf{q}_m$, and his desired level of expected active return, $E[\tilde{z}_m]$, chooses his portfolio weights, $\mathbf{w}_m$, so as to minimize active risk (tracking error volatility):

$$\mathbf{w}_m^* = \underset{\{\mathbf{w}_m, w_{0m}\}}{\operatorname{argmin}} \{(\mathbf{w}_m - \mathbf{q}_m)'\mathbf{\Sigma}(\mathbf{w}_m - \mathbf{q}_m)\}, \tag{6}$$

subject to the budget constraint (2) and

$$E[\tilde{z}_m] = (\mathbf{w}_m - \mathbf{q}_m)'\mu_m. \tag{7}$$

To solve, we construct the Lagrangian

$$L = (\mathbf{w}_m - \mathbf{q}_m)'\mathbf{\Sigma}(\mathbf{w}_m - \mathbf{q}_m) + \lambda_m(E[\tilde{z}_m] - (\mathbf{w}_m - \mathbf{q}_m)'\mu_m), \tag{8}$$

where λ_m is the Lagrangian multiplier for fund manager m. Differentiating with respect to $(\mathbf{w}_m - \mathbf{q}_m)$ obtains the first order condition:

$$(\mathbf{w}_m^* - \mathbf{q}_m) = \lambda_m \mathbf{\Sigma}^{-1}\mu_m. \tag{9}$$

It follows from the budget constraint (2), that the portfolio share allocated to the risk free asset is:

$$w_{0m}^* = -\lambda_m \mathbf{1}'\mathbf{\Sigma}^{-1}\mu_m. \tag{10}$$

Rearranging the terms in the first order condition (9), we obtain an expression for the beliefs of fund manager m:

$$\mu_m = \lambda_m^{-1}\mathbf{\Sigma}(\mathbf{w}_m^* - \mathbf{q}_m), \tag{11}$$

where the only unsolved term is λ_m.

Combining the first order condition (9) and the active return constraint (7), we obtain an expression for the Lagrange multiplier λ_m:

$$\lambda_m = \frac{(\mathbf{w}_m^* - \mathbf{q}_m)'\mathbf{\Sigma}(\mathbf{w}_m^* - \mathbf{q}_m)}{E[\tilde{z}_m]} = \frac{Var[\tilde{z}_m]}{E[\tilde{z}_m]}, \tag{12}$$

which is a constant specific to each fund manager.

Thus, given an estimate of $\mathbf{\Sigma}$, denoted by $\hat{\mathbf{\Sigma}}$, a fund manager's private beliefs, μ_m can be immediately revealed (up to a multiplicative constant, λ_m) by $(\mathbf{w}_m^* - \mathbf{q}_m)$. Since the variance-covariance matrix can be estimated from historical return data, this leads to the extraction of heterogeneous beliefs from portfolio holdings. This finding is summarized in the following result.

Result 1 *For a given investor's portfolio weights $\mathbf{w}_m$, benchmark portfolio $\mathbf{q}_m$, and estimated variance-covariance matrix $\hat{\mathbf{\Sigma}}$, the investor's private beliefs on expected returns, μ_m, are revealed up to a constant:*

$$\hat{\mu}_m = \lambda_m^{-1}\hat{\mathbf{\Sigma}}(\mathbf{w}_m - \mathbf{q}_m). \tag{13}$$

where $\hat{\mu}_m$ denotes an estimate of μ_m and λ_m is an investor-specific constant.

We term these sets of beliefs "*revealed beliefs*" because they are revealed by portfolio holdings. These revealed beliefs have several unique properties. First, they are forward-looking. That is, estimated at one point in time, they are expectations about future stock returns from that point on. Second, they are heterogeneous across fund managers. By choosing different optimal portfolios at the same time, fund managers reveal their differences in opinion. Third, they are inherently dynamic. As portfolio choices may change over time for the same fund manager, his revealed beliefs also vary over time.

3 Evaluating the Information Content of Revealed Beliefs

Since *revealed beliefs* capture a fund manager's *ex-ante* beliefs on future stock returns, how closely these beliefs match the *ex-post* returns is an intuitive measure of the manager's forecasting ability.[2] Specifically, to capture this forecasting ability, we use the correlation between the realized excess returns and the revealed beliefs about excess returns (as computed in Result 1) using all stocks in fund m's portfolio. We term this correlation the "*belief accuracy index* (BAI)."

Definition 1 *For a given set of portfolio revealed beliefs $\hat{\mu}_m$, and a set of ex-post excess*

[2] Of course, this depends on the portfolio holdings being the interior solution to the optimization program just described. In practice, the portfolio holdings may be a corner solution, a solution to a completely different optimization program (e.g. traditional mean-variance optimization), or a solution to a much more complicated program (e.g. one involving the aerodynamics of darts). In any of these cases the relation will obviously not hold. However, it will become clear that this can only weaken our results.

returns $\mathbf{r}$*, the belief accuracy index (BAI) is defined as:*

$$BAI_m \equiv corr(\hat{\mu}_m, \mathbf{r}) \tag{14}$$

where $corr(\cdot)$ *is the sample correlation function.*

If fund managers with higher BAIs possess beliefs on expected stock returns that are close to the realized returns, the extent to which these beliefs differ from the beliefs of the rest of the managers (which can be thought of as the market) indicates the amount of information embedded in the current stock price. The following thought experiment illustrates the intuition of this measure. Managers can be categorized into two categories: "informed" and "less informed." Informed managers are those who are ranked in the top thirty percent of all funds based on BAI. Given the evidence on mutual fund performance, we consider all managers in the bottom seventy percent of the BAI distribution to be less informed. We hypothesize that our informed managers have more accurate beliefs about ex post stock returns than their less informed peers. We assume that current prices are consistent with the beliefs of the less informed fund managers since they are in majority. This characterization is consistent with an equilibrium of the type proposed by Grossman and Stiglitz (1980), where the cost of obtaining information deters a certain fraction of investors from obtaining it. Since the less-informed investors' beliefs simply reflect current market prices, the difference between the informed investors' beliefs and those of the less-informed investors measures how much information held by informed investors is not yet embedded in the stock price.

For any stock, the difference between the informed and less-informed beliefs constitute a measure of information content which should predict future non-systematic price movement, as the information of the informed investors is eventually incorporated into the price.[3] Hence the differences in beliefs between these two groups of managers reveal information not embedded in the stock price: A large positive difference indicates that the positive information is not embedded into the stock price while a large negative difference suggests that the negative information is not embedded into the stock price. We refer to this measure as the "*belief difference index*" or BDI.

Definition 2 *For a given* $(m \times n)$ *matrix of portfolio revealed beliefs* $\hat{\mu}$*, and an* m*-vector of*

[3]In this analysis we only consider those stocks for which we are able to obtain portfolio revealed beliefs from both sets of managers.

BAI rankings, $BAI = [BAI_0, \ldots, BAI_m]'$*, the BDI is an* n*-vector defined as:*

$$BDI \equiv \hat{\mu}' \left(\frac{\iota}{\mathbf{1}'\iota} - \frac{\mathbf{1} - \iota}{m - \mathbf{1}'\iota} \right) \tag{15}$$

where ι *is an* m*-vector with* $\iota_i = 1$ *if fund* i *is among the top 30% of funds in terms of BAI and* $\iota_i = 0$ *otherwise.*

If BDI measures the information not embedded in the current stock prices and this missing information transmits to prices over time, we conjecture that empirically one would observe that stocks with large positive (negative) BDI statistics have larger (smaller) future returns as prices eventually incorporate the positive (negative) information.

In the rest of the paper, we extract revealed beliefs from mutual fund stock holdings and test the above two empirical implications regarding BAI and BDI.

4 Data and Methods

In this section, we begin by describing the data set and the sample selection criteria. We then describe the methodology used to extract fund managers' beliefs. Finally, we describe the tests used to evaluate our predictions.

4.1 Sample

We employ four databases: CRSP stock daily return file, CRSP stock monthly return file, CRSP mutual fund monthly return file and the stock holdings of mutual funds from the CDA/Spectrum Mutual Fund Holding database maintained by Thomson Financial from January 1980 to December 2005. The mutual fund holding database comprises mandatory SEC filings as well as voluntary disclosures by mutual funds. It is typically available quarterly. Wermers (2000) describes this database in more detail.

For this study, we focus on domestic all-equity funds. To construct the sample, we begin with quarterly fund holdings from the CDA/Spectrum Mutual Fund Holding database.[4] We restrict our sample to only those fund-quarters where the Investment Object Code reported by CDA/Spectrum is: aggressive growth, growth, growth and income, unclassified, or missing. We remove all observations where the number of shares held is missing, where the

[4]Throughout, our quarters are calendar year quarters. Any reported holding date in CDA/Spectrum that does not fall on a calendar quarter end is adjusted (into the future) so that it does.

CUSIP of the held security is missing, or where the CUSIP cannot be matched to the CRSP monthly stock return file. We also eliminate any funds that cannot be matched to a fund tracked in the CRSP monthly mutual fund file.[5]

Finally, we eliminate any fund-quarters where the fund's equity holdings amount to less than $10 million in year 2000 dollars (where the fund holdings are adjusted for inflation), or where the fund holds fewer than 20 stocks. [6]

The final sample for the mutual fund holdings contains 76,823 fund-quarters, covers the period 1980-2005 (104 quarters in total), and includes holdings for 3,616 distinct funds. A year by year summary of the sample is found in Table 1.

4.2 Extracting Beliefs

From Result 1, the extraction of fund managers' beliefs requires five elements: the manager's portfolio holdings, the variance-covariance matrix, the fund's benchmark portfolio, the fund's performance target (the expected active return), and a horizon over which to estimate expected returns. Once these elements are known, the calculation is trivial. Hence, we focus here on our handling of these information requirements.

Portfolio Holdings

Our ability to observe portfolio holdings is a major constraint on the methodology. Specifically, we are only able to observe portfolio holdings at the frequency available in the CDA/Spectrum database. Funds typically report holdings on a quarterly basis, so in the best case, our methodology is limited to generating quarterly beliefs.[7] This will obviously limit the power of our tests, as we will be observing no more than 104 quarters to extract

[5]Matching the CDA/Spectrum holdings to the CRSP monthly mutual fund file is done using the MFLINKS programs provided by Wermers (2000) In cases where a single mutual fund has multiple share classes reported in the CRSP file, the monthly returns for the fund are taken to be the value-weight returns of all the share classes. In rare cases, total net assets are not available; for these fund-months, equal-weight returns are used.

[6]The fund's equity holding and number of stocks held is calculated from the reported holdings in CDA/Spectrum. The calculation considers only those stocks that can be matched to the CRSP stock monthly file.

[7]At the end of each calendar quarter, we will generate beliefs for all funds for which recent (less than one year old) holding data is available. If multiple recent holding reports are available, we use the most recent. Some funds do not report on a quarterly basis, hence the beliefs generated for these funds may be more "stale"

beliefs for each fund manager. Another issue is that some fund holdings do not correspond to domestic equity issues for which we have readily available return data. These might be foreign securities, ADRs, bonds, commercial paper, etc.. Theory would require that these securities be included in the analysis, but this is not practical. We deal with this problem by ignoring these holdings in the analysis. This can be justified by noting that our analysis focuses on U.S. equity funds. Such funds typically have negligible holdings in these types of securities.

A similar problem stems from the fact that the portfolio holdings reports do not include information about balances of cash and cash equivalents at quarterly frequencies. This is potentially troublesome: unlike the standard portfolio optimization problem, the optimal portfolio of the benchmark-tracking fund manager is *not* a convex combination of a tangency portfolio and the risk-free asset. The consequence of ignoring cash holdings is a biased estimation of beliefs. That said, we will largely ignore this issue, since that the bias that is introduced will not be large, and it is unusual for the types of funds that we are considering to hold large cash balances. In any event, ignoring cash holdings will only bias against finding any significant predictive power in our measures.

Covariance Matrix

The second challenge in our methodology is obtaining an accurate estimate of the variance-covariance matrix. With financial data, this is always a problematic proposition. The culprit is the relatively small number of historical observations (T) given the large number of securities (N) for which covariances have to be estimated. Typically, N is on the order of a few thousand; while with ten years of monthly data, $T = 120$. When $T < N$, conventional variance-covariance estimation (using the sample covariances) will produce a matrix that is singular, not positive semi-definite, and whose eigenvalues bear little resemblance to the originals.[8] To address this issue we resort to a multiple factor model of the covariance matrix.

Multiple factor model have the advantage of being significantly simpler to estimate and—more importantly—are likely to be more representative of the covariance estimators in use by mutual fund industry. [9] Since our goal is to extract the beliefs of a mutual fund manager by

[8] Schäfer and Strimmer (2005) provide some simulation results that demonstrate the severity of the problem.

[9] Risk models sold to the mutual fund industry by vendors like MSCI Barra typically feature multi-factor covariance matrix estimators.

reverse-engineering his optimization problem, our goal is not to develop the "best" covariance matrix estimator. Rather, we are interested in using a covariance estimator that is as similar as possible to the one used by the mutual fund manager in question.

We model the covariance structure in stock returns using 53 factors. These include the three Fama French (1993) factors, excess market return, MKTRF, small-minus-big, SMB, and high-minus-low (book-to-market), HML. the momentum factor. We also include the momentum, or up-minus-down factor, UMD of Carhart (1997). Finally, we include the returns on the 49 industry portfolios available from Ken French's website. Thus, the data generating process for excess returns is taken to be

$$\tilde{\mathbf{r}} = \alpha + \beta\tilde{\mathbf{f}} + \tilde{\mathbf{e}}$$

where $\tilde{\mathbf{f}}$ is the return to the factors, with covariance matrix Σ_f, and $\tilde{\mathbf{e}}$ is the vector of idiosyncratic returns with zero mean and a diagonal variance structure ($Var[\tilde{\mathbf{e}}] = diag(\sigma_1^2, \ldots, \sigma_n^2)$). The matrix β is taken to be the factor loadings on each of the 53 factors for each of the securities. Idiosyncratic returns are assumed to be uncorrelated with the factor returns, hence the total variance of returns is:

$$\Sigma = \beta\mathbf{\Sigma_f}\beta' + diag(\sigma_1^2, \ldots, \sigma_n^2). \tag{16}$$

To estimate the factor covariance matrix, Σ_f we use the sample covariance estimator. To estimate the factor loadings, β, we regress each security's returns on the contemporanous returns of the 53 factors. We then use the coefficient estimates from these regressions as estimates of the factor loadings. To estimate $\sigma_1^2, \ldots, \sigma_n^2$ we use the sample variances of the residuals from each of the factor-loading regressions.

It would be optimistic to suppose that the variance-covariance matrix is stationary for extended periods. In our model this manifests itself via changes to the covariance of the factor returns, changes to the factor loadings for each security, and changes to the idiosyncratic volatility of each security. To address this, we use relatively short, five year windows to estimate all covariances and factor loadings. In order to maintain a sufficient number of observations we resort to higher frequency return data. Doing so introduces the problem of non-synchronous trading effects. Lo and MacKinlay (1990) show that small stocks may not react to common market news for days, or even weeks. Covariance estimates that do not take this into account understate the degree of co-movement. To deal with this issue, we follow the convention and calculate weekly (Wednesday to Wednesday) returns from the CRSP daily return file. This increases the number of observations four-fold without incurring the brunt of the the non-synchronous trading bias.

In our analysis, we generate a new variance-covariance estimate, $\hat{\mathbf{\Sigma}}$, for each quarter based on the previous five years of weekly returns. For missing weekly return observations, we assume the risk-free rate.

Benchmark Portfolio

Up to this point we have considered the benchmark portfolio, $\mathbf{q}_m$, as given. Unfortunately, empirical realities are quite a bit different. There is no reliable source for fund benchmarks over the entire period under consideration. Furthermore where such data are available there is no guarantee that they are accurate. Funds may (for various reasons) say one thing, and do quite another. Given this, our approach is to let the holdings data speak: if the fund holds or has recently (in the last five years) held some security, then that security is considered to belong to the benchmark. We set the benchmark weights based on market capitalization, thus the fund's benchmark is a value-weighed index of securities in which the fund has shown any interest in the last five years.

Our benchmark selection methodology will include all securities that are part of the true benchmark. This follows from the premise that the manager is mean-variance optimizing. If that is the case, then the solution to the optimization problem will inevitably suggest some non-zero position for every security. Although one could argue that the manager will not hold certain negative positions due to short-sale constraints, in our setting this is not a major issue. Unlike a mean-variance investor, our fund manager will rarely run up against a short-sale constraint. Loosely speaking, this is because to our fund managers, any underweighing of a security relative to his benchmark is effectively a short position.

Performance Target

A fund manager's performance target is required to fully back out his beliefs. Without this, we are only able to obtain his beliefs up to a multiplicative constant. If we only evaluate how closely a fund manager's belief correlates with ex post stock returns, this is not an issue as the correlation is unaffected by scale. However, if we compare *differences* in beliefs across fund managers, proper scaling of the beliefs is important. In our analysis we deal with this by assuming that while fund managers may have different beliefs on the returns of individual securities, they have common beliefs on the dispersion of returns—the standard deviation of beliefs is assumed constant across managers. Thus, scaling each manager's revealed beliefs, $\hat{\mu}_m$ by the standard deviation of his beliefs $((\hat{\mu}_m - \mathbf{1}'\hat{\mu}_m/n)'(\hat{\mu}_m - \mathbf{1}'\hat{\mu}_m/n)/(n-1))$ obtains a measure of belief that is comparable across all managers. We refer to these as *normalized beliefs*.

Horizon

Given that we observe funds' holdings data quarterly, we update our estimates of *revealed belief* using returns for the first month of each quarter. Assuming that a quarter ends at the end of month t, we take the funds' reported holdings at quarter end and calculate $\hat{\mu}_m$ from those data. To calculate a fund's BAI, we correlate realized returns in month $t+1$ with the fund's $\hat{\mu}_m$. To compute a stock's BDI, we take the difference between the mean of top-30% BAI ranked managers' normalized beliefs and the mean of the remaining managers' normalized beliefs for the stock. Finally, we evaluate the forecasting ability of BAI and BDI over months $t+2$, $t+3$ and $t+4$. Some of the forecasting ability of BDI appears to persist for a few months after $t+4$, but it weakens with forecast horizon.

Specifically, at the end of each calendar quarter (end of March, June, September, and December), we calculate the beliefs for each fund manager in the sample using the latest available holding data. We compute the BAI measure by correlating these beliefs with the ex-post returns in the subsequent month (end of April, July, October, and January).[10] We then rank funds into deciles by BAI, compute BDI and sort stocks into deciles by BDI. We form equal- and value-weighted decile portfolios based on BDI and evaluate their performance over the next three months (May to July, August to October, November to January, and February to April).

5 Results

The aim of our empirical analysis is primarily to determine the extent to which the information in funds' holdings can be used to predict stock returns. In this section, we first compute revealed belief for each stock in each fund' portfolio. Using these revealed beliefs, we construct BAI for each fund manager in the sample to capture their forecasting ability. We then measure the information not embedded in the prices by calculating BDI for each

[10]Note that we form BAI by correlating one-month realized returns with expected returns. This one month horizon is similarly used in Spiegel, Mamaysky, and Zhang (2007) for their back testing method.

stock using revealed beliefs and BAI. Finally we test the predictability of BDI for future stock returns.

5.1 BAI

Table 2 presents the full summary statistics for each of the BAI deciles. Funds that are in the top decile have an average correlation of 25.4% while the with the realized the returns. Since we rank funds each quarter, the minimum and maximum correlations of different decile groups frequently overlap each other.

There exists considerable heterogeneity in the composition of the deciles as shown in Table 3. Funds that find themselves in the extreme deciles tend to be smaller, both in terms of assets under management, and the number securities held. For example, in the 1980s a fund in the fifth decile held—on average—94 stocks, while the funds in the 1st and 10th decile held—on average—50 and 51 stocks respectively. A similar relationship holds for assets under management. Statistics for the period of 1990s and the period of 2000-2005 are presented in the other two panels of Table 3 and show similar patterns. These are broadly consistent with the model of Berk and Green (2004). In their model, the costs of obtaining accurate beliefs on security returns is increasing in scale. Hence, as funds grow large, they allocate more and more capital to the pursuit of passive strategies. In our context, such passive strategies correspond to investing in the benchmark portfolio, which adds noise to the belief estimation procedure.

To examine the persistence of the BAI rankings, we consider the decile-to-decile transition probabilities (initial ranking compared to subsequent ranking). Table 4 presents a contingency table of initial and subsequent quarter mutual fund rankings. A three-dimensional histogram of the contingency table is displayed in Figure 1. Table 4 and Figure 1 show that the transition probability of staying in the same decile is the highest for the bottom BAI decile funds (15%) while the top decile funds is the next (13%). Further, the rankings of the top and bottom deciles are very volatile. There are high transition probabilities for switching to the other extreme decile (15% from top to bottom decile and vice versa). By comparison, the middle deciles have a relatively stable ranking. That is, a fund in decile 5 is very likely to remain in the middle deciles (67% from decile 5 to deciles 3 to 8). This quarterly ranking transitional pattern is similar to that reported by Carhart (1997), who comments:

> "... it is apparent that winners are somewhat more likely to remain winners, and

> losers are more likely to either remain losers or perish. However, the funds in the top decile differ substantially each year, with more than 80 percent annual turnover in composition. In addition, last year's winners frequently become next year's losers and vice versa, which is consistent with gambling behavior by mutual funds ... Thus, while the ranks of a few of the top and many of the bottom funds persist, the year-to-year rankings on most funds appear largely random."

Our rankings, however, also show that losers are likely to bounce back to be winners at the quarterly frequency. The turnover can be about 85%. This difference in ranking persistence reveals that in the short-term, when the liquidation risk is not a concern, some losers' gamble can also pay off. Overall, the evidence on the persistence of BAI suggests that it contains potentially valuable information about the ability of fund managers.

5.2 Revealed Beliefs and Future Stock Returns

In this section, we turn to our primary goal—determining the extent to which the information in funds' holdings can be used to predict stock returns. To do so, we construct BDI based on BAI and revealed beliefs. We sort stocks into deciles based on BDI. From these deciles, we construct equal-weight and value-weight portfolios and evaluate the performance for three months starting with the second month of the quarter.

The performance measures we consider include excess returns, alphas from the one factor CAPM model of Sharpe (1963), three-factor model of Fama and French (1993), and four-factor model of Carhart (1997), as well as the characteristic selectivity (CS) measure of Daniel, Grinblatt, Titman, and Wermers (1997). For the various measures of alpha, we run standard time-series regressions. We use the Newey-West standard errors to deal with possible auto-correlation in the error terms.

The risk- and style-adjusted net returns for each equal-weight decile portfolio are reported in Table 5. The second column in this table reports the average returns for stocks in each BDI decile. The next column reports the excess returns (that is, returns over the risk-free rate). The next three columns report the intercepts from a time-series regression based on the one-factor CAPM model, the three-factor model of Fama and French (1993), the four-factor model of Carhart (1997), and the characteristic selectivity measure of Daniel, Grinblatt, Titman, and Wermer (1997), respectively.

Table 5 shows that future stock performance is generally monotonic in BDI. These results show that the difference in beliefs on future returns between the informed managers and the less informed managers is a good predictor of future stock returns: stocks in the top deciles outperform those in the bottom deciles. Investing in stocks in the top decile and shorting those in the bottom decile results in a significant monthly return of between 15 and 29 basis points or 1.80 and 3.48 percent per annum, depending on the measure.

These results are all statistically significant at the ten percent level, and most are significant at the five percent level. They do not seem to be explained by risk factor loadings. Strongest results obtain for the difference between top decile stocks and bottom decile stocks. Results for the difference between the nine-to-ten and one-to-two decile are somewhat weaker—19 to 23 basis points per month—but still generally significant at the five percent level.

Further, much of the return predictability in BDI is driven by stocks in the highest BDI deciles. Looking specifically at the returns reported in Table 5, it is clear that most of the positive alphas reported for deciles 9 and 10 of stocks sorted by BDI are significantly positive. This is important, since it is not possible to short open-end mutual funds and it is sometimes costly to short stocks. Apparently a significant portion of the return predictability that we document is driven by successful fund managers that choose stocks that outperform the market.

Table 6 shows the characteristics of the equal-weighted BDI decile portfolios. The loading on the market factor is very close to one for almost all the deciles except that the worst performing decile has a market beta close to 2. The factor loadings for SMB follow a "U" shape and for HML and UMD an inverted "U" shape. That is, stocks in the middle deciles—those stocks on which the informed and less informed managers agree—are bigger, have higher book-to-market, and have performed better in the past than the stocks on which there is no consensus.[11] In other words, there appears to be more disagreement about small, growth stocks, with poor performance track records. This is consistent with intuition.

Table 7 and Table 8 report the results for the value-weight decile portfolios. The difference between the top and bottom value-weight deciles is a bit larger than the difference between the top and bottom equal-weight deciles in magnitude and statistical significance for all measures except characteristic selectivity. Excluding characteristic selectivity, the results for the difference between the nine-to-ten and one-to-two deciles are between 42 and 48 basis

[11]More precisely, the stocks in the "consensus" deciles have a lower loading on the size factor, and higher loadings on the value and momentum factors.

points per month, or 5.16 and 5.91 percent per annum, approximately twice as large as the equal-weight results. This is somewhat unexpected as value-weighing tends to emphasize the large stocks for which it should be comparatively more difficult to obtain information not already known to the market. On the other hand, the characteristic selectivity measure is not statistically significant, indicating also that the value-weight portfolio performance might be more volatile than and not as robust as the equal-weight portfolio results, indicating that stock picking skills most origin from small stocks. The factor loadings for the value-weight deciles are similar to the equal-weight results,

Overall the evidence supports our conjecture that BDI contains information that is valuable for forecasting stock returns. We also conduct various robustness checks about our results. First, we use various alternative methods to estimate the variance-covariance matrix. For example, we refine the BARA estimation by constructing the factor loadings as the sum of the coefficients on the contemporaneous and the lagged factor returns. The results are similar as shown in Table 9. We also estimate revealed beliefs using a diagonal variance-variance matrix where the diagonal terms are estimated using sample variances of the corresponding stock returns. The results are weaker when we ignore the off-diagonal terms, as shown in Table 10. We also estimated the variance-covariance matrix using a shrinkage estimator as in Ledoit and Wolf (2003). The shrinkage estimator yields bigger coefficient but at a lower level of statistical significance, indicating that the shrinkage estimator is a noisier estimate for the variance-covariance matrix. Second, to weed out the noise in the BAI constructed using just one-quarter holding information, we capture the forecasting skills of a manager by computing his average BAI over the past eight quarters. The results are similar.

Finally, we also ensure that the results are not driven by short-term return continuations. Given the nature of our tests, it is possible that the results are driven by some sort of short-term return continuation rather than fund manager skill. Suppose, for example, that the managers with the highest BAI have no particular skill, but are randomly lucky in any given quarter to have a high correlation between their revealed beliefs and returns in the subsequent month. If the stocks that have done well in one month continue to have relatively high returns in the subsequent quarter, a positive correlation between BAI and subsequent fund performance may result. While this story appears to go against the literature on short-term return reversals, it is easy to check whether it drives the results.

We check the robustness of the results by delaying all the holding reports by two quarters and running the entire analysis again. That is, we treat holdings reported at quarter t as if they are reported at quarter $t+2$. We then construct BAI using these lagged portfolios. According to our model, the information embedded in these lagged portfolios is stale and should not reflect managerial skills. Consequently, none of results should obtain. Alternatively, if our results are driven by the return continuation story described above, BAI and BDI constructed using the these "random" portfolios could be related to future stock return. Our analysis confirms that the former is the case. We find that using lagged portfolio holdings to perform our test reduces the return predictability of our tests to close to zero. This gives us confidence that the predictability we document truly results from manager ability rather than from the nature of our model and estimation algorithm. For brevity, we do not report the numerical results.

6 Conclusion

This paper aims to examine the information content of revealed beliefs of mutual fund managers. The revealed beliefs are backed out by reverse-engineering fund manger's portfolio optimization problem. Specifically, to elicit the revealed beliefs, we assume that each manager rationally optimize over the risk return tradeoff relative to his own benchmark portfolio. The key idea is that managers tilt their portfolios toward stocks with better risk-return tradeoffs according to their private beliefs. Hence, observing their holdings, one can determine whether fund managers' beliefs on future returns are accurate.

Based on these revealed beliefs, we propose a new fund performance measure—BAI—based on how closely fund managers' beliefs regarding future stock returns match realized returns. We measure the differences in beliefs between minority informed funds (those with higher BAI) and the rest, which is BDI. The evidence in this paper suggests that investors could profit from extracting fund manager information about individual stock returns through the BDI measure. Further, the BAI and BDI measures appear to contain information that is not in existing measures. Theoretically, both metrics contribute to the general finance literature by extracting information contained in cross-sectional fund holdings through exploiting a portfolio optimization framework.

More fundamentally, the paper makes a unique contribution to the finance literature by introducing a revealed preference approach to measuring investor expectations. That is, instead of estimating investor expectation regarding risk and returns from historical returns, we show that it can be useful to back out investors' expectations regarding returns (and potentially risks) from their portfolio holdings. This approach may be of great empirical importance for future work. For example, various strands of the market microstructure literature are built on the assumption that investors are heterogeneously or asymmetrically informed. Without a concrete measure of investor beliefs, most empirical tests of these theories are based on equilibrium price patterns, which might suffer from endogeneity and measurement error. Having a relatively direct measure of investor beliefs might help researchers identify the degree of information asymmetry at a point in time or among a set of investors. Similarly, the asset pricing literature often involves estimation of dynamically changing expected returns. The information provided in portfolio holdings on investor belief about expected future returns might improve existing estimation techniques, and hence have important implications for empirical asset pricing.

References

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

Black, Fisher, and Bob Litterman, 1992, Global portfolio optimization, *Financial Analysts Journal* 48, 28–43.

Carhart, Mark M., 1997, On persistence in mutual fund performance, *Journal of Finance* 50, 679–698.

Cohen, Randy B., Joshua D. Coval, and L. Pastor, 2005, Judging fund managers by the company they keep, *Journal of Finance* 60, 1057–1096.

Cremers, Martijn, and Antti Petajisto, 2006, How active is your fund manager? A new measure that predicts performance, Yale International Center for Finance Working Paper.

Daniel, Kent, M. Grinblatt, Sheridan Titman, and Russ Wermers, 1997, Measuring mutual fund performance with characteristic-based benchmarks, *Journal of Finance* 52, 1035–1058.

Fama, Eugene, and Kenneth French, 1993, Common risk factors in returns on stocks and bonds, *Journal of Financial Economics* 51, 3–56.

Ferson, Wayne E., and Kenneth Khang, 2002, Conditional performance measurement using portfolio weights: Evidence for pension funds, *Journal of Financial Economics* 65, 249–282.

Goetzmann, William N., and Massimo Massa, 2006, Behavior factors in mutual fund flows, Yale International Center for Finance Working Paper.

Graham, John, and Campbell Harvey, 1996, Market timing ability and volatility implied in investment newsletters' asset allocation recommendations, *Journal of Financial Economics* 42, 397–421.

Grinblatt, M., and Sheridan Titman, 1989, Mutual fund performance: An analysis of quarterly portfolio holdings, *Journal of Business* 62, 393–416.

Grossman, Sanford J., and Joseph E. Stiglitz, 1980, On the Impossibility of Informationally Efficient Markets, *The American Economic Review* 70, 393–408 ArticleType: primary_article / Full publication date: Jun., 1980 / Copyright ? 1980 American Economic Association.

Kacperczyk, Marcin T., Clemens Sialm, and Lu Zheng, 2005, On industry concentration of actively managed equity mutual funds, *Journal of Finance* 60, 1983–2011.

Kacperczyk, Marcin T., Clemens Sialm, and Lu Zheng, forthcoming, Unobserved actions of mutual funds, *Review of Financial Studies*.

Ledoit, Olivier, and Michael Wolf, 2003, Improved estimation of the covariance matrix of stock returns with an application to portfolio selection, *Journal of Empirical Finance* 10, 603–621.

Lo, Andrew, and Jiang Wang, 2000, Trading volume: Definition, data analysis, and implications of portfolio theory, *Review of Financial Studies* 13, 257–300.

Lo, Andrew, and Jiang Wang, 2001, Trading volume: Implications of an intertemporal capital asset pricing model, MIT Sloan School of Management Working Paper.

Lo, Andrew W., and A. Craig MacKinlay, 1990, When are Contrarian Profits Due to Stock Market Overreaction?, *The Review of Financial Studies* 3, 175–205.

Markowitz, Harry, 1952, Portfolio selection, *Journal of Finance* 7, 77–91.

Roll, R., 1992, A Mean/Variance Analysis of Tracking Error, *Journal of Portfolio Management* 18, 13–22.

Schäfer, J., and K. Strimmer, 2005, A Shrinkage Approach to Large-Scale Covariance Matrix Estimation and Implications for Functional Genomics, *Statistical Applications in Genetics and Molecular Biology* 4, 32.

Sharpe, William F., 1963, A Simplified Model for Portfolio Analysis, *Management Science* 9, 277–293.

Spiegel, Matthew, Harry Mamaysky, and Hong Zhang, 2007, Improved forecasting of mutual fund alphas and betas, *Review of Finance* 11, 359–400.

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 Ji Zhao, 2007, The investment value of mutual fund portfolios, University of Maryland Working Paper.

Table 1: Summary Statistics: Sample

For each year (row) below, descriptive statistics are presented on the funds in the sample. We begin with mutual fund holdings data provided by CDA/Thomson Financial. We eliminate non-equity funds (based on IOC code), any funds that do not appear in the CRSP mutual fund monthly file, funds that hold fewer than 20 stocks, and funds that have equity holdings totaling less than $10 million (in year 2000 dollars) under management. The reported number of observations (N) is the number of fund-quarter pairs in each year. The reported number of funds (Funds) is the number of distinct funds. Summary statistics for the size of equity holdings (Holdings), and the average number of individual stocks held by a fund (# Stocks Held) are also reported.

			Holdings (000s)			# Stocks Held
Year	N	Funds	Min.	Avg.	Max.	Avg.
1980	915	264	4,816	127,940	1,527,751	57
1981	917	262	5,283	136,666	1,631,473	61
1982	768	259	5,629	146,079	1,714,679	60
1983	913	298	5,806	217,073	1,993,302	70
1984	1,032	318	6,204	202,675	2,127,687	71
1985	1,144	348	6,285	239,567	3,210,548	72
1986	1,317	411	6,414	273,084	6,161,852	75
1987	1,500	458	6,634	317,742	9,657,842	77
1988	1,579	476	6,876	262,421	8,099,248	80
1989	1,601	526	7,201	294,831	10,488,253	81
1990	1,674	551	7,615	290,079	9,570,117	83
1991	1,931	668	7,963	329,781	15,677,987	88
1992	2,221	796	8,198	400,399	15,805,327	98
1993	2,561	1,049	8,470	454,237	27,445,811	111
1994	2,785	1,215	8,611	435,673	33,620,511	122
1995	2,988	1,365	8,858	579,129	47,844,372	125
1996	3,431	1,516	9,117	680,727	43,649,913	119
1997	4,086	1,708	9,334	845,305	56,631,408	121
1998	4,573	1,809	9,472	1,042,077	73,826,198	121
1999	5,118	1,989	9,685	1,161,186	104,561,516	119
2000	5,920	2,205	10,004	1,257,767	106,348,255	131
2001	5,680	2,181	10,274	972,476	90,996,827	136
2002	6,349	2,211	10,454	834,316	76,572,887	139
2003	6,742	2,305	10,690	844,330	92,935,768	143
2004	7,655	2,333	11,022	1,094,943	105,889,838	146
2005	7,951	2,339	11,350	1,188,854	106,591,008	148
2006	7,724	2,259	11,766	1,394,911	120,049,670	146

Table 2: BAI Deciles: Summary Statistics

At the end of each calendar quarter (end of March, June, September, and December), we calculate the revealed beliefs for each fund manager in the sample using the latest available holding data. We compute the BAI measure (the correlation between revealed beliefs about returns and realized returns) by correlating these beliefs with the ex-post returns in the subsequent month (end of April, July, October, and January). We rank funds into deciles based on BAI. This table presents the mean, standard deviation, minimum, and maximum of the BAI for each decile. Decile 1 denotes the least accurate group, Decile 10 the most accurate.

	BAI			
BAI Decile	Mean	Std. Dev.	Min	Max
1	-0.259	0.092	-0.800	-0.116
2	-0.139	0.038	-0.274	-0.029
3	-0.087	0.031	-0.195	0.005
4	-0.050	0.026	-0.145	0.039
5	-0.018	0.022	-0.096	0.074
6	0.013	0.021	-0.062	0.114
7	0.046	0.024	-0.030	0.155
8	0.085	0.029	0.010	0.218
9	0.137	0.036	0.044	0.279
10	0.254	0.089	0.103	0.867

Table 3: BAI Deciles: Composition

At the end of each calendar quarter (end of March, June, September, and December), we calculate the revealed beliefs for each fund manager in the sample using the latest available holding data. We compute the BAI measure (the correlation between revealed beliefs about returns and realized returns) by correlating these beliefs with the ex-post returns in the subsequent month (end of April, July, October, and January). We rank funds into deciles based on BAI. For each BAI decile, we report summary statistics for funds that find themselves in that decile. These include the mean, standard deviation and median of the value of fund holdings and of the number of securities held. As there has been considerable growth in the mutual fund industry, we break the table into panels, with each panel covering one decade of the study period. Decile 1 (Decile 10) contains the lowest (highest) BAI-ranked funds.

	Fund Holdings ($ 000s)			Number of Stocks Held		
BAI Rank	Mean	Std Dev	Median	Mean	Std. Dev.	Median
Panel A: 1980–1989						
1	169,191	322,627	57,732	50	32	43
2	251,158	538,526	75,435	62	50	52
3	240,655	393,231	101,833	72	61	58
4	270,974	537,756	101,632	82	109	61
5	297,776	679,119	104,721	94	152	61
6	325,956	737,642	99,148	98	166	64
7	257,362	422,682	108,453	87	110	61
8	240,181	458,519	95,523	74	77	58
9	234,508	434,606	84,155	65	55	55
10	181,488	419,069	59,613	51	25	44
Panel B: 1990–1999						
1	405,158	1,413,203	94,020	59	36	50
2	635,729	1,857,830	140,063	82	73	64
3	686,764	2,244,629	146,973	103	111	74
4	842,186	2,668,491	174,675	145	246	80
5	917,113	3,532,401	188,100	176	294	86
6	891,938	2,895,567	184,064	171	286	85
7	954,140	3,508,663	183,672	153	282	79
8	908,626	3,162,529	168,886	113	134	73
9	758,521	2,693,410	145,379	94	112	66
10	416,553	1,110,079	95,099	62	42	51
Panel C: 2000–2005						
1	781,229	3,009,239	135,058	70	61	54
2	1,001,563	3,654,687	187,550	100	115	73
3	1,160,779	4,423,039	184,590	144	195	84
4	1,125,531	4,179,440	189,756	170	295	87
5	1,050,595	3,726,874	184,379	197	368	91
6	1,002,403	3,249,338	194,572	184	341	91
7	1,100,510	4,019,846	181,073	160	271	85
8	1,178,919	4,925,113	179,830	146	241	81
9	1,004,754	3,239,279	178,571	111	131	73
10	817,052	2,924,540	147,853	83	113	56

Table 4: Transition Probabilities: Initial vs. Subsequent BAI Decile

At the end of each calendar quarter (end of March, June, September, and December), we calculate the revealed beliefs for each fund manager in the sample using the latest available holding data. We compute the BAI measure (the correlation between revealed beliefs about returns and realized returns) by correlating these beliefs with the ex-post returns in the subsequent month (end of April, July, October, and January). For each quarter (starting in Q1 1980 and ending in Q4 2005) we sort funds into deciles based on our BAI measure. Each BAI rank is matched to the subsequent ranking of the same fund. The rows (columns) in the table correspond to the initial (subsequent) decile rank, where Decile 1 (Decile 10) contains the lowest (highest) BAI-ranked funds. Each cell contains the probability of the subsequent ranking conditional on the initial ranking. Most frequently, the subsequent ranking is from the subsequent quarter, however this need not be the case (some funds do not report at a quarterly frequency).

	1	2	3	4	5	6	7	8	9	10
1	0.15	0.11	0.09	0.077	0.076	0.078	0.08	0.09	0.10	0.15
2	0.10	0.10	0.10	0.1	0.09	0.094	0.096	0.10	0.11	0.10
3	0.081	0.098	0.10	0.11	0.10	0.11	0.11	0.10	0.10	0.084
4	0.08	0.097	0.11	0.11	0.11	0.11	0.11	0.10	0.095	0.076
5	0.073	0.094	0.11	0.11	0.12	0.11	0.11	0.11	0.092	0.071
6	0.074	0.098	0.11	0.11	0.12	0.11	0.10	0.10	0.10	0.074
7	0.079	0.099	0.098	0.12	0.11	0.12	0.10	0.11	0.096	0.078
8	0.081	0.097	0.10	0.11	0.11	0.098	0.11	0.11	0.11	0.084
9	0.10	0.11	0.1	0.095	0.098	0.096	0.099	0.096	0.098	0.11
10	0.15	0.1	0.09	0.087	0.078	0.08	0.087	0.092	0.10	0.13

Table 5: Equal-Weight BDI-Decile Stock Portfolios: Performance

One month into each quarter we sort stocks into deciles based on the difference in beliefs between the "informed" fund-managers (those above the 70th percentile according to our BAI measure) and the rest using the previous quarter's holding reports. Based on these (BDI) deciles we create equal-weight portfolios that we buy at the start of the second month of the quarter. We hold the portfolio for 3 months, repeating the procedure in the subsequent quarter. For each decile portfolio, we present the monthly average returns, the excess returns (over the risk free rate), the CAPM, Fama-French, and Carhart alphas, as well as the characteristic selectivity (CS) measure of Daniel, Grinblatt, Titman, and Wermers (1997). We also provide results for long-short portfolios constructed by buying the top one (two, five) decile(s) and selling the bottom one (two, five) decile(s). Decile 1 (Decile 10) contains the lowest (highest) BDI-ranked stocks.

Decile	Average Return	Excess Return	CAPM Alpha	Fama-French Alpha	Carhart Alpha	DGTW CS
1	0.0115	0.0068	-0.0018	-0.0032	0.0010	0.0003
	(0.0041)**	(0.0041)*	(0.0021)	(0.0014)*	(0.0014)	(0.0013)
2	0.0120	0.0073	-0.0000	-0.0021	0.0006	-0.0001
	(0.0033)**	(0.0033)*	(0.0017)	(0.0011)*	(0.0009)	(0.0008)
3	0.0132	0.0085	0.0011	-0.0008	0.0013	0.0011
	(0.0032)**	(0.0032)**	(0.0016)	(0.0010)	(0.0010)†	(0.0007)†
4	0.0135	0.0088	0.0018	-0.0004	0.0013	0.0006
	(0.0030)**	(0.0030)**	(0.0015)	(0.0009)	(0.0009)†	(0.0006)
5	0.0134	0.0087	0.0018	-0.0007	0.0008	0.0009
	(0.0029)**	(0.0029)**	(0.0015)	(0.0009)	(0.0008)	(0.0006)†
6	0.0135	0.0087	0.0017	-0.0007	0.0005	0.0008
	(0.0030)**	(0.0030)**	(0.0015)	(0.0008)	(0.0007)	(0.0005)†
7	0.0132	0.0085	0.0017	-0.0007	0.0008	0.0004
	(0.0030)**	(0.0030)**	(0.0015)	(0.0008)	(0.0008)	(0.0005)
8	0.0139	0.0091	0.0020	-0.0002	0.0017	0.0010
	(0.0031)**	(0.0031)**	(0.0015)†	(0.0008)	(0.0009)*	(0.0006)†
9	0.0137	0.0089	0.0015	-0.0005	0.0020	0.0008
	(0.0033)**	(0.0033)**	(0.0016)	(0.0009)	(0.0010)*	(0.0006)
10	0.0136	0.0089	0.0006	-0.0003	0.0032	0.0018
	(0.0039)**	(0.0039)*	(0.0019)	(0.0011)	(0.0014)*	(0.0012)†
Top 10% -	0.0021	0.0021	0.0024	0.0029	0.0021	0.0015
Bottom 10%	(0.0011)*	(0.0011)*	(0.0012)*	(0.0012)**	(0.0011)*	(0.0012)†
Top 20% -	0.0019	0.0019	0.0020	0.0023	0.0017	0.0012
Bottom 20%	(0.0008)*	(0.0008)*	(0.0009)*	(0.0009)**	(0.0009)*	(0.0009)†

Table 6: Equal-Weight BDI-Decile Stock Portfolios: Four Factor Loadings

For each decile portfolio in Table 5, we present the intercept and the "betas" for market, small-minus-big (SMB), high book-to-market minus low book-to-market (HML), and momentum (UMD) portfolios. Decile 1 (Decile 10) contains the lowest (highest) BDI-ranked stocks.

Decile	N	Alpha	Market	SMB	HML	UMD
1	312	0.0010	1.1643	0.8914	0.0125	-0.4049
		(0.0014)	(0.0347)**	(0.0726)**	(0.0752)	(0.0618)**
2	312	0.0006	1.0688	0.6740	0.1679	-0.2610
		(0.0009)	(0.0284)**	(0.0607)**	(0.0539)**	(0.0415)**
3	312	0.0013	1.0850	0.5924	0.1746	-0.1984
		(0.0010)†	(0.0222)**	(0.0586)**	(0.0550)**	(0.0418)**
4	312	0.0013	1.0568	0.5439	0.2220	-0.1599
		(0.0009)†	(0.0217)**	(0.0612)**	(0.0497)**	(0.0386)**
5	312	0.0008	1.0699	0.5216	0.2733	-0.1479
		(0.0008)	(0.0240)**	(0.0610)**	(0.0599)**	(0.0328)**
6	312	0.0005	1.0887	0.5344	0.2779	-0.1150
		(0.0007)	(0.0199)**	(0.0532)**	(0.0511)**	(0.0319)**
7	312	0.0008	1.0470	0.5735	0.2639	-0.1439
		(0.0008)	(0.0245)**	(0.0587)**	(0.0478)**	(0.0400)**
8	312	0.0017	1.0674	0.6207	0.2165	-0.1764
		(0.0009)*	(0.0229)**	(0.0487)**	(0.0440)**	(0.0375)**
9	312	0.0020	1.0757	0.7105	0.1592	-0.2330
		(0.0010)*	(0.0207)**	(0.0488)**	(0.0396)**	(0.0391)**
10	312	0.0032	1.0957	0.9033	-0.0425	-0.3297
		(0.0014)*	(0.0291)**	(0.0466)**	(0.0604)	(0.0616)**
Top 10% -	312	0.0021	-0.0686	0.0119	-0.0550	0.0752
Bottom 10%		(0.0011)*	(0.0294)**	(0.0640)	(0.0661)	(0.0360)*
Top 20% -	312	0.0017	-0.0308	0.0242	-0.0318	0.0516
Bottom 20%		(0.0009)*	(0.0254)	(0.0437)	(0.0517)	(0.0299)*

Table 7: Value-Weight BDI-Decile Stock Portfolios: Performance

One month into each quarter we sort stocks into deciles based on the difference in beliefs between the "informed" fund-managers (those above the 70th percentile according to our BAI measure) and the rest using the previous quarter's holding reports. Based on these (BDI) deciles we create value-weight portfolios that we buy at the start of the second month of the quarter. We hold the portfolio for 3 months, repeating the procedure in the subsequent quarter. For each decile, we present the monthly average returns, the excess returns (over the risk free rate), the CAPM, Fama-French, and Carhart alphas, as well as the characteristic selectivity (CS) measure of Daniel, Grinblatt, Titman, and Wermers (1997). We also provide results for long-short portfolios constructed by buying the top one (two, five) decile(s) and selling the bottom one (two, five) decile(s). Decile 1 (Decile 10) contains the lowest (highest) BDI-ranked stocks.

Decile	Average Return	Excess Return	CAPM Alpha	Fama-French Alpha	Carhart Alpha	DGTW CS
1	0.0081	0.0034	-0.0042	-0.0034	-0.0025	-0.0007
	(0.0031)**	(0.0031)	(0.0014)**	(0.0015)*	(0.0015)*	(0.0012)
2	0.0105	0.0058	-0.0012	-0.0011	-0.0003	-0.0010
	(0.0028)**	(0.0028)*	(0.0013)	(0.0013)	(0.0013)	(0.0010)
3	0.0107	0.0060	-0.0008	-0.0008	-0.0010	-0.0012
	(0.0026)**	(0.0026)**	(0.0008)	(0.0008)	(0.0009)	(0.0006)*
4	0.0126	0.0079	0.0016	0.0014	0.0007	0.0008
	(0.0023)**	(0.0023)**	(0.0008)*	(0.0008)*	(0.0008)	(0.0006)†
5	0.0124	0.0077	0.0013	0.0004	0.0008	0.0003
	(0.0024)**	(0.0024)**	(0.0007)*	(0.0006)	(0.0006)	(0.0005)
6	0.0126	0.0079	0.0016	0.0009	0.0004	0.0005
	(0.0025)**	(0.0025)**	(0.0008)*	(0.0007)	(0.0007)	(0.0005)
7	0.0118	0.0070	0.0007	0.0006	0.0005	0.0000
	(0.0025)**	(0.0025)**	(0.0009)	(0.0009)	(0.0009)	(0.0006)
8	0.0104	0.0057	-0.0011	-0.0008	-0.0001	-0.0003
	(0.0028)**	(0.0028)*	(0.0012)	(0.0009)	(0.0009)	(0.0007)
9	0.0117	0.0070	0.0001	0.0001	0.0009	0.0004
	(0.0027)**	(0.0027)**	(0.0008)	(0.0008)	(0.0009)	(0.0007)
10	0.0129	0.0082	0.0002	0.0010	0.0017	0.0001
	(0.0034)**	(0.0034)**	(0.0014)	(0.0014)	(0.0014)	(0.0011)
Top 10% -	0.0048	0.0048	0.0044	0.0044	0.0042	0.0009
Bottom 10%	(0.0020)**	(0.0020)**	(0.0021)*	(0.0023)*	(0.0023)*	(0.0016)
Top 20% -	0.0030	0.0030	0.0029	0.0028	0.0027	0.0011
Bottom 20%	(0.0015)*	(0.0015)*	(0.0017)*	(0.0017)*	(0.0018)†	(0.0012)

Table 8: Value-Weight BDI-Decile Stock Portfolios: Four Factor Loadings

For each decile portfolio in Table 7, we present the intercept and the "betas" for market, small-minus-big (SMB), high book-to-market minus low book-to-market (HML), and momentum (UMD) portfolios. Decile 1 (Decile 10) contains the lowest (highest) BDI-ranked stocks.

Decile	N	Alpha	Market	SMB	HML	UMD
1	312	-0.0025	1.0735	0.1439	-0.1552	-0.0864
		(0.0015)*	(0.0513)**	(0.0851)*	(0.0790)*	(0.0397)*
2	312	-0.0003	1.0374	0.0191	-0.0408	-0.0806
		(0.0013)	(0.0435)**	(0.0790)	(0.0847)	(0.0536)†
3	312	-0.0010	1.0422	-0.0463	0.0228	0.0160
		(0.0009)	(0.0248)**	(0.0444)	(0.0675)	(0.0424)
4	312	0.0007	0.9905	-0.1078	0.0619	0.0697
		(0.0008)	(0.0178)**	(0.0334)**	(0.0474)†	(0.0286)**
5	312	0.0008	1.0200	-0.0454	0.1303	-0.0339
		(0.0006)	(0.0197)**	(0.0390)	(0.0338)**	(0.0219)†
6	312	0.0004	1.0199	-0.1081	0.1305	0.0475
		(0.0007)	(0.0253)**	(0.0265)**	(0.0365)**	(0.0199)**
7	312	0.0005	0.9866	-0.0901	0.0249	0.0088
		(0.0009)	(0.0253)**	(0.0367)**	(0.0581)	(0.0326)
8	312	-0.0001	1.0182	-0.0816	-0.0578	-0.0634
		(0.0009)	(0.0255)**	(0.0409)*	(0.1088)	(0.0602)
9	312	0.0009	1.0139	0.0826	-0.0237	-0.0762
		(0.0009)	(0.0256)**	(0.0291)**	(0.0445)	(0.0287)**
10	312	0.0017	1.1008	0.3103	-0.1708	-0.0625
		(0.0014)	(0.0304)**	(0.0404)**	(0.0663)**	(0.0443)†
Top 10% -	312	0.0042	0.0273	0.1664	-0.0157	0.0240
Bottom 10%		(0.0023)*	(0.0559)	(0.0990)*	(0.1136)	(0.0578)
Top 20% -	312	0.0027	0.0019	0.1150	0.0007	0.0142
Bottom 20%		(0.0018)†	(0.0502)	(0.0824)†	(0.0993)	(0.0541)

Table 9: Robustness Check using Lagged BARA: Equal-Weight BDI Decile Stock Portfolios

BDI and BAI are computed using a variance-covariance matrix that is estimated using lagged BARA approach. In this approach, the factor loadings of stocks are the sum of the coefficients on the contemporaneous and lagged factor returns. One month into each quarter we sort stocks into deciles based on the difference in beliefs between the "informed" fund-managers (those above the 70th percentile according to our BAI measure) and the rest using the previous quarter's holding reports. Based on these (BDI) deciles we create value-weight portfolios that we buy at the start of the second month of the quarter. We hold the portfolio for 3 months, repeating the procedure in the subsequent quarter. For each decile, we present the monthly average returns, the excess returns (over the risk free rate), the CAPM, Fama-French, and Carhart alphas, as well as the characteristic selectivity (CS) measure of Daniel, Grinblatt, Titman, and Wermers (1997). We also provide results for long-short portfolios constructed by buying the top one (two, five) decile(s) and selling the bottom one (two, five) decile(s). Decile 1 (Decile 10) contains the lowest (highest) BDI-ranked stocks.

Decile	Average Return	Excess Return	CAPM Alpha	Fama-French Alpha	Carhart Alpha	DGTW Excess
1	0.0119	0.0071	-0.0011	-0.0027	0.0009	-0.0001
	(0.0039)**	(0.0039)*	(0.0020)	(0.0012)*	(0.0013)	(0.0010)
2	0.0131	0.0084	0.0009	-0.0012	0.0014	0.0006
	(0.0034)**	(0.0034)**	(0.0018)	(0.0010)	(0.0010)†	(0.0007)
3	0.0137	0.0090	0.0019	-0.0002	0.0018	0.0010
	(0.0031)**	(0.0031)**	(0.0016)	(0.0009)	(0.0010)*	(0.0006) †
4	0.0142	0.0095	0.0024	0.0001	0.0020	0.0013
	(0.0030)**	(0.0030)**	(0.0015)†	(0.0009)	(0.0008)**	(0.0006) *
5	0.0137	0.0089	0.0019	-0.0002	0.0012	0.0009
	(0.0029)**	(0.0029)**	(0.0015)†	(0.0010)	(0.0010)†	(0.0006) †
6	0.0137	0.0090	0.0021	-0.0003	0.0010	0.0009
	(0.0029)**	(0.0029)**	(0.0015)†	(0.0009)	(0.0009)	(0.0006) †
7	0.0128	0.0081	0.0011	-0.0014	0.0002	-0.0002
	(0.0031)**	(0.0031)**	(0.0016)	(0.0008)†	(0.0008)	(0.0006)
8	0.0136	0.0088	0.0018	-0.0004	0.0011	0.0006
	(0.0031)**	(0.0031)**	(0.0016)	(0.0008)	(0.0008)†	(0.0006)
9	0.0138	0.0091	0.0017	-0.0002	0.0020	0.0012
	(0.0034)**	(0.0034)**	(0.0017)	(0.0009)	(0.0009)*	(0.0007) *
10	0.0136	0.0088	0.0008	-0.0000	0.0031	0.0021
	(0.0039)**	(0.0039)*	(0.0020)	(0.0011)	(0.0015)*	(0.0010) *
Top 10% -	0.0017	0.0017	0.0019	0.0027	0.0023	0.0022
Bottom 10%	(0.0009)*	(0.0009)*	(0.0009)*	(0.0009)**	(0.0009)**	(0.0009)**
Top 20% -	0.0012	0.0012	0.0013	0.0019	0.0014	0.0014
Bottom 20%	(0.0007)*	(0.0007)*	(0.0007)*	(0.0007)**	(0.0007)*	(0.0007) *

Table 10: Robustness Check using Diagonal Variance-Covariance Matrix: Equal-Weight BDI Decile Stock Portfolios

BDI and BAI are computed using a diagonal variance-covariance matrix. In this approach, the diagonal terms are estimated using the sample variances of the corresponding stock returns of the last eight quarters. One month into each quarter we sort stocks into deciles based on the difference in beliefs between the "informed" fund-managers (those above the 70th percentile according to our BAI measure) and the rest using the previous quarter's holding reports. Based on these (BDI) deciles we create value-weight portfolios that we buy at the start of the second month of the quarter. We hold the portfolio for 3 months, repeating the procedure in the subsequent quarter. For each decile, we present the monthly average returns, the excess returns (over the risk free rate), the CAPM, Fama-French, and Carhart alphas, as well as the characteristic selectivity (CS) measure of Daniel, Grinblatt, Titman, and Wermers (1997). We also provide results for long-short portfolios constructed by buying the top one (two, five) decile(s) and selling the bottom one (two, five) decile(s). Decile 1 (Decile 10) contains the lowest (highest) BDI-ranked stocks.

Decile	Average Return	Excess Return	CAPM Alpha	Fama-French Alpha	Carhart Alpha	DGTW Excess
1	0.0105	0.0057	-0.0040	-0.0029	-0.0010	-0.0006
	(0.0044)**	(0.0044)†	(0.0020)*	(0.0013)**	(0.0014)	(0.0013)
2	0.0122	0.0074	-0.0007	-0.0016	0.0007	0.0001
	(0.0036)**	(0.0036)*	(0.0018)	(0.0012)†	(0.0013)	(0.0009)
3	0.0134	0.0086	0.0014	-0.0006	0.0020	0.0011
	(0.0032)**	(0.0032)**	(0.0017)	(0.0012)	(0.0011)*	(0.0008)†
4	0.0133	0.0086	0.0017	-0.0009	0.0015	0.0006
	(0.0031)**	(0.0031)**	(0.0016)	(0.0009)	(0.0008)*	(0.0007)
5	0.0130	0.0083	0.0017	-0.0007	0.0022	0.0008
	(0.0029)**	(0.0029)**	(0.0016)	(0.0011)	(0.0012)*	(0.0008)
6	0.0137	0.0090	0.0023	-0.0001	0.0026	0.0018
	(0.0030)**	(0.0030)**	(0.0016)†	(0.0010)	(0.0012)*	(0.0008)**
7	0.0133	0.0086	0.0019	-0.0006	0.0018	0.0013
	(0.0031)**	(0.0031)**	(0.0016)	(0.0010)	(0.0012)†	(0.0007)*
8	0.0129	0.0082	0.0012	-0.0011	0.0010	0.0011
	(0.0030)**	(0.0030)**	(0.0016)	(0.0010)	(0.0012)	(0.0007)†
9	0.0133	0.0086	0.0011	-0.0003	0.0012	0.0010
	(0.0033)**	(0.0033)**	(0.0015)	(0.0008)	(0.0010)	(0.0006)†
10	0.0131	0.0084	-0.0005	-0.0007	0.0008	0.0017
	(0.0040)**	(0.0040)*	(0.0019)	(0.0010)	(0.0013)	(0.0010)†
Top 10% -	0.0027	0.0027	0.0035	0.0022	0.0018	0.0022
Bottom 10%	(0.0014)*	(0.0014)*	(0.0015)*	(0.0014)†	(0.0015)	(0.0011)*
Top 20% -	0.0019	0.0019	0.0027	0.0018	0.0011	0.0016
Bottom 20%	(0.0012)†	(0.0012)†	(0.0013)*	(0.0012)†	(0.0013)	(0.0009)*

Figure 1: Contingency table of initial and subsequent BAI rankings.

Here, the cells of Table 4 are presented using a heat-plot. Lighter (darker) regions represent higher (lower) transition probabilities. Decile 1 (Decile 10) contains the lowest (highest) BAI-ranked funds.