THE JOURNAL OF FINANCE • VOL. LXIV, NO. 5 • OCTOBER 2009

Do Hedge Fund Managers Misreport Returns? Evidence from the Pooled Distribution

NICOLAS P.B. BOLLEN and VERONIKA K. POOL*

ABSTRACT

We find a significant discontinuity in the pooled distribution of monthly hedge fund returns: The number of small gains far exceeds the number of small losses. The discontinuity is present in live and defunct funds, and funds of all ages, suggesting that it is not caused by database biases. The discontinuity is absent in the 3 months culminating in an audit, suggesting it is not attributable to skillful loss avoidance. The discontinuity disappears when using bimonthly returns, indicating a reversal in fund performance following small gains. This result suggests that the discontinuity is caused at least in part by temporarily overstated returns.

IN JUNE 2007, two Bear Stearns hedge funds with large positions in subprime mortgage-related instruments collapsed, precipitating a crisis in credit markets. Barclays Bank, the sole investor in one of the funds, filed a lawsuit in December 2007, claiming that Bear Stearns fraudulently misrepresented the value of the fund's assets. The episode illustrates the discretion available to hedge fund managers when valuing illiquid securities in their portfolios, and the ability of managers to temporarily conceal losses.[1] While isolated cases of hedge fund fraud have surfaced in recent years, it is unclear whether purposeful misreporting is widespread.

We investigate the prevalence of misreporting in the hedge fund industry by studying a discontinuity in the distribution of monthly hedge fund returns, pooled across funds and over time. Discontinuities in distributions have been used in a variety of contexts, including earnings management and corruption in sports, where highly asymmetric incentives bracket a fixed hurdle. The histogram of returns reported to the Center for International Securities and Derivatives Markets (CISDM) database from 1994 to 2005 contains a sharp

*Nicolas Bollen is with the Owen Graduate School of Management at Vanderbilt University and Veronika K. Pool is with Indiana University. The authors thank Cliff Ball, Jim Byrne, Mila Getmansky, Craig Lewis, Bing Liang, Neil Ramsey, Jacob Sagi, Paul Schultz, Fei Xie, and seminar participants at Georgetown University, Indiana University, the University of Massachusetts, Vanderbilt University, and the 2008 WFA annual meeting for helpful discussions. Research support was provided by the Financial Markets Research Center.

[1] As Skeel and Partnoy (2007) argue, credit-sensitive securities can be so complex, and so reliant on subjective inputs, that model values are vulnerable to manipulation. Managers can also distort returns by opportunistically selecting favorable broker quotes. For a report of this activity, see "Does it all add up? Worries grow about the true value of repackaged debt" *Financial Times* (June 28, 2007, p. 11).

discontinuity at zero: The number of small gains is significantly greater than expected, whereas the number of small losses is significantly lower. One interpretation of this anomaly is that some hedge fund managers distort monthly returns to avoid reporting losses. There are many others, however, that must be ruled out before accepting the discontinuity as evidence of misreporting. Perhaps hedge fund managers skillfully avoid realizing losses through their trading acumen; perhaps hedge funds hold securities that generate option-like payoffs; or perhaps database biases result in losses being underrepresented in the observed distribution of returns. We present two types of evidence supporting the conclusion that the discontinuity is due at least in part to misreporting.

First, the hypothesis that the discontinuity is due to misreporting predicts that subsequent fund performance will weaken since overstatements must eventually be reversed. Testing for reversals in performance is difficult because a manager has the option to wait for some unspecified number of months before reversing a prior overstatement. We therefore conduct three tests. The first examines the properties of bimonthly returns under the assumption that overstatements are reversed the following month. There is no trace of a discontinuity in bimonthly returns. The second test measures the serial correlation of fund returns conditioned on the magnitude of lagged returns. Getmansky, Lo, and Makarov (2004) establish that hedge fund returns are unconditionally positively serially correlated. However, we detect negative serial correlation when lagged returns are just above zero. The third test separates funds into two groups: those that feature a discontinuity and those that do not. We find that the performance of the former deteriorates substantially out of sample. Taken together, these results indicate that the discontinuity is due to temporarily overstated returns.

Second, we eliminate a number of alternative explanations by measuring the discontinuity in carefully constructed subsamples of the data. The discontinuity is absent on audit dates and the 2 months leading up to them, a result consistent only with the misreporting explanation. The discontinuity is also absent in the returns of Commodity Trading Advisors (hereafter CTAs) and in the returns of factors commonly used to proxy for hedge fund strategies, suggesting that neither dynamic trading strategies nor nonlinearities in underlying asset returns are the cause. The discontinuity is present in both live and defunct funds, as well as subsamples formed by fund age, suggesting that it is not simply a reflection of survivorship bias or backfill bias. One common attribute of subsamples that do feature a discontinuity in the pooled distribution of returns is illiquidity in the funds' assets. For example, funds in the Distressed Securities category have the discontinuity, whereas funds in the Equity Market Neutral category do not. This is perhaps no surprise: Managers have more discretion when valuing illiquid securities, and hence the opportunity to misreport is greater.

Our results have at least three important implications. First, we estimate that approximately 10% of returns in the database we use are distorted. This suggests that misreporting returns is a widespread phenomenon. Though small distortions in returns do not directly put a fund's investors at risk, they may indicate more serious violations of a manager's fiduciary duty. Second, when returns are overstated, fund net asset values are as well, and new investors will overpay for entry to the fund. We estimate that overpayments totaled between \$1 billion and \$2 billion over the years 1994 to 2005. Third, if some managers avoid reporting losses, investors may underestimate the risk in hedge funds they are considering for investment and overestimate managerial performance. As a consequence, investors in aggregate may allocate more capital to hedge funds than is warranted.

The rest of this paper is organized as follows. Section I describes the data and presents the discontinuity in the distribution of hedge fund returns. Section II relates our study to earnings management and discusses prior research on misreported returns. Section III reviews the empirical methods employed in the analysis. Section IV presents the main results and Section V discusses their economic significance. A brief summary is provided in Section VI.

I. Data

We use the CISDM hedge fund database in our empirical analysis. The sample period is January 1994 through December 2005. The CISDM database includes live and defunct hedge funds, funds of funds, CTAs, commodity pool operators, and indices. We focus attention on the monthly returns of hedge funds since the returns of a fund of funds reflect the behavior of a number of individual hedge fund managers as well as the fund of funds manager. We also examine CTAs for comparison.

We eliminate all returns of 0.0000 and consecutive returns of 0.0001. On the one hand, these observations may reflect fund managers' attempt to avoid reporting negative returns, while on the other, they may simply represent missing data or conservative managers' choice not to change portfolio values when reliable market prices are not available. After applying these exclusionary criteria, 215,930 hedge fund return observations and 64,562 CTA return observations remain in our sample.

Figure 1 displays a histogram of the hedge fund returns. The two solid black vertical bars include observations just below and just above zero. Two features are readily apparent. First, there appears to be a sharp discontinuity in the distribution at zero. This is objectively corroborated by the bottom graph, which plots the value of a standard normal test statistic measuring whether the height of each bar is different from what would be expected, as described in Section III. The frequency of returns just below zero is significantly lower than expected, whereas the frequency of returns just above zero is significantly higher than expected. Second, the entire distribution to the left of zero appears deflated relative to the corresponding mass to the right of zero. Our challenge is to interpret these curious empirical phenomena.

Part of our analysis involves linear factor models for hedge fund returns. We collect a set of 10 factors that are used in the existing hedge fund literature to proxy for the trading strategies employed by hedge fund managers. These factors are drawn from three sources. The three Fama–French factors, that is, the excess return of the market and the returns of size and value portfolios,

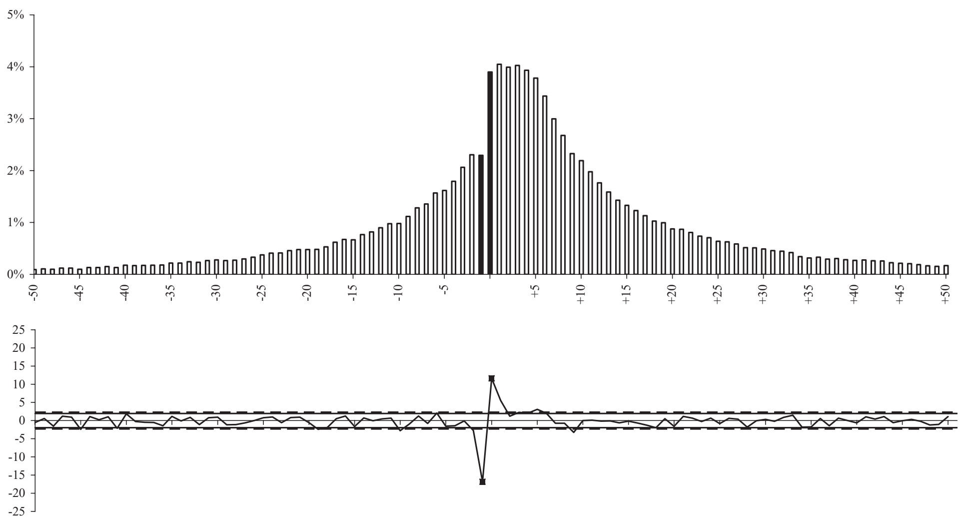

Figure 1. Full sample, 1994–2005. The figure includes monthly returns of all hedge funds in the 2005 CISDM database, $n = 215{,}930$. The top graph is a histogram of returns. Bold vertical bars indicate returns bracketing zero. Bins are 19 basis points wide with 101 bins displayed, centered on the first bin to the right of zero. The bottom graph shows the value of the following standard normal test statistic measuring the smoothness of the histogram

$$(x - np)/\sqrt{np(1-p)},$$

where x is the number of observations in the bin, n is the total number of observations, and p is the probability that a given observation resides in a bin as estimated by a kernel estimate of a smooth distribution matched to the empirical distribution. Horizontal lines indicate 95% and 99% critical values.

are from Kenneth French's website. Five trend-following factors, which are the returns of portfolios of options on bonds, foreign currencies, commodities, short-term interest rates, and stock indexes, are obtained from David Hsieh's website. The change in the yield of a 10-year Treasury note and the change in the credit spread, defined as the yield on 10-year BAA corporate bonds less the yield of a 10-year Treasury note, are obtained from the U.S. Federal Reserve's website. To estimate the factor model, we also require a risk-free rate, and for this we use the 1-month T-Bill rate from Kenneth French's website.

For robustness, we test for a discontinuity in the distribution of stock returns and equity mutual funds over the same time period. In neither case is there room for distortion: Stock returns are based on trade prices and mutual fund returns are computed from net asset values that are confirmed daily by custodians. We use NYSE/Amex and NASDAQ monthly stock returns from CRSP, eliminating those that are exactly zero, and CRSP mutual funds with Investment Company Data Inc. (ICDI) objective codes AG, GI, and LG, corresponding to Aggressive Growth, Growth and Income, and Long-Term Growth, respectively, and all domestic equity Standard & Poor's style codes.

Table I

Discontinuity by Year

Listed are the values of test statistics that measure the smoothness of histograms of monthly returns. Each year, data are pooled and a smooth distribution is estimated using a kernel density. Histograms are then created with the bin size for each year's histogram chosen following the algorithm in Silverman (1986). The number of observations in the two bins bracketing zero, labeled Negative and Positive, are then compared to the expected number given the kernel density estimate. The test statistics have independent standard normal distributions under the null hypothesis of no discontinuities in the histogram. Values in bold indicate significance at the 1% level. Mutual fund data are from CRSP. Hedge fund data are from CISDM.

	Mutual Funds		Hedge Funds	
Year	Negative	Positive	Negative	Positive
1994	−2.16	0.58	**−2.92**	**5.64**
1995	−0.81	−0.17	**−4.53**	**2.98**
1996	−1.57	−0.28	**−5.19**	0.84
1997	−1.64	0.12	**−6.90**	**4.87**
1998	−0.98	2.14	**−6.59**	**9.24**
1999	−1.66	−0.50	**−6.41**	**7.21**
2000	1.03	0.11	**−6.67**	**4.15**
2001	−1.61	1.33	**−6.11**	**10.29**
2002	2.34	1.56	**−4.94**	**8.11**
2003	**−4.25**	**−3.87**	**−5.49**	**8.07**
2004	**−4.40**	−1.10	**−5.48**	**4.30**
2005	−0.94	−0.24	**−4.56**	**3.57**

The distribution of stock returns may reflect the effects of market microstructure. For this reason, we construct histograms of stock returns over three subperiods defined by tick size: January 1995 through May 1997 (8^{ths}), July 1997 through August 2000 (16^{ths}), and April 2001 through December 2005 (decimals).[2] The two gaps between the three subperiods are excluded to avoid mixing data from different tick size regimes. In both the NASDAQ and the NYSE/Amex samples, the frequency of observing returns in the two bins around zero is significantly lower than surrounding bins. This is likely due to trades occurring at the bid price at the beginning or end of the month and the ask price on the other date. Indeed, as the tick size shrinks, and bid-ask spreads narrow, the discontinuity in the distribution becomes insignificant. Even when present, the discontinuity in the stock return distributions is qualitatively different from that of hedge funds in Figure 1, since returns in both bins bracketing zero appear underrepresented.

Table I lists the value of the test statistics for the two bins bracketing zero, labeled "Negative" and "Positive," respectively, when the discontinuity is measured on yearly subsamples of the hedge fund and mutual fund data. For the

[2] Results of this analysis, and all others described in the text but not reported in tables or figures, are available in an Internet Appendix at the Supplements and Data Set Section of the Journal of Finance website, http://www.afajof.org/supplements.asp.

hedge funds, in all years, returns in the bin below zero are statistically significantly underrepresented. In all years except 1996, returns in the bin above zero are statistically significantly overrepresented. Thus, the discontinuity is a pervasive phenomenon in the pooled distribution of hedge fund returns. For the mutual funds, in only 2003 and 2004 are the number of returns in the bins bracketing zero significantly different from what would be expected, and in no case is there a discrete increase at zero. These results indicate that the discontinuity is not a feature of mutual fund returns.

II. The Setting

This section motivates our analysis. Subsection A links our study to the accounting literature on earnings management. Subsection B contrasts our work to prior research on misreported returns.

A. Earnings Management

The discretion available to hedge fund managers when valuing illiquid securities and computing portfolio returns is similar to the choices faced by CFOs when measuring earnings. Both CFOs and hedge fund managers have strong incentives to report superior performance. According to survey evidence in Graham, Harvey, and Rajgopal (2005), CFOs state that earnings are by far the most important measure of a company's performance to external constituents. Companies with a string of increased earnings are rewarded with higher stock prices, whereas failing to meet analyst forecasts leads to a significant price reduction. CFOs benefit from equity-based compensation contracts as well as from enhancing and protecting their reputation as successful managers. Similarly, hedge fund managers benefit from reporting high returns: Standard compensation contracts include incentive fees equal to 20% of fund profits. Furthermore, a positive relationship between fund returns and investor capital flows may put pressure on managers to avoid reporting monthly losses. Outflows lower managerial compensation because a smaller asset base leads to smaller management fees and incentive fees. Also, satisfying redemptions may require a manager to close losing positions before arbitrage profits are realized, as in the model of Shleifer and Vishny (1997), resulting in lower returns and hence lower incentive fees.

Prior studies report strong evidence of a positive relation between mutual fund performance and the subsequent flow of investor capital, which Berk and Green (2004) interpret as a rational response to updated beliefs about managerial ability.[3] But what about an incremental sensitivity to the number of prior monthly losses, even after controlling for cumulative returns? Anecdotal evidence suggests that zero is a powerful quantitative anchor. Waring and Siegel (2006), for example, argue that many institutional investors pursue "absolute

[3] See, for example, Chevalier and Ellison (1997), Sirri and Tufano (1998), Busse (2001), and Del Guercio and Tkac (2002).

return" strategies based partly on the desire to consistently achieve positive returns in any market environment. Also, the "cockroach theory" implies that investors will overreact to the slightest bit of bad news, such as a negative monthly hedge fund return because they fear that more bad news lurks.

To examine whether there is an incremental sensitivity of investor fund flow to reported losses, we run the following regression:

$$F_{i,t} = \alpha_1 + \alpha_2 Y_{i,t-1} + (\beta_1 + \beta_2 Y_{i,t-1}) R_{i,t-1} + (\beta_3 + \beta_4 Y_{i,t-1}) NPOS_{i,t-1} + \varepsilon_{i,t}, \tag{1}$$

where $F_{i,t}$ is the percentage fund flow for fund i in year t, $Y_{i,t-1}$ is an indicator variable that equals 1 if fund i was age 3 years or less in year $t-1$ and 0 otherwise, $R_{i,t-1}$ is the cumulative annual return, and $NPOS_{i,t-1}$ is the number of months with positive returns. We allow for a shift in parameters for young funds to allow for differential sensitivity to lagged performance, as found in Chevalier and Ellison (1997).[4] We estimate fund flow as the difference between successive observations of total net assets, adjusted for returns, as follows:

$$F_{i,t} = [TNA_{i,t} - TNA_{i,t-1}(1 + R_{i,t})]/TNA_{i,t-1}, \tag{2}$$

where $TNA_{i,t}$ is the total net assets of fund i in year t.[5]

Table II presents the results. In panel A, the coefficient on lagged returns is 0.1506, meaning that a 1% increase in return leads to a 15 basis point incremental fund inflow. The estimate is significant at the 5% level using heteroskedasticity-consistent standard errors. The estimate for β_2 is not significant, indicating that there is no incremental response to lagged returns for young funds. The estimates for β_3 and β_4 are 0.0665 and 0.0540, respectively, meaning that for every month fund returns are positive, inflows increase by 6.65% the following year for funds greater than 3 years old and 12.05% for funds less than or equal to 3 years old. Both estimates are significant at the 1% level. These results show that fund flow is strongly related to the number of positive returns, especially for younger funds.[6]

Panel B of Table II presents the results for the case in which $NPOS_{i,t-1}$ is replaced by $NSP_{i,t-1}$, the number of months with returns greater than or equal to the S&P 500 return. The estimates for β_3 and β_4 are 0.0805 and 0.0311, respectively, both significant at the 1% level. This suggests that investors may use other benchmarks when deciding whether to invest in or withdraw from a hedge fund. Note, though, that the coefficients on lagged returns are no longer significant in panel B, suggesting that the S&P 500 benchmark is substituting for the magnitude of lagged returns, whereas in panel A the zero benchmark

[4] In (1), heightened sensitivity to young funds is expected if investors use recent performance to update prior beliefs about managerial ability, since investors presumably have more diffuse priors about younger funds and hence put more weight on recent performance.

[5] The computation in (2) assumes fund flow occurs at the end of the year. Results are robust to the assumption that all activity occurs at the beginning of the year.

[6] Agarwal, Daniel, and Naik (2007a) find similar results but do not allow sensitivity to differ by fund age.

Table II

Flow-Performance Relation

Panel A lists the results of the regression

$$F_{i,t} = \alpha_1 + \alpha_2 Y_{i,t-1} + (\beta_1 + \beta_2 Y_{i,t-1})R_{i,t-1} + (\beta_3 + \beta_4 Y_{i,t-1})NPOS_{i,t-1} + \varepsilon_{i,t},$$

where $F_{i,t}$ is the percentage fund flow for fund i in year t, $Y_{i,t-1}$ is an indicator variable that equals 1 if fund i was age 3 years or less in year $t-1$ and 0 otherwise, $R_{i,t-1}$ is cumulative annual return, and $NPOS_{i,t-1}$ is the number of months with positive returns. Panel B lists the results of

$$F_{i,t} = \alpha_1 + \alpha_2 Y_{i,t-1} + (\beta_1 + \beta_2 Y_{i,t-1})R_{i,t-1} + (\beta_3 + \beta_4 Y_{i,t-1})NSP_{i,t-1} + \varepsilon_{i,t},$$

where $NSP_{i,t-1}$ is the number of months with returns greater than or equal to the S&P 500 return. Panel C lists the results of

$$F_{i,t} = \alpha_1 + \alpha_2 Y_{i,t-1} + (\beta_1 + \beta_2 Y_{i,t-1})R_{i,t-1} + (\beta_3 + \beta_4 Y_{i,t-1})NPOS_{i,t-1} + (\beta_5 + \beta_6 Y_{i,t-1})NSP_{i,t-1} + \varepsilon_{i,t}.$$

Superscripts * and ** indicate significance at the 5% and 1% level, respectively. In all panels, year and strategy dummies are included, but coefficients are not reported.

Parameter	Estimate	t-Statistic
	Panel A: Number of Returns ≥ 0	
β_1	0.1506	2.0201*
β_2	−0.0306	−0.2921
β_3	0.0665	10.1794**
β_4	0.0540	4.9511**
Adjusted R^2	0.0857	
	Panel B: Number of Returns $\geq$ S&P 500	
β_1	0.0722	1.3087
β_2	0.0120	0.1430
β_3	0.0805	9.7051**
β_4	0.0311	2.7324**
Adjusted R^2	0.0761	
	Panel C: Both	
β_1	0.0241	0.4463
β_2	0.0105	0.1470
β_3	0.0524	8.3311**
β_4	0.0487	4.7794**
β_5	0.0525	6.2865**
β_6	0.0213	1.9541
Adjusted R^2	0.0918	

appears to be used in conjunction with lagged returns. Furthermore, the correlation between *NPOS* and *NSP* is 0.17, so they are providing investors with different information.

Panel C reports the results for the case in which both $NPOS_{i,t-1}$ and $NSP_{i,t-1}$ are included. Coefficients related to $NPOS_{i,t-1}$ are again significant at the 1% level; hence, it appears that investors consider the number of months with positive returns when deciding whether to invest in or withdraw from a hedge fund.

Another similarity between earnings management and hedge fund reporting is that, for the calculation of both quarterly corporate earnings and monthly hedge fund returns, one can make a distinction between acceptable discretion and fraud. While there have been dramatic cases of fraudulent corporate reporting, as in revenue fabrication at Enron and Global Crossing, there are also a variety of ways to manage earnings within the confines of GAAP. These include discretionary accruals, such as shifting revenue recognition across quarters and changing assumptions regarding pension funding, and real actions taken by the firm, such as repurchasing stock and changing discretionary spending on R&D or marketing. Similarly, there are clear examples of fraud in computing hedge fund returns. One widely covered example is the case of Bayou Management, whose founder and CFO both plead guilty to fraud. In 2003, their hedge fund's $49 million loss was reported as a $43 million gain. While this behavior is clearly illegal, more subtle return manipulation may not necessarily violate the SEC's antifraud rules.[7] For example, positions in highly illiquid securities could be justifiably valued using their purchase price, an updated model value, or one or more broker quotes.

Despite these similarities, hedge fund misreporting may be more pernicious than earnings management. While CFOs can influence the stock price by managing earnings, hedge fund managers determine the value of the fund when computing returns. This has a direct impact on investor fund flow, as studied in Section V. In addition, when CFOs use discretionary accruals to affect earnings, they document these decisions in financial statements, and thus investors can reverse their impact on earnings and determine "true" earnings. In contrast, there is much greater information asymmetry in hedge funds. Investors often are unaware of the precise holdings of a fund, let alone how positions are valued, hence it is difficult for them to assess the impact of their manager's discretion.

In the accounting literature on corporate earnings, several studies have documented evidence of a discontinuity in earnings or changes in earnings around zero, including Hayn (1995), Burgstahler and Dichev (1997), and Degeorge, Patel, and Zeckhauser (1999). Subsequent research questions whether the discontinuities are evidence of earnings management. Dechow, Richardson, and Tuna (2003) find a discontinuity in the distribution of earnings deflated by market capitalization, as do Burgstahler and Dichev (1997). However, they argue that since discretionary accruals are used to the same extent by firms with earnings just below zero as firms just above zero, it is unlikely that the discontinuity is due to earnings management. They discuss several alternative explanations including increased effort, as manifested in real actions taken by the firm, and the impact of deflation. Durtschi and Easton (2005) focus on the

[7] Though hedge funds are usually structured to satisfy exemptions of the SEC's Investment Company Act and Investment Advisers Act, they are still subject to the antifraud provisions therein, including a prohibition on engaging in "any act, practice, or course of business which is fraudulent, deceptive, or manipulative."

alternative explanations in more detail. They find, for example, that earnings per share do not possess the discontinuity, while earnings per share deflated by price do, and they argue that the deflation induces a selection bias by requiring firms to have a beginning-of-period price. The papers by Dechow et al. (2003) and Durtschi and Easton (2005) highlight the importance of exhausting all alternative explanations before concluding that the discontinuity in earnings or hedge fund returns is evidence of manipulation. Section IV contains a battery of tests to rule out alternatives in our study.

B. Smoothing Returns and Calendar Effects

Getmansky et al. (2004) show that if a hedge fund manager purposefully smoothes returns, the fund's volatility will be biased downward, the fund's Sharpe ratio will be biased upward, and fund returns will be positively serially correlated. Serial correlation is not necessarily indicative of misreporting, however, as it can also be the result of innocuous marking to model when funds are invested in illiquid securities.

Bollen and Pool (2008) conjecture that a manager would be more likely to smooth losses than gains, resulting in greater serial correlation when funds perform poorly. Cross-sectional analysis indicates that the propensity for funds to feature conditional serial correlation is positively related to proxies for the risk of capital flight. Through simulation, we find that conditional smoothing can generate a discontinuity at zero if the decision rule uses zero as the trigger point for smoothing. However, we also find that other reporting algorithms, including borrowing from the future and precautionary savings, can generate a discontinuity as well. Thus, while Bollen and Pool (2008) search for evidence of one specific reporting algorithm, we can simultaneously detect the effect of many when testing for discontinuous distributions. Furthermore, while Bollen and Pool (2008) find little evidence to suggest that conditional smoothing affects the pooled distribution of hedge fund returns, we show a clear, measurable pattern consistent with widespread misreporting.

Carhart et al. (2002) study the daily returns of equity mutual funds around quarter-ends and year-ends. They find that funds with the highest year-to-date returns tend to feature large returns on the last day of a quarter or a year, and that these returns are reversed the following day. We study reversals in hedge fund returns in Section IV. Carhart et al. present evidence that some mutual fund managers temporarily inflate the value of fund assets by adding to their positions of illiquid stocks on the last day of a quarter or a year. Buying pressure increases the trade prices, and the entire position can be re-valued upward. Next-day reversals provide convincing evidence that the year-end performance is distorted, since the impact of the trading activity on the last day of the year is temporary. Similarly, Agarwal et al. (2007a) find that average hedge fund returns are higher in December than all other months. They measure the incentive for a particular fund manager as the "delta" of his compensation contract, and find that the December pattern is more pronounced for those managers with higher incentives.

We check whether the discontinuity is related to these calendar effects by constructing a histogram and computing associated test statistics for two subsamples of the data formed using all observations of monthly returns in December and January, respectively. The discontinuity at zero is significant in both months, with similar magnitudes. In fact, the results are qualitatively identical in all months. Thus, the pattern we document is different from any year-end effect because it is apparent year-round. To the extent that fund managers use the history of their reported returns as a means of advertising their funds, they would be concerned with returns that stand out as being particularly poor regardless of when they occur during the year.

Another difference between our analysis and that of Carhart et al. (2002) and Agarwal et al. (2007a) is that they focus on the mean of the distribution of returns, whereas we study the shape of the entire distribution. Of particular interest is the frequency of returns around zero. Thus, our analysis does not rely on a factor model to compute abnormal returns, which is important because the existing literature has not come to a consensus on appropriate risk adjustments.[8]

III. Empirical Methodology

Our empirical methodology uses histograms of hedge fund returns to test whether the underlying densities possess significant discontinuities.[9] The most important parameter that determines the statistical properties of a histogram, as a density estimator of some underlying true distribution, is the bin width. When the bin is too small, a histogram can feature discontinuities when none exist in the underlying distribution due simply to sampling variability. When the bin is too large, the histogram may appear continuous, even if the true distribution features discontinuities. In selecting the bin width, we minimize the mean squared error between the true distribution and the histogram and follow Silverman (1986) by setting bin width equal to

$$\alpha \times 1.364 \min(\sigma, Q/1.340) N^{-\frac{1}{5}}, \tag{3}$$

where σ is the empirical distribution's standard deviation, Q is its interquartile range, N is the number of observations, and α is a scalar that depends on the type of underlying distribution assumed. Devroye (1997) shows through simulation that the definition in (3) is robust to alternative distributional assumptions. To proceed, we set $\alpha = 0.776$, corresponding to a normal distribution.

Figure 2 illustrates the impact of bin width by plotting the histogram of the full sample of hedge fund returns using three different bin widths. The three

[8] Work by Fung and Hsieh (2004), Mitchell and Pulvino (2001), and Agarwal and Naik (2004) all offer different risk factors that may or may not be relevant for a given individual fund.

[9] Our test is related to Abdulali's (2006) bias ratio, which compares the number of positive returns to the number of negative returns within one standard deviation of zero. The bias ratio has been implemented in filtering software produced by Riskdata Inc. to flag suspicious patterns in reported returns.

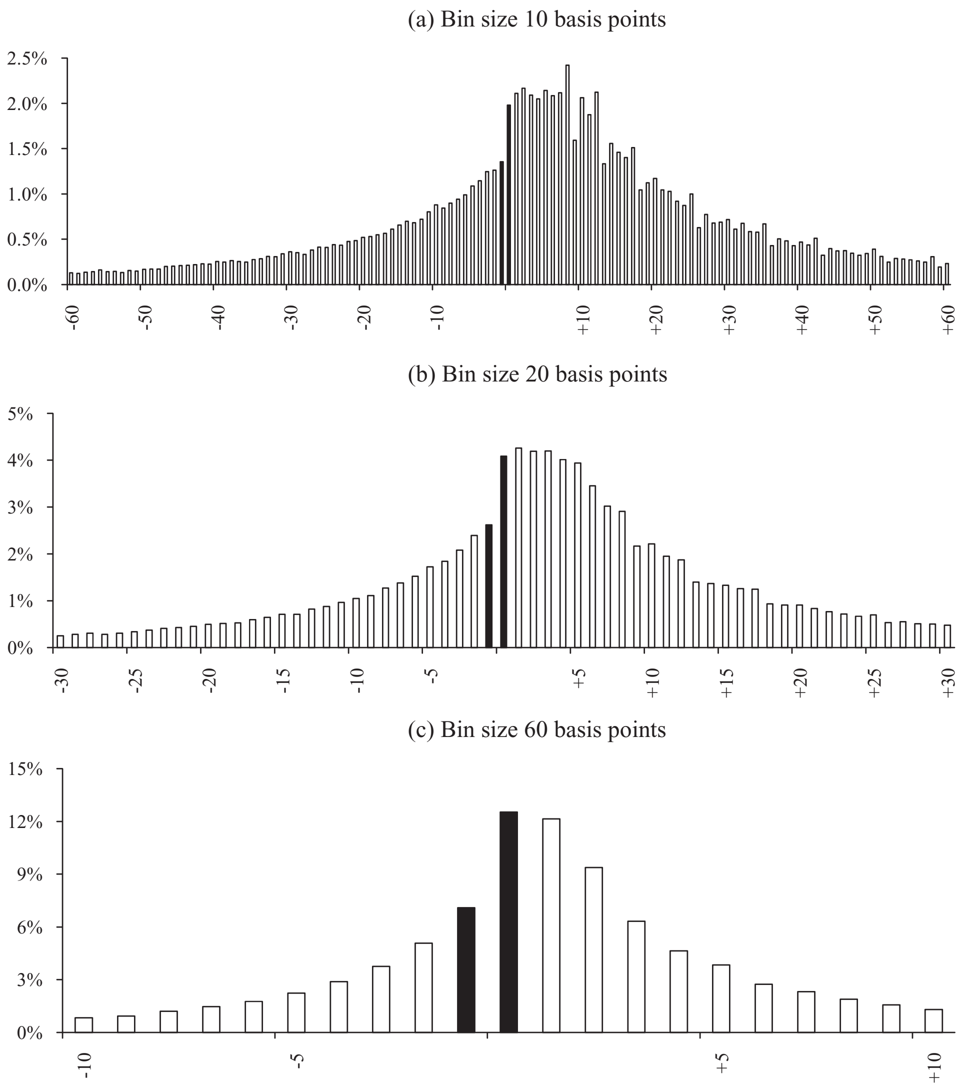

Figure 2. Bin width. The figures display histograms of monthly returns for all hedge funds in the CISDM database, $n = 215{,}930$. Bold bars indicate bins that bracket zero. The number of bins in each figure is set so that the coverage is -6% to $+6\%$.

figures cover a range of -6% to $+6\%$. Bold bars indicate bins bracketing zero. From (3), the optimal bin width for the full sample is 19 basis points. Figure 2A uses a bin width of 10 basis points, resulting in an erratic pattern for positive returns. Figure 2B uses a bin width of 20 basis points, and the discontinuity at zero is clear. In contrast, Figure 2C uses a bin width of 60 basis points, and much of the general shape of the distribution is lost.

While the histogram provides a visual evaluation of our hypothesis, we require a rigorous statistical test for a discontinuity in the distribution. Burgstahler and Dichev (1997) and Degeorge et al. (1999) determine whether the height of a particular bin differs from the height of neighboring bins. While this approach is appealing because it is related to the idea of taking numerical derivatives, it only uses a small subset of the data and it makes a potentially restrictive assumption regarding the shape of the underlying density under the null hypothesis, namely, that it is approximately linear in the neighborhood of the bin in question. We propose a new approach that uses every observation to estimate a flexible smooth density, fitted to the data, to establish the null hypothesis.[10]

The first step in our test design is to identify a smooth distribution that captures the salient features of the empirical distribution, save for any discontinuities in the density function. The smooth distribution serves as a reference, and we use it to estimate the expected number of observations in each bin of the histogram, and the corresponding standard deviation, under the null hypothesis that no discontinuity exists. It is crucial in our approach that the reference distribution fits the empirical distribution well. While this is difficult to achieve with a parametric distribution, we fit nonparametric kernel densities using a Gaussian kernel. The resulting density estimate at a point t is defined as

$$\hat{f}(t;h) = \frac{1}{Nh}\sum_{i=1}^{N}\phi\left(\frac{x_i - t}{h}\right), \tag{4}$$

where h is the bandwidth of the kernel, N is the number of observations, x is the data, and ϕ is the standard normal density. The choice of bandwidth is analogous to the choice of bin width for a histogram. Since we use the Gaussian kernel to construct the reference distribution, and we assume that the underlying distribution is Gaussian when choosing the optimal bin width, the optimal bandwidth is identical to the optimal bin width.

The second step in our test design uses an estimate of sampling variation in the histogram to determine whether the actual number of observations in a given bin is significantly different from what would be expected under the null hypothesis of a smooth underlying distribution. We estimate sampling variation in two ways.

First, we integrate the kernel density along the boundary of each bin to compute the probability that an observation will reside in it. Let p denote this probability and N the total number of observations in a sample. The DeMoivre–Laplace theorem states that the actual number of observations that will reside in the bin is asymptotically normally distributed with mean Np and standard deviation $\sqrt{Np(1-p)}$. We use this distribution to compute the value of a standard normal test statistic to determine whether the actual number of observations in a bin is significantly different from what would be expected.

[10] In unreported diagnostic tests, we find that our approach yields a superior combination of size and power relative to Burgstahler and Dichev (1997).

Second, for robustness, we generate random samples from the fitted kernel density, construct corresponding histograms, and then determine the range of bin heights to establish the cutoff levels for rejecting the null hypothesis. We adopt the algorithm summarized in Hörmann and Leydold (2000). The algorithm draws from the original sample with replacement and adds noise to the resampled data. This smoothed bootstrap creates random variates that are centered on sample data points, hence mimicking the idea of the kernel density method. See the Appendix for details. We generate 1,000 simulated samples in this manner, each of length equal to the actual sample. For each of the simulated samples, we record the number of observations that fall in each bin, and then compute the average and standard deviation across the simulations. These are then used to construct critical values as described above. In all cases, the two approaches result in critical values that are extremely similar and provide qualitatively identical inference.

To illustrate, Figure 3 displays the histogram of our sample of monthly hedge fund returns, with tails omitted to focus attention on returns near zero. The discontinuity is clear. Three curves are plotted, which indicate the frequency of observing returns in each bin according to the kernel estimate of the underlying density, as well as the upper and lower 99% confidence bands generated from 1,000 simulations. The two bins bracketing zero breach the confidence bands by a wide margin.

IV. Results

As noted before, Figure 1 reveals an intriguing feature of the pooled distribution of hedge fund returns. The number of returns in the bin just below zero is significantly less than expected, whereas the number of returns in the bin just above zero is significantly more than expected. One explanation for this result is that some hedge fund managers purposely avoid reporting losses. Subsection A tests this explanation by measuring reversals in fund returns: Misreporting predicts lower subsequent performance whereas the presence of skill predicts performance persistence. Subsection B tests whether option-like payoffs from underlying securities or dynamic trading can explain the discontinuity. Subsection C examines alternative explanations based on database biases.

A. *Skill*

Perhaps the discontinuity in the pooled distribution of hedge fund returns is due to managerial skill in actually avoiding losses. Ackermann, McEnally, and Ravenscraft (1999) find that fund Sharpe ratios are positively related to the size of incentive fees, and suggest that the level of incentive stimulates the level of managerial effort and improved performance as in the model of Starks (1987). Similarly, Agarwal, Daniel, and Naik (2007b) present evidence that fund managers with the greatest incentives, as measured carefully by the incremental dollar fee per incremental percentage return, report higher returns. The classic method of distinguishing luck from skill is to measure

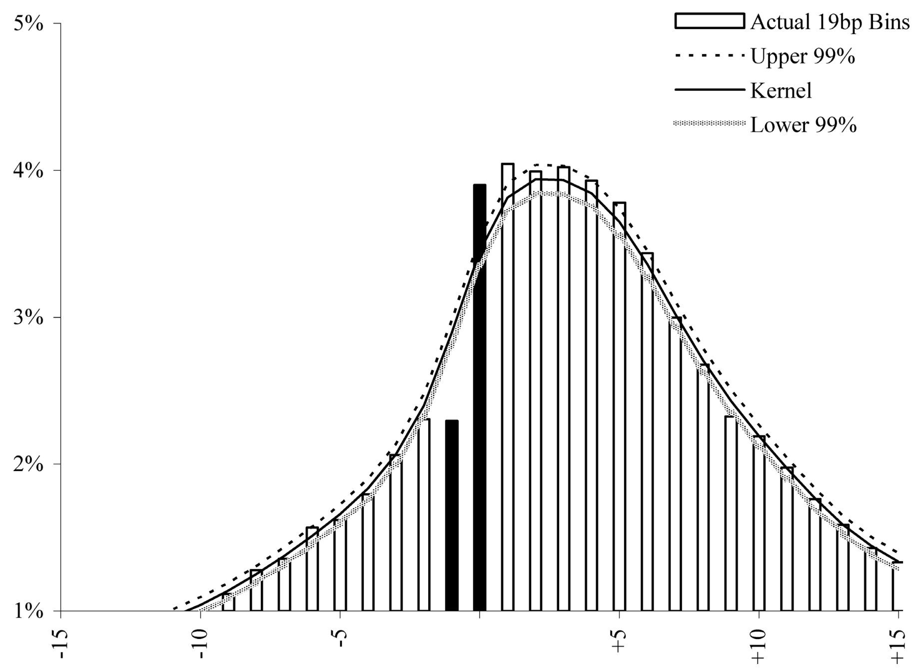

Figure 3. Kernel density. The figure displays a portion of the histogram of monthly returns for all hedge funds in the 2005 CISDM database, n = 215,930. Bold vertical bars indicate returns bracketing zero. Bins are 19 basis points wide. Also depicted is a kernel estimate of a smooth distribution matched to the empirical distribution, as well as 99% confidence bands of simulated distributions based on the kernel estimate.

persistence in performance.[11] The evidence of persistence in Brown, Goetzmann, and Ibbotson (1999), Agarwal and Naik (2000), and Avramov et al. (2007) is mixed, suggesting an alternative explanation for the link between incentives and subsequent performance: Some managers might distort returns to achieve their goals.

One way to test whether skill at avoiding losses can generate the discontinuity is to look for performance reversals. If managers create the discontinuity through their skill, then we might expect some degree of performance persistence. If instead some managers create the discontinuity through purposeful misreporting, then we might expect performance to deteriorate over time as inflated assets must at some point be written down to their true value. A difficulty in testing this hypothesis is that we have no way of knowing how long an

[11] Barras, Scaillet, and Wermers (2008) develop a new approach to separating luck from skill that does not require checking persistence. They achieve this by using a flexible null hypothesis that, unlike most prior papers, does not assume that all funds have zero alphas. Specifically, the proportion of zero-alpha funds is estimated rather than assumed a priori, and then this estimate is used to further estimate the proportion of truly skilled (positive alpha) or truly unskilled (negative alpha) funds in various segments of the alpha distribution.

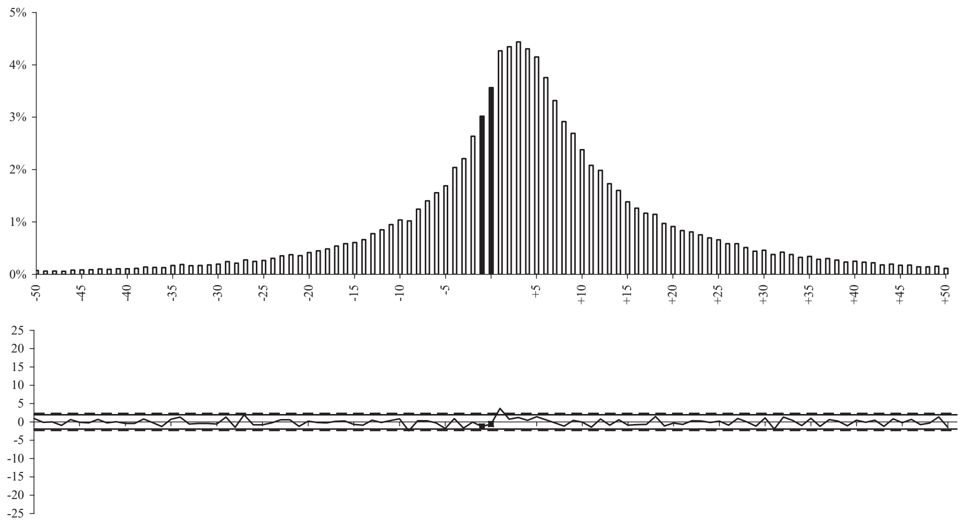

Figure 4. Bimonthly returns. The figure includes bimonthly returns of all hedge funds in the 2005 CISDM database, $n = 211{,}908$. The top graph is a histogram of returns. Bold vertical bars indicate returns bracketing zero. Bins are 36 basis points wide with 101 bins displayed, centered on the first bin to the right of zero. The bottom graph shows the value of the following standard normal test statistic measuring the smoothness of the histogram

$$x - np/\sqrt{np(1-p)},$$

where x is the number of observations in the bin, n is the total number of observations, and p is the probability that a given observation resides in a bin as estimated by a kernel estimate of a smooth distribution matched to the empirical distribution. Horizontal lines indicate 95% and 99% critical values.

overstatement is carried. Consequently, we design our tests assuming different horizons.

We assume that at least some overstatements are reversed the following month, and convert the monthly returns to 105,954 observations of nonoverlapping bimonthly returns. A small number of monthly observations are lost since we drop the last month from those funds with odd numbers of months in their history. Figure 4 shows the histogram of bimonthly returns using their optimal bin width of 36 basis points. The distribution appears quite smooth; none of the bins generate test statistics that reject the null hypothesis.[12] This result is consistent with overstatements at the monthly frequency that are reversed the following month. Reversals are inconsistent with the positive serial correlation that is a general feature of many hedge funds, especially those invested in illiquid securities, as studied in Getmansky et al. (2004). Reversals

[12] The bin width is substantially wider than the 19 basis points used for monthly returns since the number of observations is less and the dispersion of bimonthly returns is greater. In robustness tests, reported in the Internet Appendix, we find that the distribution is smooth for a wide range of bin widths, including the 19 basis points used for monthly returns.

are also inconsistent with an explanation for the discontinuity based on skillful avoidance of losses, to the extent that skill persists from month to month.

Another way to test for reversals is to measure the serial correlation of fund returns conditioned on the magnitude of lagged returns. Consider the two bins bracketing zero—let "Negative" denote the bin with returns just below zero and "Positive" the bin with returns just above zero. Positive serial correlation in hedge fund returns, documented by Getmansky et al. (2004), predicts that the subsequent returns of funds with returns in the Positive bin will be higher than the subsequent returns of funds with returns in the Negative bin. If, however, the Positive bin contains inflated returns, we would expect subsequent performance to weaken as overstatements are reversed. Two issues potentially camouflage any reversals that might occur. First, the timing of the reversals is unknown. We therefore analyze cumulative subsequent returns with horizons ranging from 1 to 6 months. Second, the magnitudes of overstatements and subsequent reversals likely vary across funds and over time. To control for heteroskedasticity, we focus on returns scaled by cross-sectional standard deviation computed each month.

We compute the mean of the cumulative, scaled returns following an observation in the Negative bin and compare to the corresponding mean following an observation in the Positive bin. For a bin size of 50 basis points, the mean subsequent return of observations in the Positive bin is lower than that of the mean subsequent return of observations in the Negative bin, significant at the 1% level, for all horizons studied.[13] The results are also clear when the subsequent returns are plotted for all bins—there is an upward slope in the relation between lagged bin on the horizontal axis and subsequent return on the vertical axis, but the relation flattens or even reverses at the Positive bin, with results varying with horizon and bin width. These results conflict with the unconditional positive serial correlation generally associated with hedge fund returns, but are consistent with reversals.

We also examine whether those funds featuring discontinuous returns exhibit less persistence than other funds over the course of their lives. In particular, we split fund histories into two equal halves, compute Sharpe ratios over each, and measure changes in Sharpe ratios. In order to identify individual funds that feature a discontinuity, we devise a fund-specific measure, labeled *Kink*, which is the percentage of observations just to the right of zero minus the percentage of observations that are just negative, computed over the first half of each fund's life.[14] The null hypothesis is that the probability of drawing a return from either bin is the same, that is, that the value of *Kink* is zero. We use a standard test of proportions that controls for the number of observations. Funds that feature a significantly positive *Kink* are labeled Offenders and other funds are labeled Non-offenders. We then compare changes in the performance of Offenders and

[13] The results weaken as the bin size is reduced. With a larger bin size, we are more likely to capture the reversals, since the magnitudes of the reversals are likely to vary.

[14] We choose a bin size of 50 basis points—the small number of observations in many funds makes computation of the optimal bin size in equation (3) infeasible.

Non-offenders to determine whether a reversal is present for the funds with a discontinuity.

Table III lists cross-sectional means of summary statistics of the Offenders and the Non-offenders. The statistics are computed over the first half *Pre* and second half *Post* of each fund's life, and the p-value corresponds to a t-test for difference in means. In panel A1, the live Non-offenders feature significant reductions in average return, standard deviation, Sharpe ratio, skewness, and excess kurtosis, whereas *Kink* shows a small but significant increase. In panel A2, the Offenders also display a reduction in average return. The row labeled

Table III

Reversals

Listed are cross-sectional means of summary statistics: average monthly return μ, standard deviation of monthly return σ, Sharpe ratio *Sharpe*, skewness *Skew*, excess kurtosis *Kurt*, and a measure of discontinuity *Kink*. *Kink* equals the percentage of a fund's observations that are between 0 and 50 basis points minus the percentage of observations that are negative and no less than −50 basis points. Statistics are computed separately over each half of each fund's life. Cross-sectional means are then taken over the first halves' *Pre* and second halves' *Post* of funds' lives. The significance of a t-test for differences in means is listed as p-value. Panel A lists the results for live funds in the CISDM database. Panel B lists the results for dead funds. Funds are further split into Offenders, for whom the percentage of observations between 0 and 50 basis points is, at the 5% level, significantly greater than the percentage of negative observations no less than −50 basis points, and Non-offenders. *Difference* indicates the significance of a t-test comparing the mean change in each fund's statistics from *Pre* to *Post* for Non-offenders versus Offenders.

No. of Funds		μ	σ	*Sharpe*	*Skew*	*Kurt*	*Kink*
			A. Live Funds				
			A1. Non-offenders				
1,258	*Pre*	0.0136	0.0426	0.3857	0.1147	1.9715	0.0234
	Post	0.0103	0.0311	0.3227	0.0171	1.2621	0.0377
	p-value	0.0000	0.0000	0.0005	0.0110	0.0000	0.0001
			A2. Offenders				
185	*Pre*	0.0101	0.0188	0.6120	−0.0798	3.6191	0.2356
	Post	0.0067	0.0142	0.3447	−0.1029	1.8805	0.1701
	p-value	0.0000	0.0200	0.0000	0.8633	0.0024	0.0016
	Difference	0.9260	0.0000	0.0000	0.5487	0.0586	0.0000
			B. Dead Funds				
			B1. Non-offenders				
1,478	*Pre*	0.0148	0.0508	0.2998	0.0576	1.7801	0.0223
	Post	0.0038	0.0546	0.0416	−0.1083	1.8605	0.0426
	p-value	0.0000	0.0313	0.0000	0.0000	0.5305	0.0000
			B2. Offenders				
144	*Pre*	0.0133	0.0386	0.3322	0.2287	3.6085	0.2441
	Post	0.0036	0.0237	−0.1206	−0.1657	3.0947	0.1285
	p-value	0.0081	0.3177	0.0001	0.0183	0.3531	0.0000
	Difference	0.7006	0.2138	0.0573	0.1230	0.2771	0.0000

Difference shows the p-value of a t-test for difference in means between the changes of the Non-offenders and the changes of the Offenders. The p-value is 0.9260, indicating the average reduction in mean is the same for the two groups of funds. The Offenders also feature a significant reduction in volatility, from 1.88% monthly to 1.42%. In both *Pre* and *Post*, the Offenders have a much lower volatility than the Non-offenders, which perhaps is no surprise since a significant *Kink* requires a substantial number of observations close to zero. Note, though, that the Non-offenders feature a significantly larger reduction, from 4.26% to 3.11%. Consequently, the Sharpe ratio of the Offenders falls far more, from 0.6120 in *Pre* to 0.3447 in *Post*, than the drop for Non-offenders, 0.3857 to 0.3227. This result indicates that Offenders do exhibit an economically and statistically significant reversal of performance.

Two other points merit a mention. First, note that the Offenders have negative skewness, on average, in both *Pre* and *Post*, whereas the average for Non-offenders is positive. Furthermore, in both *Pre* and *Post*, the Offenders have higher excess kurtosis. Taken together, these results are consistent with Offenders that report small positive returns when they actually suffer losses. Overstatements subsequently must be reversed, which could result in larger losses if the reversal occurs during a month with a loss or in smaller gains if the reversal occurs during a month with a gain. Second, note that *Kink* is reduced from *Pre* to *Post* for the Offenders, from 23.56% to 17.01%. This result suggests that a reporting algorithm of rounding up small losses to small gains is sustainable to some degree, but that the discontinuity shrinks, perhaps as a result of some managers abandoning the practice.

For the dead funds in panel B, the results are qualitatively similar. Naturally, average returns and Sharpe ratios fall far more for the dead funds. The Sharpe ratios of Non-offenders and Offenders both drop substantially, and, though the Offenders drop more, the average reductions are much closer than for the live funds.

The extremely low volatility and high Sharpe ratios of live Offenders in the *Pre* period could be interpreted as evidence that *Kink* is simply a way of identifying fund managers that artificially smooth returns. Getmansky et al. (2004) show that smoothing results in a downward bias of volatility and an upward bias of the Sharpe ratio. To determine whether *Kink* offers incremental information, we first measure the degree of smoothness in fund returns by regressing returns on their lag during the *Pre* period. Then we estimate the relation between Sharpe ratios measured in the *Pre* and *Post* periods in the following regression:

$$Sharpe_{i,Post} = \alpha + \beta_0 Sharpe_{i,Pre} + \beta_1 I_{i,Pre} Sharpe_{i,Pre} + \varepsilon_i, \tag{5}$$

where I is an indicator that equals 1 if the fund is an Offender, as defined using *Kink*, and 0 otherwise. Next, we redefine the indicator variable to equal 1 if the fund has an AR(1) coefficient that is positive and significant at the 5% level and rerun the regression.

The results are listed in Table IV. Panel A uses an indicator defined by *Kink*. For live funds, Sharpe ratios are somewhat persistent, with a β_0 coefficient of

Table IV

Persistence

Listed are results of cross-sectional regressions of hedge fund Sharpe ratios on their lags. Each fund's history is split into halves and Sharpe ratios are computed over each subperiod. Regressions are of the form

$$Sharpe_{i,Post} = \alpha + \beta_0 Sharpe_{i,Pre} + \beta_1 I_{i,Pre} Sharpe_{i,Pre} + \varepsilon_i\,,$$

where I is an indicator variable. In Panel A, $I = 1$ if the fund's percentage of observations between 0 and 50 basis points is significantly greater, at the 5% level, than the percentage of negative observations no less than -50 basis points. In Panel B, $I = 1$ if the fund has a positive AR(1) coefficient that is significant at the 5% level. Results are further split into live and dead funds.

			α	β_0	β_1
		A. Sorted by Kink			
		A1. Live Funds			
Adjusted R^2	29.0%	Estimate	0.1223	0.5147	−0.1306
No. of observations	1443	p-value	0.0000	0.0000	0.0008
		A2. Dead Funds			
Adjusted R^2	6.0%	Estimate	−0.0601	0.2599	0.2931
No. of observation	1622	p-value	0.0001	0.0000	0.0011
		B. Sorted by AR(1)			
		B1. Live Funds			
Adjusted R^2	28.5%	Estimate	0.1248	0.4725	0.0335
No. of observation	1443	p-value	0.0000	0.0000	0.2846
		B2. Dead Funds			
Adjusted R^2	5.5%	Estimate	−0.0570	0.2972	−0.0712
No. of observation	1622	p-value	0.0002	0.0000	0.1782

0.5147. In addition, Sharpe ratios of Offenders are significantly less persistent, with a β_1 coefficient of -0.1306. Thus, after controlling for the initial Sharpe ratio, funds with a discontinuity in the first half of their lives feature lower Sharpe ratios in the second half of their lives than other funds. For dead funds, the regression adjusted R^2 is only 6.0%, compared to 29.0% for live funds; hence the relation between Sharpe ratios in the first and second halves of funds' lives is much weaker. Panel B shows results using an indicator defined by the AR(1) coefficient of each fund, estimated over the first half of its life. The indicator variable does not provide information about the persistence of the Sharpe ratio for either live or dead funds. Hence, *Kink* provides more information than a measure of unconditional smoothing. Our results indicate that the discontinuity is not the result of skilled loss avoidance, since those funds that feature a kinked distribution subsequently suffer a substantial drop in performance.

An alternative way to test the skill-based explanation is to compare the distribution observed in the months surrounding audits to the distribution during other months. There should be no difference if the discontinuity is due to skillful

Table V

Audits and Fund Attributes

Panel A lists the average monthly return, fund size in \$000, and age for subsets of funds grouped by whether the auditor or last audit date is listed in the CISDM database (Audited) or not (Non-audited). A fund's size is the assets under management averaged over the fund's observations in the database. A fund's age is the age of the fund from inception averaged over the fund's observations in the database. Also listed are t-statistics for a difference in means. Significance is assessed using Satterthwaite's correction. Superscripts * and ** indicate significance at the 5% and 1% level, respectively. Panel B lists coefficient estimates and significance levels from a probit regression in which the dependent variable takes a value of 1 if the fund has an auditor or last audit date listed.

Panel A: Univariate Statistics

Attribute	Audited	Non-audited	t-Statistic
Monthly return	1.06%	0.74%	5.3643**
Size (in \$000)	192,922	95,516	1.2839
Age	5.04	2.47	26.1894**

Panel B: Probit Regression Coefficients

Parameter	Estimate	z-Statistic
Constant	−1.1419	−4.4752**
Active	0.3557	7.2033**
Age	0.1636	14.8648**
Ln(size)	0.0485	3.1606**
Onshore	0.1077	2.3485*
Monthly return	4.9241	3.3834**
McFadden R^2	0.1353	

avoidance of losses, since there is no reason to expect a fund manager's ability to diminish on or near audit dates.[15] Of the 4,286 hedge funds, 2,213 have an audit date listed in the CISDM database. Those funds with no audit date listed likely consist of two groups of funds—those that have been audited but for which no information was provided to the database and those that have not been audited. Our conjecture is that, taken as a group, the funds with no audit date listed have less oversight than the funds with an audit date listed.

To determine whether the two groups of funds have other systematic differences, Table V lists in panel A the average monthly return, fund size (in \$000), and age of the fund for the two groups. A fund's size is measured as the average assets under management over a fund's history in the database. Similarly, a fund's age is the average time between fund inception and the observations of fund returns. Audited funds have higher monthly return, size, and age than nonaudited funds. Though the audited funds are on average twice as large, approximately \$193 million compared to \$95 million, the difference is not significant given the large variation across funds. Panel B lists the results of a probit analysis to determine the relation between fund attributes and the likelihood

[15] Liang (2003) presents evidence that audited funds feature more accurate returns.

that a fund is audited. Again, older, larger, and more successful funds are more likely to be audited. In addition, funds that are domiciled in the United States (Onshore) and that are still being reported to the database as of December 2005 (Active) are more likely to be audited.

Table VI summarizes results for six different tests of alternative explanations. For each category of funds, the table displays test statistics for 11 bins centered on the bin just above zero. "Audit" lists the test statistics from the distribution of returns of funds on their last reported audit date and the prior 2 months. There is no significant discontinuity in the distribution since no test statistic is significant. "No Audit" lists the test statistics from the distribution of a matched sample of observations. The sample is formed by first identifying all unique combinations of date, strategy, and size quintile reflected in the data in the Audit subsample. Then, all observations from funds that do not have any audit dates listed in the CISDM database that match the date, strategy, and size quintile of one of the combinations identified in the Audit subsample are collected. A significant discontinuity exists at zero for this matched sample of funds with no audit dates listed, with a test statistic at the bin left of zero of −4.07, and a test statistic at the bin right of zero of 3.86, both significant at the 1% level. This result suggests that the discontinuity disappears when funds are close to an audit date, which does not seem consistent with a skill-based explanation for the shape of the pooled distribution of hedge fund returns.

B. Nonlinearities in Underlying Assets and Strategies

Perhaps the discontinuity is due to nonlinearities in returns generated either by the assets in which hedge funds invest or the dynamic trading strategies that hedge fund managers employ.[16] We examine the pooled distribution of returns of asset-based style factors that mimic some of the dynamic strategies employed by some fund managers, and are constructed directly from market prices of underlying assets. These distributions are therefore free of distortions. If the strategies inherently feature discontinuous distributions, then these asset-based style factors should feature distributions similar to those of the hedge funds. For each fund with at least 24 contiguous returns, we regress excess fund returns on a subset of the 10 factors described in Section I. The subset has a maximum of three factors and is selected to minimize the Bayesian Information Criterion. We then construct fitted returns by adding the risk-free rate to the sum product of factor loadings and factor returns.

Table VI lists the test statistics for the distribution of fitted and raw returns from the 3,067 funds included in this part of the analysis. The fitted return distribution is smooth: No bin rejects the null hypothesis that the observed distribution conforms to the kernel density estimate. In contrast, the raw return distribution has test statistics of −15.27 and 15.02 for the two bins bracketing zero. To the extent that the underlying risk factors, chosen from prior

[16] Fung and Hsieh (1997) argue that hedge fund returns feature option-like payoffs relative to the return of underlying assets, consistent with dynamic trading strategies.

Table VI
Alternative Explanations

Listed in each column are standard normal test statistics that determine whether a histogram of returns is consistent with an estimated underlying smooth distribution. Values in bold indicate significance at the 1% level. Statistics for 11 bins are reported, where bin Zero is the first with positive observations. "Audit" consists of monthly returns for funds on their audit date and the 2 months leading up to it. "No Audit" consists of a matched set of returns for funds without an audit date listed and that match the size, strategy, and date of the "Audit" sample. "Fitted" consists of fitted returns from an optimal factor model estimated for each fund with at least 24 observations. "Raw" are the corresponding raw returns. "Non-put Sellers" consists of returns of funds with nonnegative exposure to the payoffs of out-of-the-money index put options. "Put Sellers" consists of returns of funds with negative exposure to the payoffs of out-of-the-money index put options. "Equity Neutral" consists of returns of funds labeled as Equity Market Neutral funds in the CISDM database. "Distressed Securities" consists of returns of funds labeled as such in the CISDM database. "Low Smooth" consists of returns of funds in the lowest quartile as ranked by their Getmansky et al. (2004) smoothing coefficient. "High Smooth" consists of funds in the highest quartile. "Live" consists of returns of live funds according to the CISDM database. "Dead" consists of returns of defunct funds. The bin width for each histogram is determined separately as

$$1.058\min(\sigma, Q/1.340)N^{-\frac{1}{5}},$$

where σ is the sample's standard deviation, Q is the interquartile range, and N is the number of observations. The underlying smooth distribution is a kernel density estimate using a Gaussian kernel. The test statistic compares the number of observations in each bin to the expected number given the underlying smooth distribution.

Bin	Audit	No Audit	Fitted	Raw	Non-put Sellers	Put Sellers	Equity Neutral	Distressed Securities	Low Smooth	High Smooth	Live	Dead
−5	−1.05	0.18	0.38	−1.74	0.75	−0.10	0.19	0.66	−1.22	−0.98	−0.26	−1.25
−4	−0.20	−2.19	−0.89	−2.51	**−2.99**	0.46	−0.74	1.62	−1.41	−1.36	−0.57	−0.11
−3	−0.73	−1.14	0.01	1.48	1.60	−1.46	−1.22	−1.18	−0.45	0.13	**−2.97**	−2.55
−2	−1.09	0.02	−0.07	**−3.81**	**−3.87**	−1.92	0.22	−1.67	−0.73	−1.64	−0.12	**−2.85**
−1	−1.32	**−4.07**	−0.59	**−15.27**	**−12.24**	**−10.58**	−0.09	**−5.00**	**−2.71**	**−8.33**	**−8.61**	**−10.35**
Zero	0.33	**3.86**	0.12	**15.02**	**13.84**	**9.13**	**3.62**	**3.78**	**6.73**	**7.11**	**8.01**	**15.03**
+1	2.26	1.82	0.73	2.09	−0.30	2.24	−1.24	0.07	0.52	1.74	**5.36**	2.11
+2	1.51	1.42	0.58	**3.05**	1.33	0.54	−0.50	1.38	2.04	−0.13	−1.68	−0.57
+3	1.60	0.87	1.85	−0.12	0.69	0.34	2.07	1.18	1.17	0.64	**3.77**	2.11
+4	−1.22	1.21	0.35	−0.70	2.07	1.02	0.14	**2.77**	−0.24	**3.17**	0.31	1.54
+5	−0.67	−0.62	1.01	**5.19**	1.68	0.47	1.48	−0.78	0.02	0.59	**3.77**	0.13

research in modeling hedge fund returns, mimic important hedge fund trading strategies, these results suggest that nonlinearities therein cannot explain the discontinuity.[17]

Positions in options generate asymmetric payoffs; perhaps they can generate the discontinuity in fund returns. The payoffs to writing out-of-the-money (OTM) index put options, for example, would supply many small positive returns and rare large negatives, which together are consistent with the discontinuity we observe. To test this conjecture, we regress excess fund returns on the four series created by Agarwal and Naik (2004) to mimic positions in puts and calls on the S&P 500. We split funds into two categories: those with negative coefficients on OTM puts, which we refer to as "Put Sellers," and all other funds, labeled "Non-put Sellers." Then we examine the histogram of these two categories.

In Table VI, both Put Sellers and Non-put Sellers feature the discontinuity: The test statistics are -12.24 and 13.84 for the two bins bracketing zero for the Non-put Sellers and -10.58 and 9.13 for the same bins for Put Sellers. These results are counter-intuitive because the discontinuity is sharper for Non-put Sellers, but this is likely caused by a substantial difference in the number of observations between the two groups. In any case, this test indicates that selling index puts does not appear to explain the discontinuity.

Another way to test an explanation based on dynamic trading is to compare the pooled distribution of reported returns from hedge funds to the pooled distribution of reported returns from CTAs. Managers in both classes of funds are free to use dynamic strategies; hence, we might expect to observe the same kind of discontinuity in both distributions. If, however, the *kink* is due to distortions in reported returns, then we might expect to see a less significant discontinuity in the CTAs. The reason is that CTA managers primarily invest in highly liquid future contracts, which implies that distorting reported returns would be more difficult.

Figure 5A shows the histogram of CTA returns. The CTAs feature a symmetric distribution, with no discontinuity below zero. The category just to the right of zero happens to be the peak of the distribution and hence does reject smoothness in our test. In contrast, the hedge fund distribution in Figure 1 displays the discontinuity below zero and appears to have a deflated left tail. These results suggest that the discontinuity in the hedge fund distribution is not the result of dynamic trading.

A likely reason why hedge fund returns feature a discontinuity at zero but CTA returns do not is that CTA managers invest in more liquid securities, in which case there is no opportunity for CTA managers to exercise discretion when computing fund returns. We examine the relationship between liquidity and discontinuities in return distributions in two other ways.

[17] The average adjusted R^2 in the factor regressions is 27.7%. While this level of fit is consistent with existing research, for example, Bollen and Whaley (2008), it suggests that commonly used factors can only partly capture the trading of hedge fund managers.

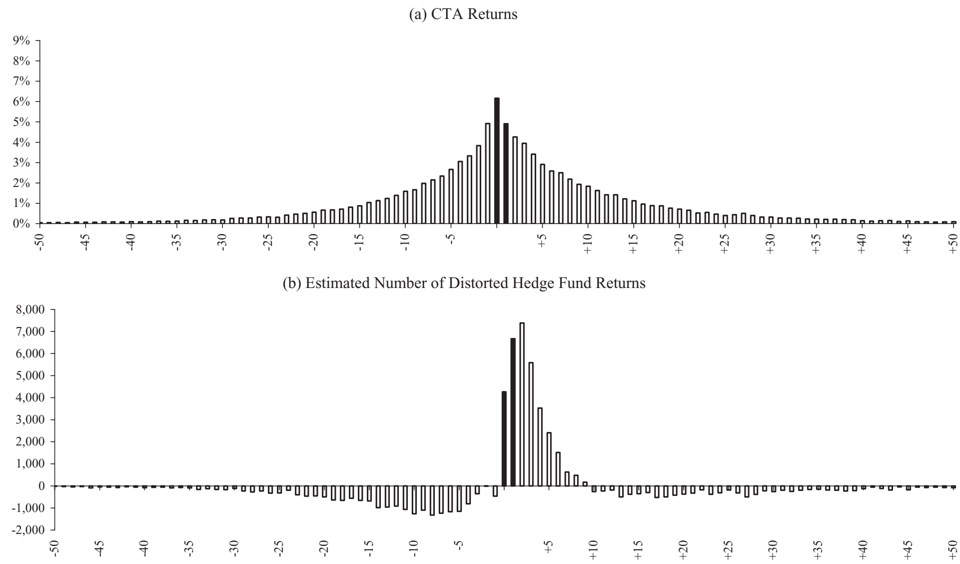

Figure 5. CTA returns. A: A histogram of CTA returns. B: Shows for each bin the estimated number of hedge fund returns that may have been distorted. Estimates are computed by comparing the percentage of returns in different bins of the CTA return distribution to the percentage of returns in the bins of the hedge fund return distribution. Data include monthly returns of all hedge funds and CTAs and in the 2005 CISDM database. $n = 215{,}930$ and $n = 64{,}562$, respectively. Bold vertical bars indicate returns bracketing zero.

First, we examine the discontinuity for subsets of funds formed by strategy. In particular, Equity Market Neutral funds invest in stocks, which are extremely liquid relative to the securities held by Distressed Securities Funds. Table VI shows that while the bin to the right of zero rejects the null for both the Distressed Securities funds and Equity Market Neutral funds, with test statistics of 3.78 and 3.62, respectively, only the Distressed Securities funds reject also at the bin to the left of zero. The test statistics there are -5.00 and -0.09 for the two groups of funds. This result is again consistent with the notion that avoidance of reporting losses is more prevalent when managers have greater discretion over valuation.

Second, we estimate the Getmansky et al. (2004) smoothing coefficient for each fund. Funds with more smoothing can be interpreted as those with more illiquid assets. To quantify smoothness in returns, Getmansky et al. estimate the following MA(2) model for demeaned returns:

$$X_t^i = \theta_0^i \eta_t^i + \theta_1^i \eta_{t-1}^i + \theta_2^i \eta_{t-2}^i, \tag{6}$$

where $X_t^i = R_t^i - \mu_i$ is the demeaned reported return for hedge fund i at time t and $\eta_t^i \sim N(0, \sigma_i^2)$. To identify the system, Getmansky et al. apply the following constraint:

$$1 = \theta_0^i + \theta_1^i + \theta_2^i. \tag{7}$$

The estimated smoothing parameter, $\hat{\theta}_0^i$, represents the fraction of current true returns reported in a given month on average, and is therefore inversely related to liquidity. We use maximum likelihood to estimate the MA(2) model for each hedge fund with at least 24 monthly returns. We then sort funds by their estimated smoothing parameter.

In Table VI, "Low Smoothing" corresponds to the pooled distribution of liquid funds, that is, the top quartile of funds, as ranked by the smoothing parameter, whereas "High Smoothing" corresponds to the illiquid funds in the bottom quartile. Both categories reject the null hypothesis of a smooth distribution at the two bins bracketing zero, although the discontinuity is much more pronounced in funds with more smoothing, that is, the funds with more illiquid securities. For the bin below zero, for example, the test statistics for the Low Smoothing funds is -2.71 versus -8.33 for the High Smoothing funds. Getmansky et al. (2004) argue that the smoothing coefficient can be interpreted as a proxy for the illiquidity of the assets in which a fund invests, and hence they have difficulty distinguishing between purposeful smoothing and innocuous smoothing resulting from marking assets to model. In contrast, since there is no apparent reason why marking to model would result in the discontinuity featured in Table VI, purposeful distortion seems to be the more likely explanation.

C. Database Biases

Perhaps the discontinuity that we have documented is simply another manifestation of well-known biases in hedge fund databases. The survivorship studied in Brown et al. (1999), for example, may result in a distribution featuring fewer observations of poor returns than expected, even if the reported returns are accurate. Ackermann et al. (1999) argue that both superior performers and inferior performers may stop reporting, and present evidence that these groups somewhat offset each other in the pooled distribution of returns. Ackermann et al. also discuss back-fill bias, in which only successful funds initiate reporting to a database and report superior historical returns. While both survivorship bias and back-fill bias might result in fewer observations of poor returns than expected, they do not necessarily imply a discrete break in the distribution at zero, but rather a shift in location of the entire distribution. Nevertheless, we test for the possible impact of survivorship bias and back-fill bias by first discarding the first 12 months of observations, eliminating the impact of back-fill, and then searching for a discontinuity in live funds and defunct funds separately. Table VI lists test statistics for the distributions of live and defunct funds separately. The discontinuity at zero is present in both distributions, indicating that neither back-fill bias nor survivorship can explain the presence of the discontinuity.

We further examine the role of back-fill by analyzing the histograms of pooled distributions of subsets formed by fund age. This allows us to investigate whether the discontinuity is only present during the beginning of a fund's reporting history, consistent with back-fill bias, or more prevalent during later years in a fund's reporting history, consistent with survivorship bias. If, however, the discontinuity is present in all subsamples stratified by fund age, then managerial distortion appears to be a more likely explanation.

Four categories are formed, corresponding to observations of all funds during their first 2 years of reporting history (age ≤ 1), their 4$^{\text{th}}$ year (age = 3), their 7$^{\text{th}}$ year (age = 6), and their 11$^{\text{th}}$ year (age = 10). In unreported tests, we find that the discontinuity exists at zero return in all categories, and in fact is larger for the observations from older funds. This result is consistent with the flow-performance results in Section II. Since investors reward funds with fewer negative returns by directing more capital to them, one would expect that funds that survive longer have a greater tendency to avoid monthly losses. The discontinuity is statistically significant in all cases. The value of the test statistic for bins bracketing zero falls in magnitude as funds age as a result of a substantial drop in the number of observations.

V. Economic Significance

The results described in Section IV indicate that hedge funds invested in illiquid securities feature a robust discontinuity in their return distribution. We have demonstrated the statistical significance of the discontinuity by comparing the actual number of observations in each bin to an expected number given by an underlying smooth density. Below we present our analysis of the economic significance of the discontinuity.

A. Number of Misreported Returns

Comparing the distribution of CTA returns depicted in Figure 5A to hedge fund returns allows us to estimate the number of observations that may be affected by distortion. Specifically, at each bin we multiply the difference between the frequencies of the two distributions by the size of the hedge fund sample. There might be differences between the CTA distribution and the unobserved "true" hedge fund distribution due to differences in the assets in which the two vehicles invest. However, we have no way of estimating a distortion-free distribution of hedge fund returns, and so require some type of reference distribution. Figure 5B plots the difference between the actual number of hedge fund observations in each bin and the expected number given the CTA distribution. Negatives to the left of zero and the sharp positive peak just to the right of zero indicate that negative returns of all magnitudes may have been reported as small, positive returns. In total, over 20,000 observations to the left of zero appear to be "missing," approximately 10% of the entire sample. Interestingly, there are negatives well to the right of zero as well, which could be due to hedge fund managers saving for a rainy day, or reversing prior overstatements.

B. Wealth Transfers from Misreported Returns

When hedge fund managers misreport returns, investor capital flow is distorted in the same way that stale net asset values distort capital flow in mutual funds. Inflows are likely more susceptible to overstated returns, since hedge fund managers, using fund redemption policies, can delay outflows until overstatements are reversed. For this reason, we estimate the aggregate wealth transfer that occurs when investors subscribe to funds with overstated returns.

We label funds with discontinuous distributions "Offenders" and identify them using a test similar to the one used in our analysis of reversals in Section IV. Using a fund's entire history, we compute the percentage of returns in two 50-basis point bins bracketing zero. Funds with significantly more small positives than small negatives are flagged as Offenders; all other funds are labeled "Non-offenders." For the Offenders, we compute the dollar wealth transfer each month by comparing the actual cost incoming shareholders pay to obtain fund shares to what the cost should be at fair value. We first measure total actual cost in month t by dollar fund inflow. Inflows are calculated as follows:

$$F_{i,t} = TNA_{i,t} - TNA_{i,t-1}(1 + R^R_{i,t}), \tag{8}$$

where $F_{i,t}$ is fund flow in month t, $TNA_{i,t}$ is total net assets, and $R^R_{i,t}$ is reported return. The number of shares purchased equals the dollar flow divided by the fund's reported net asset value per share, which we label $NAV^R_{i,t}$. The wealth transfer $W_{i,t}$ equals the number of shares purchased multiplied by the amount of distortion in the net asset value:

$$W_{i,t} = \frac{F_{i,t}}{NAV^R_{i,t}}\left[NAV^R_{i,t} - NAV^A_{i,t}\right], \tag{9}$$

where $NAV^A_{i,t}$ is our estimate of the fair net asset value. We estimate $NAV^A_{i,t}$ by grossing up the lagged reported net asset value by a proxy for the true return of the fund. Note that this calculation assumes that the prior net asset value is correct, which is consistent with the evidence of monthly reversals in Section IV.

We use two approaches to proxy for true returns. The underlying assumption in both cases is that small positive returns are reported to hide negative true returns. The first approach takes a random draw (with replacement) from the fund's own negative returns to proxy for true returns. One disadvantage of this approach is that when funds manage their returns, the entire return distribution is distorted. Moreover, we lose those funds in our sample that have no negative returns in their histories. The second approach is to draw from the negative returns of Non-offender funds. While these distributions are not distorted by assumption, a disadvantage of this approach is that Offender and Non-offender true distributions may differ. To mitigate this issue, we match Offenders to Non-offenders by strategy, whether the fund is active, age, and size.

Table VII reports the estimated aggregate wealth transfer using several different proxies for true returns. There are 3,682 observations of Offenders that experience fund inflows on months that feature small positive returns and have available fund flow and asset value observations. For each of these observations, we calculate the extra cost incoming shareholders face based on equation (9),

Table VII

Impact on Fund Flow

Listed are estimates of the wealth transfer between new and existing hedge fund investors when fund inflows occur in funds with discontinuous return distributions. Funds are split into Offenders, for which the percentage of observations between 0 and 50 basis points is, at the 5% level, significantly greater than the percentage of negative observations no less than −50 basis points, and Non-offenders. For each month that an Offender has a fund inflow and a return between 0 and 50 basis points, wealth transfer is estimated as shares purchased times the difference between the reported net asset value and a proxy for the fair net asset value. The fair net asset value equals lagged reported net asset value scaled by a proxy for true return. "Own" is a proxy using a random draw from the fund's negative returns. Other proxies are draws from the negative returns of a matched set of Non-offender funds. Listed are the different combinations of matching criteria used.

Return Proxy						
Own	Strategy	Age	Size	Active	No. of Observation	Wealth Transfer
x					3,474	424,829,545
	x	x	x	x	3,092	537,204,555
	x	x	x		3,325	623,628,681
	x	x		x	3,423	611,142,143
	x	x			3,547	917,055,102
	x		x	x	3,585	554,778,422
	x		x		3,595	670,343,299

and report the sum in the table. We use several subsets of the matching variables. Using the own return proxy, the total loss is estimated at 425 million dollars. For the Non-offender proxies, the range is between 537 million and 917 million dollars. Since the CISDM database has approximately 40% coverage of the hedge fund industry, this result implies an aggregate wealth transfer between 1 billion and 2 billion dollars.[18]

VI. Conclusions

This paper documents a robust feature of the pooled cross-sectional, time-series distribution of hedge fund returns: A discontinuity exists at zero. The discontinuity is not present during the 3 months culminating in an audit, in factor returns commonly used to proxy for trading strategies employed by funds, or in subsets of funds that invest in liquid securities. These results suggest that some managers distort returns when possible, for example, when fund returns are at their discretion and when their reported returns are not closely monitored. The discontinuity is present in both live and defunct funds, indicating that it is not a function of survivorship. Further, the discontinuity persists as funds age. Taken together, our results suggest the purposeful avoidance of reporting losses.

[18] See Agarwal et al. (2007b) for an analysis of the coverage of the CISDM database.

A possible alternative explanation for the discontinuity is that managers are optimistic in their valuations of illiquid securities held in their portfolios.[19] Managers might be prone to overvalue their own securities in the same way that retail investors have been shown to be overconfident in their abilities to pick winners, as in Barber and Odean (2001). We leave additional study of this behavioral explanation for future research.

If some managers are in fact purposely avoiding reporting losses, then investors may underestimate the potential for losses in the future and may overestimate the ability of hedge fund managers. We estimate that approximately 10% of observations in the database are distorted. This may be biased downward because we focus on one point of discontinuity—returns near zero. Though we show that investors respond to the frequency of positive returns, motivating our use of this specific discontinuity, some managers may be rounding up returns to achieve specific return targets related to high water marks and other benchmark returns. Since these other hurdles are fund specific and time varying, they are likely much more difficult to identify, especially in the aggregate distribution.

If hedge fund returns are distorted at the frequency we estimate, then our results have several implications for investors and regulators. Investors should be wary when using performance metrics based on the number of positive returns in a fund's history, as this measure appears to be unreliable. Also, investors who withdraw capital following a month or two of return inflation would benefit from somewhat overvalued fund shares, whereas investors who deposit capital would suffer. Regulators who argue that the low number of fraud cases prosecuted by the SEC means that additional oversight is unwarranted may find the robust discontinuity we document indicative of more widespread violations.

Appendix

The simulation algorithm includes the following steps to generate a sample of size n from a kernel density:

1. Draw n independent random integers, denoted by $I_1, \ldots, I_{\mathrm{n}}$, that are uniformly distributed on $\{1, 2, \ldots, n\}$,
2. generate n independent random variates, $W_1, \ldots, W_n$, that are distributed $k(\cdot)$, where k is the relevant kernel density,
3. the i^{th} element of the simulated sample from the kernel density is created as follows: $y_i = x(I_i) + b \cdot W_i$, where $x(I_i)$ is the $I_i{}^{\text{th}}$ element of the original sample and b is the bandwidth of the kernel density estimation.

The random variates drawn in step 2 are determined by the choice of the kernel. For instance, for the Gaussian kernel, the error distribution is the normal distribution, while for the Epanechnikov kernel, the corresponding distribution is the symmetric beta distribution. We choose the Gaussian kernel because of its speed and simplicity. As noted in Hörmann and Leydold (2000) and Silverman

[19] The authors thank Bing Liang for this suggestion.

(1986), the difference in efficiency between the optimal kernel and most other kernels is very small. Therefore, computational cost is frequently considered a leading criterion in the kernel choice. The variance-corrected algorithm, which ensures that the variance of the simulated sample equals the variance of the original data, replaces step 3 above by the following:

In step 3, the I^{th} element of the simulated sample from the kernel density is created as follows: $y_i = \bar{x} + (x(I_i) - \bar{x} + b \cdot W_i) \cdot c_b$, where $\bar{x}$ is the sample mean and $c_b = 1/\sqrt{1 + b^2\sigma_k^2/s^2}$, where σ_k^2 is the variance of the kernel and s is the standard deviation of the original sample.

REFERENCES

Abdulali, Adil, 2006, The bias ratio: Measuring the shape of fraud, *Protégé Partners* 30, 1–8.

Ackermann, Carl, Richard McEnally, and David Ravenscraft, 1999, The performance of hedge funds: Risk, return, and incentives, *Journal of Finance* 54, 833–874.

Agarwal, Vikas, Naveen D. Daniel, and Narayan Y. Naik, 2007a, Why is Santa so kind to hedge funds? The December return puzzle!, Working Paper, Georgia State University.

Agarwal, Vikas, Naveen D. Daniel, and Narayan Y. Naik, 2007b, Role of managerial incentives and discretion in hedge fund performance, Working Paper, Georgia State University.

Agarwal, Vikas, and Narayan Y. Naik, 2000, Multi-period performance persistence analysis of hedge funds, *Journal of Financial and Quantitative Analysis* 35, 327–342.

Agarwal, Vikas, and Narayan Y. Naik, 2004, Risks and portfolio decisions involving hedge funds, *Review of Financial Studies* 17, 63–98.

Avramov, Doron, Robert Kosowski, Narayan Y. Naik, and Melvyn Teo, 2007, Investing in hedge funds when returns are predictable, Working Paper, University of Maryland.

Barber, Brad M., and Terrance Odean, 2001, Boys will be boys: Gender, overconfidence, and common stock investment, *Quarterly Journal of Economics* 116, 261–292.

Barras, Laurent, Olivier Scaillet, and Russell Wermers, 2008, False discoveries in mutual fund performance: Measuring luck in estimated alphas, Working Paper, Swiss Finance Institute.

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

Bollen, Nicolas P. B., and Veronika K. Pool, 2008, Conditional return smoothing in the hedge fund industry, *Journal of Financial and Quantitative Analysis* 43, 267–298.

Bollen, Nicolas P. B., and Robert E. Whaley, 2008, Hedge fund risk dynamics: Implications for performance appraisal, *Journal of Finance* 64, 985–1035.

Brown, Stephen J., William N. Goetzmann, and Roger G. Ibbotson, 1999, Offshore hedge funds: Survival and performance, 1989–95, *Journal of Business* 72, 91–117.

Burgstahler, David, and Ilia Dichev, 1997, Earnings management to avoid earnings decreases and losses, *Journal of Accounting and Economics* 24, 99–126.

Busse, Jeffrey A., 2001, Another look at mutual fund tournaments, *Journal of Financial and Quantitative Analysis* 36, 53–73.

Carhart, Mark M., Ron Kaniel, David K. Musto, and Adam V. Reed, 2002, Leaning for the tape: Evidence of gaming behavior in equity mutual funds, *Journal of Finance* 57, 661–693.

Chevalier, Judith, and Glenn Ellison, 1997, Risk taking by mutual funds as a response to incentives, *Journal of Political Economy* 105, 1167–1200.

Dechow, Patricia M., Scott A. Richardson, and Irem Tuna, 2003, Why are earnings kinky? An examination of the earnings management explanation, *Review of Accounting Studies* 8, 355–384.

Degeorge, François, Jayendu Patel, and Richard Zeckhauser, 1999, Earnings management to exceed thresholds, *Journal of Business* 72, 1–33.

Del Guercio, Diane, and Paula A. Tkac, 2002, The determinants of the flow of funds of managed portfolios: Mutual funds versus pension funds, *Journal of Financial and Quantitative Analysis* 37, 523–557.

Devroye, Luc, 1997, Universal smoothing factor selection in density estimation, theory and practice, *Test* 6, 223–320.

Durtschi, Cindy, and Peter Easton, 2005, Earnings management? The shapes of the frequency distributions of earnings metrics are not evidence ipso facto, *Journal of Accounting Research* 43, 557–592.

Fung, William, and David A. Hsieh, 1997, Empirical characteristics of dynamic trading strategies: The case of hedge funds, *Review of Financial Studies* 10, 275–302.

Fung, William, and David A. Hsieh, 2004, Hedge fund benchmarks: A risk based approach, *Financial Analysts Journal* 60, 65–80.

Getmansky, Mila, Andrew W. Lo, and Igor Makarov, 2004, An econometric model of serial correlation and illiquidity in hedge fund returns, *Journal of Financial Economics* 74, 529–609.

Graham, John R., Campbell R. Harvey, and Shiva Rajgopal, 2005, The economic implications of corporate financial reporting, *Journal of Accounting and Economics* 40, 3–73.

Hayn, Carla, 1995, The information content of losses, *Journal of Accounting and Economics* 20, 125–153.

Hörmann, W., and J. Leydold, 2000, Automatic random variate generation for simulation input, in J. A. Joines, R. R. Barton, K. Kang, and P. A. Fishwick, eds.: *Proceedings of the 2000 Winter Simulation Conference* (Winter Simulation Conference, Monterey, CA), 675–682.

Liang, Bing, 2003, The accuracy of hedge fund returns, *Journal of Portfolio Management* 29, 111–122.

Mitchell, Mark, and Todd Pulvino, 2001, Characteristics of risk and return in risk arbitrage, *Journal of Finance* 56, 2135–2175.

Shleifer, Andrei, and Robert W. Vishny, 1997, The limits of arbitrage, *Journal of Finance* 52, 35–55.

Silverman, B. W., 1986. *Density Estimation for Statistics and Data Analysis* (Chapman and Hall, New York, NY).

Sirri, Erik R., and Peter Tufano, 1998, Costly search and mutual fund flows, *Journal of Finance* 53, 1589–1622.

Skeel, David A., and Frank Partnoy, 2007, The promise and perils of credit derivatives, *University of Cincinnati Law Review* 75, 1019–1051.

Starks, Laura T., 1987, Performance incentive fees: An agency theoretic approach, *Journal of Financial and Quantitative Analysis* 22, 17–32.

Waring, M. Barton, and Laurence B. Siegel, 2006, The myth of the absolute-return investor, *Financial Analysts Journal* 62, 14–21.