

Prospect Theory and Institutional Investors[†]

Paul G. J. O'Connell[‡]

FDO Partners, LLC and
State Street Associates, LLC

Melvyn Teo[§]

Singapore Management University and
FDO Partners, LLC

The Draft: October 2003

There is ample evidence that past performance affects the trading decisions of individual investors. This paper looks at this issue using a detailed database of currency trading decisions of institutional investors. Past performance manifestly affects currency risk-taking in this group, but the sign and magnitude of the effect runs counter to much of the existing theory and evidence. There is no evidence whatsoever of disposition effects; rather, the dominant characteristic is aggressive risk reduction in the wake of losses. This effect is more prominent later in the year, and among older and more experienced funds. A modified version of the loss aversion model of Barberis, Huang and Santos (2001) offers the best hope of adequately accounting for the observed behavior.

JEL Classification Numbers: G10, G11, G30

[†]Our thanks are due to State Street Corporation for help in obtaining the data. We thank Nick Barberis, Ken Froot, and seminar participants at State Street Associates, LLC, for helpful suggestions. The views expressed here are ours, and we are responsible for any errors or inaccuracies.

[‡]Corresponding address: 5 Revere Street, Cambridge MA 02138, USA. Tel. +1 (617) 234-9414. Fax: +1 (617) 864-5548. E-mail: poconnell@fdopartners.com.

[§]Corresponding address: School of Business, Tanglin PO Box 257, Singapore 912409. Tel. +65 822-0735. E-mail: melvynteo@fdopartners.com.

The link between risk-taking and past performance has recently come to prominence in the finance literature, and for good reason. On the empirical side, there is ample evidence that past trading performance can affect an investor's future trading decisions. For example, Odean (1998) has shown that individual investors prefer to sell past winners rather than past losers, and Coval and Shumway (2001) find that market makers in Treasury Bond futures contracts at the Chicago Board of Trade are far more likely to take on additional risk following morning losses than morning gains. On the theoretical side, a growing body of work looks at the impact of past performance on investor psychology. The resulting theories of disposition, overconfidence and dynamic loss aversion offer powerful explanations for a variety of asset pricing anomalies.

Notwithstanding the attention such "performance dependence" has received, there are wide gaps in the literature, and much remains to be understood. With some exceptions, the bulk of empirical research has looked at the equity trading of individual investors. A principal goal is to extend the empirical analysis of past performance to include other investor classes, and in particular, institutional investors. Institutional asset holdings now dwarf directly-held individual holdings in G7 countries, especially the U.S. and the U.K. It may be that institutional investors mimic individual investors in their sensitivity to past performance, but there are good reasons why this might not be the case. Professional money managers manage "other people's money," and so face substantially different incentives and reward structures from individual investors. This could alter the effect of past performance on their risk-taking, raising the question as to whether past performance matters for equilibrium pricing.

More important than simply establishing the prevalence of performance dependence, however, is identifying the precise channels through which it operates. To some extent, the theoretical literature has gotten ahead of the data here. Models based on the disposition effect (Grinblatt and Han (2002)), overconfidence (Daniel, Hirshleifer and Subrahmanyam (2001)), prospect theory (Barberis, Huang and Santos (2001)) and the like show great promise, yet these models will remain vulnerable to the "Fama critique" until the specific links between past performance and risk taking that they posit can be shown to have empirical content.¹ As discussed in the roundup of the literature in Section 2, some progress has been made in this direction, either through direct testing on individual investors (Glaser and Weber, 2003) or examination of the aggregate market predictions of a theory (Grinblatt and Han (2002)), but there is much that remains untested.

In this paper we take up the challenge of moving forward on both of these fronts. Our medium is a proprietary dataset encompassing the complete currency trades of 512 large institutional funds

over the period 1994–2002. Our main findings are striking. Past performance manifestly affects currency risk-taking, but the sign and magnitude of this effect differs substantially from what has been observed for individual investors. There is no evidence whatsoever of disposition effects: rather, the dominant feature of the behavior is aggressive risk-cutting in the wake of losses. We term this the *stop-loss effect*. Profits do bring some increase in risk-taking, but this increase reverses within a calendar quarter. The effect is pervasive across the major currencies, and characteristic of both foreign exchange and bond funds, though not pure equity funds. It is also more prominent later in the year, and among older and more experienced funds.

In teasing out an explanation for these patterns, we argue that disposition effect theories are simply not relevant. Both overconfidence theories and models of changing loss aversion offer a reasonable explanation for the increase in risk following profits, but neither does a good job explaining the scale of the stop-loss effect. Our conclusion is that overconfidence and loss aversion theories, while consistent with the evidence, need to be modified in important ways if they are to adequately account for the observed investor behavior.

The remainder of this paper is structured as follows. In the next section, we review the current state of theoretical and empirical knowledge on performance dependence. A description of the nature and characteristics of our data follows in Section II. Section III presents the main empirical results on the link between risk-taking and past performance, both unconditionally and conditional on long-horizon performance, age and experience. Section IV summarizes our conclusions.

I. Performance dependence

A. The current state of play

The link between risk-taking and past performance receives short shrift in the traditional, rational finance literature. As Coval and Shumway (2001) put it, in a setting where traders have standard Von Neumann-Morgenstern expected utility, profit opportunities are uncorrelated across the trading day, wealth effects are negligible, margin effects are unimportant and traders are fully rational, profits are not related to future trading activity.

Appealing though this view is, it is at odds with much of the available evidence. As already mentioned, Odean (1998) shows that individual investors are apt to sell past winners before past losers, a phenomenon earlier dubbed the disposition effect by Shefrin and Statman (1985). Further

evidence of the disposition effect has come from many sources, so much so that Dhar and Zhu (2002) term it one of the widely documented biases in investor behavior.² Other forms of performance dependence are also manifest. Coval and Shumway (2001) find that market makers in Treasury Bond futures contracts at the Chicago Board of Trade are far more likely to take on additional risk following morning losses than morning gains. Linnainmaa (2003) finds that day traders in Finland look at recent rather than total trading losses in deciding whether to continue their day-trading activities.

Lakonishok and Smidt (1986) catalog some rational reasons why performance dependence might occur. An investor may elect to sell his winners to maintain a desired asset allocation balance, or because the fundamental value he was seeking at the time he put on the trade has been realized.³ For taxable investors, tax structure provides an incentive towards disposal of assets with short-term capital losses. Accounting structure may create incentives for so-called “window-dressing.” Company size may be a determinant of portfolio membership, either exogenously or through its impact on transactions costs, creating a link between price and trading.⁴

A problem with many of these explanations is that, when brought to the data, they don’t seem to account for the performance dependence that is observed. Lakonishok and Smidt (1986) report that tax incentives are a secondary influence on trading. Odean (1998) writes that the strong preference to dispose of winners rather than losers displayed by the investors in his dataset cannot be attributed to portfolio rebalancing, subsequent portfolio performance, transactions costs or tax considerations. A striking finding by Odean (1998) is that stocks that are sold tend to outperform those that are not, suggesting that beliefs in mean-reversion to targets are irrational. Overall, Barberis and Thaler (2002) conclude that it is hard to account for the disposition effect on rational grounds.

A growing body of research, therefore, attempts to explain performance dependence in terms of investor psychology. There are two strands of work here. One strand focuses on investor *beliefs*, and argues that investors become overconfident in their ability to assess the moments of asset returns in the aftermath of investment success. This in turn may lead to increased risk-taking. Gervais and Odean (2001) develop a model in this vein in which periods of profitability are followed by periods of higher trading, a correlation observed amongst individual investors by Barber and Odean (2000), and Statman, Thorley and Vorkink (2003). Their model also predicts that experience will temper overconfidence. Locke and Mann (2001) confirm that, among professional traders on the floor of the Chicago Mercantile Exchange, traders with more experience are less likely to take more risk after a period of abnormally good profits. Glaser and Weber (2003) argue that it is differences in beliefs

about first moments, as opposed to second moments, that is at the root of such findings.

The second strand focuses on investor *preferences*. Based on experimental evidence, Kahneman and Tversky (1979) provide a description of how people depart from expected utility theory when offered a single risky gamble. Their description, called prospect theory, has three elements: (a) gains and losses matter, not the level of final wealth; (b) people are more sensitive to losses than gains; and (c) people are risk-averse over gains, and risk-seeking over losses. (a) and (b) appear to be common sense, yet together they represent a profound departure from traditional finance theory.⁵ We, like Barberis and Thaler (2002) use the term “loss aversion” to refer to (a) and (b).⁶ In the absence of a generally accepted term, we use the label “value inflection” to refer to (c), since it implies concavity of the utility function over gains and convexity of the function over losses.

This theory of a single risky gamble has been applied by a number of authors to explain the disposition effect. In particular, value inflection implies that investors would prefer to take the risk that their losing position improves rather than sell it now for a certain loss. Other theories of performance dependence are based on an extension of prospect theory to sequences of gambles. Barberis, Huang and Santos (2001) (hereafter BHS) point out that, if an investor cumulates his gains and losses, value inflection would seem to imply that he is more likely to take risk after a series of bad outcomes, and less likely after a series of good outcomes. This runs counter to the experimental evidence. Thaler and Johnson (1990), for example, show that individual willingness to take risk *increases* following recent success, an effect termed the “house-money” effect. To square prospect theory with such “mental accounting” (Thaler, 1990), BHS look to loss aversion rather than value-inflection. Specifically, they suggest that loss aversion increases as losses cumulate, and falls as gains cumulate. We label this “dynamic loss aversion.”

The importance of these theoretical developments based on performance dependence cannot be underestimated. If true, the theories offer potential explanations for some of the most enduring asset pricing anomalies. BHS argue that if changes in the value of holdings matter to investors, then the effect of prior outcomes goes some way towards explaining the three main puzzles associated with aggregate stock market behavior: the equity premium, excess volatility and long-horizon predictability. Grinblatt and Han’s (2002) model of the disposition effect shows that it can give rise to price momentum. Goetzmann and Massa (2003) show that a disposition effect factor constructed from the Grinblatt-Han model should be priced as a risk factor. In short, performance dependence matters because it can affect equilibrium prices.

B. The gaps in the literature

Taken together, this represents an impressive and promising body of knowledge. However, there are two obvious gaps that need to be filled. First, in order to know whether performance dependence is a market-wide phenomenon, there is a need to extend empirical research to include other investor classes. In particular, the empirical work needs to encompass institutional investors. As mentioned earlier, assets under institutional management exceed direct holdings of equity and fixed income securities by a good margin, especially in the Anglo-Saxon world. As of 1997, the ratio of institutional to direct holdings was 1.5 across G7 households (Davis, 2000).⁷ Despite this, virtually no empirical work has been done on this dominant investor class.

It is not just the scale of institutional assets that necessitates their study. The more potent motivation is that there are good reasons why performance dependence might take on a very different character within this investor class. Consider first incentives. Institutional investors typically receive a fixed percentage of assets under management as compensation. This creates a more complex relationship between performance and manager wealth than is true for the individual investor. A manager who loses money on behalf of his clients will see his compensation fall in direct proportion to assets under management, but he also faces the risk of redemptions from his fund, which would further erode his stream of income. The opposite is true, of course, when a manager achieves good performance. Going beyond fixed fees, institutional portfolio managers often receive performance-related compensation in the form of a bonus or direct participation in profits. Option-like payout structures are apt to further complicate the link between past performance and risk-taking.

Leaving aside incentives, there is the possibility that professional managers simply behave differently. There is some evidence to suggest this. In their investigation of brokerage investors in the Israeli market, Shapira and Venezia (2001), find that those who trade professionally are less prone to disposition effects than independent investors. Dhar and Zhu (2002) show that while individual investors exhibit the disposition effect *on average*, fully one fifth of the investors do not. Investor characteristics that temper the disposition effect include income level, professional occupation and trading experience. Cutting the other way is evidence from Griffin and Tversky (1992) that experts tend to be more overconfident than relatively inexperienced individuals.

The very promise of the theoretical models discussed above creates the second gap that needs to be filled: empirical identification of the most relevant theories. Consider the competing explanations offered for the equity momentum puzzle, the tendency of stock returns to persist over horizons of

a quarter to a year. Grinblatt and Han (2002) argue that the disposition effect is responsible. A relative willingness to close profitable trades before loss-making ones creates an excess supply of stocks with aggregate capital gains, and an excess demand for stocks with aggregate capital losses. In the equilibrium, this can generate momentum. Contrast this with the theory of dynamic loss aversion offered by BHS. In their framework, risk tolerance is directly related to past profitability, so that investors' willingness to take on risk *decreases* in the wake of losses. This leads to follow-on purchases of stocks that do well, and follow-on sales of stocks that do poorly, generating momentum. In a similar vein, theories of overconfidence predict that profits will lead investors to overestimate the precision of their expected returns, creating excess demand for stocks that have performed well.

How can these theories be empirically distinguished and validated? The linchpin is the sign of the performance dependence. Theories based on the disposition effect rely on *increased* willingness to take on risk in the wake of losses, relative to profits. Theories of dynamic loss aversion and overconfidence, by contrast, predict a *decreased* tolerance for risk in the wake of losses. Thus, at a basic level, measuring the sign and magnitude of the influence of past performance will help to discriminate among these theories. This is but one simple example, albeit an important one, of how more data analysis is needed in order to circumvent the Fama critique.

C. The contributions of this paper

This paper focuses exclusively on institutional investors. In particular, it looks at the daily currency trading activity of 512 large institutional funds over the period 1994–2002. While many of these funds also manage equities and fixed income securities, there are some good reasons to look in the first instance at their currency activity. More so than equity or fixed income trades, currency trades are driven by the fund manager rather than the underlying stakeholder. If the stakeholders in a technology mutual fund sell after declines in net asset value, the fund manager himself sells the constituent stocks of the fund. In order to capture the direct effect of past performance at the institutional level, one needs to be able to identify the actively managed piece of the fund. In other words, the question is whether the fund manager's choice between, say, Dell and Gateway stock is influenced by his own performance in allocating across these stocks.⁸ Second, forward currency contracts are derivatives, in zero net supply. This eliminates the possibility of aggregate capital gains or losses at the level of each currency, which alters the pricing implications that come from theories such as the Grinblatt and Han (2002) model.

[Table I here]

The level of detail in the dataset allows us to go some way towards identifying the relevant theories of performance dependence for this investor class. Table I gives the taxonomy of questions that we address using both the cross-sectional and time-series variation in the data, together with the rationale for each. The questions are broadly divided into three categories. In the first, “Basic dynamics,” the goal is to size up the degree of performance dependence that is present. Any link between risk-taking and lagged P&L represents a departure from the bulk of traditional finance theory. A finding that lagged performance impacts risk-taking negatively would tend to support theories of disposition effects, whereas evidence to the contrary would lend support to theories of overconfidence and dynamic loss aversion. We also wish to know whether any performance dependence measured is economically relevant. In the second category, “Conditional dynamics,” the questions investigate the difference between past gains and losses, and also the extent to which the effects measured are sensitive to *cumulative* of losses. This gets at the question of whether investors integrate outcomes across sequences of trades, or treat them independently. In addition, we look for the presence of calendar effects in the data. It is often noted that investment manager performance bonuses are paid annually, typically on a calendar-year cycle. This could well give rise to different degrees of performance dependence early and late in the year.

In the third category, “Cross-sectional features,” we use the latitude of the dataset to look at a variety of potentially important cross-sectional characteristics: currency, fund type, fund age and fund experience matter. To the extent that it is present, behavioral tendencies are expected to attenuate as investors gain in experience. This tempering is central to the model of Barber and Odean (2000), and Locke and Mann (2001) use it to provide an identification scheme for empirically distinguishing between overconfidence and dynamic loss aversion. Finally, we gauge the scope of framing. The theory of narrow framing (Redelmeier and Tversky (1992)) implies that performance dependence will operate at the level of the individual currency—there will be no cross-asset or portfolio effects. Hence it is of interest to know whether it is single-currency or portfolio losses that are at the root of performance dependence.

A critical dimension of all of these tests is the time-period over which gains and losses are measured, and over which they exert an influence on future decisions. Here the extant theory offers us less guidance. Benartzi and Thaler (1995) argue that, with Kahneman and Tversky-type loss aversion, the equity premium can be reconciled with the outstanding supply of stocks and bonds

if investors evaluate their gains and losses once a year. This may make sense at the level of the individual investor: as Barberis and Thaler (2002) note, we receive our most comprehensive mutual fund reports once a year, and do our taxes once a year. For institutional investors, however, this time frame may not be appropriate. Coval and Shumway (2001) write that: “...when fund managers are averse to losses, it is not clear whether their aversion relates to monthly, quarterly or annual horizons.” They argue that the most important advantage of their dataset is that the time horizon is clear: CBOT market makers have incentives that encourage them to evaluate their performance once a day. Rather than take a stance on a particular time horizon that is relevant, we consider a number of fixed forecast horizons, and let the data choose the appropriate lag structure.

II. Data

A. Raw inputs

The data used in the analysis is provided by State Street Corporation, one of the world’s largest investor services providers. State Street clients are primarily large institutional money managers, and the total of all funds serviced by the Corporation is currently USD 8.4 trillion, approximately 16 percent of total global assets. Our sample covers the period December 31st, 1993–January 1st, 2003, and comprises over 8 million individual trade records undertaken by some 8,500 anonymous funds. Each record provides us with the currency pair traded, the exchange rate, and the tenor or duration of the contract.

Given the distributional assumptions needed for estimation, quality of the data series is important. Hence the analysis is restricted to the larger funds in the universe, as these tend to have more frequent, continuous trading. Moreover, only trades in the 11 major currencies are included.⁹ Mindful of survivorship bias, the requirement for inclusion in our sample is that a fund be in the 95th percentile of trading volume in one or more of 6 regularly sampled weeks over the nine-year sample period. This criterion selected a subset of 512 funds that account for an average of 72 percent of the volume across the 11 currencies.

There are a number of important fund characteristics to look at. The first is fund life. Although specific information on fund life is not available in the database, an examination of currency holdings makes it clear that most of the funds are not active in the currency markets for the entire sample. Indeed only two percent of the funds have nonzero currency holdings on every day of the sample.

Of course, a fund manager may make an active decision to hold no open currency positions, so zero holdings may not imply that a fund is “dead.” Cognizant of this, one way to proceed is to measure the life of each fund from the first day of nonzero holdings to the last day of nonzero holdings, and then to gauge the likelihood that this is a biased estimate from the incidence of zero holdings during this estimated life. Calculated in this way, the mean fund life is about 4.5 years, while the incidence of zero exposure throughout fund life is only 12 percent, suggesting that the lifespan estimates are reasonable. A second important fund characteristic is base currency, since measured currency risk ought to exclude base-currency holdings. The breakdown by base currency is as follows: U.S. dollar, 67 percent; Australian dollar 12 percent; Canadian dollar 6 percent; euro 3 percent; Japanese yen 3 percent; British pound 3 percent; others 6 percent. Finally, it will be of interest to consider the underlying type of each fund. The database includes comprehensive information on the total holdings of each fund by asset class for the year 2001. Based on this, the funds are classified as fixed income, equity or currency for that year.¹⁰ The resulting categorization comprises 158 fixed income funds, 71 equity funds and 149 currency funds.

B. Basic series

The first step is to construct flow and holdings series for each fund across the currencies. Each day, net flows by currency, fund and tenor are measured.¹¹ All flows on date t with tenor s are converted to dollars by dividing by the appropriate forward currency exchange rate f_t^s , where f is units of foreign currency per dollar. Holdings are built up by cumulating these flows, after adjusting for mark-to-market gains and losses on each day’s pre-existing positions. For a position with tenor s on date $t - 1$, the marked-to-market gross return between date $t - 1$ and t is f_{t-1}^s / f_t^{s-1} , reflecting the fact that it is one day closer to maturity. It is these mark-to-market gains and losses that provide the key profit-and-loss (P&L) series that are used to measure performance. Any currency holdings that do come to maturity—that is, reach a tenor of zero—are treated as delivered, and removed from holdings on value date. This would occur, for example, if a fund purchased and took delivery of spot local currency to facilitate the purchase of an underlying equity or fixed income security. Such transactions are common for fixed income and equity funds, so negative serial correlation at short horizons is likely to be observed in the holdings series for such funds.

With holdings in hand, it is a simple matter to calculate the second key series—a measure of risk exposure. Let \mathbf{h}_{it} be the vector of currency holdings for fund i on date t . Risk is measured

as the standard quadratic form $\mathbf{h}'_{it}\boldsymbol{\Sigma}\mathbf{h}_{it}$, where $\boldsymbol{\Sigma}$ is the covariance matrix of annualized currency returns constructed from exponentially-weighted daily currency returns¹² The relevant $\boldsymbol{\Sigma}$ matrix differs according to the base currency of each fund. For example, a euro position held by a dollar-based fund entails much more risk than the same position held by a Scandinavian fund, relative to base currency.

Figure 1 plots the holdings series for each of our currencies aggregated across all 512 funds, grouped into four rough regions: North America, Japan and Antipodes, Europe and Scandinavia. There is a large amount of variation in the raw holdings numbers, and so to render them comparable, they are measured in units of trading days. For example, if a fund is long \$5 million against the euro, and the fund's average daily EUR/USD volume is \$1m, then it is counted as having 5 trading days worth of holdings. Figure 1 illustrates that, throughout the sample, funds have tended to be long the dollar and short other currencies. However, towards the end of the sample, this tendency waned considerably.

Holdings tell only part of story, however, a fact that becomes abundantly clear when risk is examined. Figure 2 plots the aggregate risk exposure held by the funds in each of the currencies. Notice in particular that the exposure to the Japanese yen and British pound has remained quite high in the recent period. This implies that, individually, the funds in the universe continue to maintain large exposures to these currencies. Some funds are long and some funds are short, with the positions netting out to give an aggregate holding of close to zero. In other words, there is a considerable amount of disagreement across the fund positions. This cross-sectional richness contributes to the statistical power of the data sample.

[Figure 1 here]

[Figure 2 here]

C. Persistence

It is well-established that portfolio flows in underlying assets such as equities tend to be persistent

(see Froot, O’Connell and Seasholes (2001)). The question arises as to whether the same is true for institutional currency flows. Figure 3 plots the sample autocorrelation function for daily currency risk exposures out to 20 lags, together with 95 percent confidence bands. The functions are plotted for three different levels of aggregation. Panel (a), “Aggregated by currency and fund,” adds up the total risk of all funds across all currencies to arrive at a single time series. Panel (b), “Aggregated by currency,” adds up the total risk across all funds in each currency separately, and shows the autocorrelation estimates for the currency panel. Analogously, Panel (c), “Aggregated by fund,” adds up the total risk across all currencies for each fund, and shows the autocorrelation estimates for the fund panel. At the aggregate and individual currency level, there is evidence of positive serial correlation at the 1-day and 5-day frequencies. Interestingly, however, there is no such persistence at the individual fund level.¹³ Individual funds are not persistent in their actions, but funds tend to mimic one another. A substantially similar picture emerges from examination of weekly risk autocorrelations. Overall, this echoes the Froot and Tjornhom (2002) finding of statistically significant cross-fund lags in equity flows to developed and emerging markets.

Turning to performance, Figure 4 plots similar sample autocorrelation functions for P&L. Here there is no evidence of serial correlation, indicating that the lead-lag effects in risk-taking do not engender persistent performance. Again, the same is true at weekly frequencies. The interesting implication is that managers do not undergo cycles in profitability—for the most part, profits are independent from one period to the next.

[Figure 3 here]

[Figure 4 here]

III. The evidence

A. Basic dynamics

The tool we use to address the questions under the “basic dynamics” heading is an unrestricted

vector autoregression. Analogously to the serial correlation analysis presented above, we estimate panel VARs for risk and P&L at the aggregate, the currency and the fund level. Figure 5 shows the essential information that comes out of this exercise. The model allows for heteroskedasticity across currencies and funds, and the lag length for each model is set at 13 weeks, the value selected by the Bayes-Schwartz Information Criterion for the panel fund regression. The first column of plots in the figure shows the impact that a unit-standard deviation shock to P&L has on risk, while the second column shows the impact that a unit-standard deviation shock to risk has on P&L. The effects measured on the vertical axes are also scaled in standard deviation units, and 90 percent confidence intervals based on the maximum likelihood standard errors are sketched in lighter weight around each function. The own-equation effects are similar to those conveyed by Figures 3 and 4, and so are omitted.

[Figure 5 here]

Performance dependence is manifest in the data. At all levels of aggregation, past performance exerts a positive and statistically significant effect on risk-taking, and the impact persists for between six and eight weeks. The economic impact is significant too: for the panel fund regression, a one-standard deviation shock to P&L produces a one-standard deviation change in risk-taking after four weeks. In dollar terms, this means that a \$1 million dollar profit produces an increase in currency holdings of approximately \$0.3 million over the subsequent four weeks. Importantly, the serial correlation estimates for P&L calculated earlier make clear that this result is not simply due to persistence in profits or losses.¹⁴ Turning to the second column of plots, there is no appreciable effect in the other direction: as might have been expected, increases in risk-taking do not have a meaningful effect on profits. There is some indication that returns improve with risk-taking, though naturally risk rises in tandem with this.

An important feature of these results is the relatively short horizon over which the effects play out. True, a model estimated on weekly changes is hardly well-suited to capturing long-horizon phenomena, but we find that the effects are no more durable when the model is re-fitted at the monthly horizon. This is significant, because if performance dependence is to stand as a viable explanation for the equity premium puzzle, excess volatility, long-horizon predictability and the like, as its proponents argue, the effects must be long-lived. The dynamics measured here suggest

that, at least in an unconditional sense, institutions have shorter memories. It remains to be seen whether this is true in the various conditional cases which we look at below.

The most striking aspects of the results is the complete absence of any evidence of disposition effects. Rather, risk-taking is directly proportional to performance, lending support to theories that predict a positive relationship such as overconfidence or dynamic loss aversion. It turns out that this conclusion is only strengthened when gains and losses are examined separately, and it is to this that we now turn.

B. Conditional dynamics

There is some existing evidence that gains affect risk-taking in a manner different from losses. Coval and Shumway (2001) find that, among Chicago Board of Trade proprietary traders, risk-taking responds strongly to losses, but only weakly to profits, and Odean (1999) reports that while losses are equally likely to produce buying or selling, gains are apt to lead to selling. Theory also predicts some asymmetries. Consider the model of BHS. The baseline utility function they adopt is shown as the heavier central line in Figure 6. This function, defined over gains and losses, is almost identical to the loss averse function calibrated by Kahneman and Tversky, absent value inflection. The innovation of the Barberis, Huang and Santos approach is in their modelling of how this function changes in response to gains and losses. They conjecture that, after a gain, the function slides down and to the left, becoming less concave, while after a loss, the function pivots at the origin, becoming more concave.

[Figure 6 here]

[Figure 7 here]

To address this issue, we distinguish between the dynamic effects of gains and losses. Figure 7 plots separate impulse response functions for gains and losses, estimated from the fund-by-fund data panel.¹⁵ There is in fact a striking difference in the two response functions. Gains produce transitory increases in risk-taking that taper off after about six weeks. Beyond that there is evidence

of “take-profit” activity as the impulse response function turns statistically negative. By contrast, the effects of losses are both stronger and more permanent. We label this phenomenon the *stop-loss effect*. Note that the impulse response function sketched in the lower left-hand corner of Figure 5 is simply an average of these two functions. The implication is that the relatively short-lived average effect illustrated there masks separate effects of gains and losses that appear to be durable.

The critical point to make here is that the shape of these impulse functions runs counter to much of the existing empirical and theoretical work. The pattern of a modest increase in risk-taking following gains coupled with aggressive risk-cutting following losses is not evident in any of the existing empirical work. On the theoretical side, it is not consistent with any theoretical variant of the disposition effect. Moreover, it does not square with the standard models of overconfidence. To see this latter point, note that the asymmetry in overconfidence models goes the other way. As Gervais and Odean (2001) write, overconfidence theory is premised on the psychological finding of biased self-attribution: when people succeed, they believe that the success was due to their personal abilities; when they fail, they attribute their failure to chance and outside factors.¹⁶ Thus overconfidence would lead managers to extend their positions after gains, but not to cut their positions after losses. Finally, the pattern in the impulse response functions is not consistent with the theory of dynamic loss aversion proposed by BHS, at least in its vanilla form. Once again, in their model, the asymmetry goes the other way. The reason is that, with the parameters they calibrate, the cost of risk—the difference between the expected value of a risky prospect and its certainty-equivalent outcome—changes more after a gain than after a loss for all but the most risky trades.¹⁷

Overconfidence theory could be used to explain the stop-loss effect if it were generalized to allow for an ebb and flow of confidence, rather than the ratchet effect of successes and failures envisioned in the traditional formulation. However confidence would have to be modelled in a way that makes it much more susceptible to losses than gains. Such a theory, which we might label *underconfidence*, would represent a departure from the theory of biased self-attribution that underlies overconfidence theory.

For the dynamic loss aversion model of BHS to explain the stop-loss effect, it needs to be rejiggered to allow the cost of risk to respond much more strongly to losses rather than gains if it is to explain our results. One possibility is to allow for a larger slope change in the left arm of the utility function in Figure 6 following losses. BHS do consider steeper slopes, but point out that these raise the average level of loss aversion to the point where it may be unrealistic. Moreover,

to achieve even a symmetric response to gains and losses would require the pivot response of the utility function to be stronger than the highest response contemplated by BHS. Instead, to alter the relative impact of gains and losses, we propose the two alternatives sketched in Figure 8. The first modification, labelled (a), shows the utility function sliding up and to the right after a loss, rather than pivoting. Under this alternative, losses coming on the heels of earlier losses are no more painful than before, but gains coming on the heels of earlier losses are more rewarding. The second modification, labelled (b), shows the left arm of the utility function pivoting up after a gain. In this instance, losses are less painful following gains, but gains are always equally rewarding. These modifications, either separately or together, would yield risk-taking behavior that is much more consistent with what we observe for institutional currency trading.

[Figure 8 here]

As described in Table I, a second level of conditioning that is informative to consider is that with respect to prior losses. The notion here is that the response to gains and losses documented above may depend on the cumulation of profits prior to each realization. This bears on the question of whether managers integrate gains and losses over sequences of trades. In other words, if a manager sustains a loss L , does he assess the risky payoff π from a subsequent trade as $U(L + \pi)$, or simply as $U(\pi)$. Integration of outcomes is required if value inflection and in particular the convexity of the Kahneman-Tversky utility function is to cause disposition effects. BHS interpret the evidence of Thaler and Johnson (1990) to mean that people do not integrate sequential outcomes. What of the investors in our data sample?

To answer this, we measure the impact of gains and losses separately, after conditioning on the previous one-week and one-month loss. Figure 9 presents the one-week and one-month coefficient estimates, with standard errors in parentheses. If institutional managers come to the market carrying pre-existing losses, they are less likely to cut risk after a further loss than they would be if they were carrying pre-existing gains. This difference is statistically significant at the monthly level. If they are carrying pre-existing losses and experience a gain, then they are likely to increase risk-taking, whereas if they are carrying pre-existing gains, they are likely to take profit. Thus there is some evidence of integration of outcomes. This gives a sense for the dynamics of the utility functions in Figure 8. Sequences of profits do not produce a continuing upward pivot of the utility function.

Rather, after a period, the pivot reverses. By contrast, sequences of losses do produce continuing shifts of the utility function, albeit at a slower rate than caused by the initial loss.

[Figure 9 here]

The final question of interest in this category is whether there are calendar effects present in the data. The rationale here is that the incentives faced by many fund managers vary throughout the year. For example, a manager who receives a performance-related bonus is likely to be more risk-tolerant early in the performance measurement period, reflecting the option-like structure of his payoffs. Does this carry over to his sensitivity to past gains and losses? Figure 10 shows that it does. It measures the impulse response functions shown in Figure 7 separately for each half of the calendar year.¹⁸ It's clear that managers are *conditionally* more risk-tolerant in the first half of the year. Gains in the first half of the year lead to incremental risk-taking, but there is no such evidence in the second half of the year. Correspondingly, losses in the first half produce very little stop-loss activity: it is only in the second half of the year that managers systematically cut risk following losses. The clear message is that managers husband their portfolios to a greater degree in the latter half of the year.

[Figure 10 here]

C. Cross-sectional features

Having investigated the dynamic relationship between risk-taking and P&L, we now turn to the cross-sectional features of the data. As shown in Table I, we are interested in understanding whether the effects identified are pervasive, in the sense that they apply across currencies, fund types and so on.

[Table II here]

Table II shows the effect of the first eight lags of P&L on risk-taking across each currency and across the three fund types discussed in Section II. Looking first at the currencies, the basic pattern observed in the full panel is seen to characterize seven of the ten currencies, the exceptions being Denmark, Sweden and New Zealand. Trading volume in these currencies is 1.25 percent, 2.43 percent and 1.46 percent of total volume respectively. Thus the patterns measured earlier apply to the bulk of currency trading in our sample. Among fund types, the black sheep is the equity category. FX and bond funds display essentially the same sensitivity to past P&L as was documented earlier for the full group. Equity funds, by contrast, display a somewhat random response to past performance that is statistically insignificant. This accords with the folk wisdom that equity fund managers simply care less about the currency component of their returns. In fairness, it must also be said that the statistical power of the equity sample is lower, since the number of equity funds, at 71, is about half the number of currency or bond funds in the sample.

Much has been made in the empirical literature about how investor age and experience can influence behavioral biases. According to Barber and Odean (2000) and Dhar and Zhu (2002), older and more experienced retail investors are less overconfident than younger and less experienced retail investors. As already mentioned, Locke and Mann (2001) use this fact to empirically discriminate between overconfidence and dynamic loss aversion. We look for a similar pattern among institutional investors. The sample is split into a formation period—December 31, 1993–December 31, 1999—and an evaluation period—January 1, 2000–January 1, 2003. A fund's age is proxied by the fund's first trade date during the formation period, and experience is gauged by the numbers of days during the formation period that the fund actually traded. Then we use a simple two-step procedure. In step one, the sensitivity of each fund to lagged P&L is measured across the evaluation period. Then in step two, the cross-section of coefficients is regressed on fund age and fund experience. Table III reports the results for the first lag of the regression coefficient on total profits, total gains, and total losses. Both age and experience exert a statistically significant mitigating effect on the total profit coefficient at the first lag. This suggests that the performance dependence we have observed is sensitive to learning, as in the confidence model of Gervais and Odean (2001). More interesting, though, is the fact that, once again, the effect is asymmetric for gains and losses. Age and experience tend to decrease the magnitude of the coefficient on lagged gains, but to *increase* the magnitude of the coefficient on lagged losses. So the older, wiser funds eschew added risk in the wake of gains, but cut risk more aggressively in the wake of losses.

[Table III here]

Finally, there is the question of framing. To our knowledge, no existing study has been able to address this important issue. Does performance dependence operate at the level of a single security, in which case investors are engaged in *narrow framing*, or is it driven by the performance of the portfolio. It is one thing if investors respond to changes in wealth, rather than wealth itself, as prospect theory would suggest. It is another if changes in wealth are narrowly defined asset-by-asset. Table IV presents the results. The coefficients are naturally much smaller than the own P&L coefficients shown in Table II, since conjugate P&L is a much larger quantity on average than own-P&L. Interestingly, no clear pattern emerges from the coefficients. If we focus on the major currencies, there is some mild evidence from the point estimates that portfolio profits increase risk-taking in the British pound, the Australian dollar and Japanese yen, but this doesn't appear to be statistically significant. Only for the euro is there significant evidence of an effect, which appears to be a negative one. Overall, these results suggest that these institutional currency managers are narrow framers.

[Table IV here]

IV. Conclusion

The sheer level of detail in our dataset has allowed us to learn much about performance dependence among institutional investors. One could summarize what we have learnt as follows. Past performance manifestly affects currency risk-taking. The sign and magnitude of this effect, however, runs counter to much of the evidence gleaned from data on individual investors. There is no evidence whatsoever of disposition effects: rather, the dominant characteristic is aggressive stop-loss trading in the wake of losses. Gains do tend to elicit a mild increase in risk-taking, but this increase reverses within a calendar quarter. These findings are pervasive across the major currencies, and foreign exchange and bond funds. However, they do not seem to characterize equity funds.

The patterns observed facilitate discrimination between the various theories of performance dependence. Disposition effect theories are not relevant for this investor class. Both overconfidence

theories and models of changing loss aversion offer a reasonable explanation for the increase in risk following profits, but neither does a good job explaining the stop-loss behavior. Our conclusion is that modifications of overconfidence and dynamic loss aversion models that permits losses to have much stronger effects than gains offers the best hope of adequately accounting for the observed investor behavior.

We also learn a considerable amount from conditioning the results on a variety of variables. Time-of-year matters, in the sense that these investors appear to be more risk tolerant in the first-half of the year. We conjecture that this owes to the incentives faced by fund managers. In addition, age and experience matter profoundly: older and wiser funds do not increment their risk-taking following gains, and are assiduous in cutting risk once losses occur. Finally, all of these effects appear to be narrow, in the sense that they operate at the level of the individual currency rather than the portfolio.

References

- Barber, Brad, and Terrance Odean, 2000, Trading is hazardous to your wealth: The common stock investment performance of individual investors, *Journal of Finance* 55, 773-806.
- Barber, Brad, and Terrance Odean, 2001, Boys will be boys: Gender, overconfidence and common stock investment, *Quarterly Journal of Economics* 116, 261-292.
- Barberis, Nicholas, and Richard Thaler, 2002, A survey of behavioral finance, NBER Working Paper 9222.
- Barberis, Nicholas, Ming Huang, and Tano Santos, 2001, Prospect theory and asset prices, *Quarterly Journal of Economics* 116, 1-53.
- Benartzi, Shlomo, and Richard Thaler, 1995, Myopic loss aversion and the equity premium puzzle, *Quarterly Journal of Economics* 110, 75-92.
- Coval, Joshua D., and Tyler Shumway, 2001, Do behavioral biases affect prices? Unpublished working paper, University of Michigan.
- Davis, E. Philip., 2000, Implications of the growth of institutional investors for the financial sectors, Discussion Paper PI-0001, The Pensions Institute, Birkbeck College.
- Davis, E. Philip and Benn Steil, 2001, *Institutional investors* (Cambridge, MIT Press).
- Daniel, Kent, David Hirshleifer and Avanidhar Subrahmanyam, 1998, Investor psychology and security market under- and over-reactions, *Journal of Finance* 53, 1839-1886.
- Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam, 2001, Overconfidence, arbitrage, and equilibrium asset pricing, *Journal of Finance* 56, 921-965.
- Dhar, Ravi, and Ning Zhu, 2002, Up close and personal: An individual level analysis of the disposition effect, Yale International Center for Finance Working Paper No. 02-20.
- Fama, Eugene, 1998, Market efficiency, long-term returns, and behavioral finance. *Journal of Financial Economics* 49, 283-306.
- Ferris, Stephen, Robert Haugen and Anil Makhija, 1988, Predicting contemporary volume with historic volume at differential price levels: Evidence supporting the disposition effect, *Journal of Finance* 43, 677-697.
- Froot, Kenneth, Paul G. J. O'Connell, and Mark S. Seasholes, 2001, The portfolio flows of institutional investors, *Journal of Financial Economics* 59, 151-193.
- Froot, Kenneth, and Jessica Tjornhom, 2002, Decomposing the persistence of international equity flows, NBER Working Paper No W9079.
- Genesove, David, and Christopher Mayer, 2001, Loss aversion and seller behavior: Evidence from the housing market, *Quarterly Journal of Economics* 116, 1233-1260.
- Gervais, Simon, and Terrance Odean, 2001, Learning to be overconfident, *Review of Financial Studies* 14 (1), 1-27.
- Glaser, Markus and Martin Weber, 2003, Overconfidence and trading volume, CEPR Discussion Paper 3941.
- Goetzmann, William, and Massimo Massa, 2003, Disposition matters: Volume, volatility and price impact of a behavioral bias, Yale ICF Working Paper No 03-01.
- Grinblatt, Mark, and Bing Han, 2002, The disposition effect and momentum, NBER Working Paper 8734.
- Griffin, Dale, and Amos Tversky, 1992, The weighing of evidence and the determinants of confidence, *Cognitive Psychology* 24, 411-435.
- Harris, Lawrence, 1988, Discussion of predicting contemporary volume with historic volume at different price levels: Evidence supporting the disposition effect, *Journal of Finance* 43, 698-699.

- Hastorf, Albert H., David J. Schneider, and Judith Polefka, 1970, *Person Perception* (Addison-Wesley, Reading MA.).
- Heath, Chip, Stephen Huddart and Mark Lang, 1999, Psychological factors and stock option exercise, *Quarterly Journal of Economics* 114, 601–627.
- Kahneman, Daniel, and Amos Tversky, 1979, Prospect theory: An analysis of decision under risk, *Econometrica* 47, 263–291.
- Lakonishok, Josef, and Seymour Smidt, 1986, Volume for winners and losers: Taxation and other motives for stock trading, *Journal of Finance* 41, 951–974.
- Linnainmaa, Juhani, 2003, The anatomy of day traders, Unpublished working paper, Helsinki School of Economics.
- Locke and Mann, 2001, House money and overconfidence on the trading floor, Unpublished working paper, George Washington University.
- Odean, Terrance, 1998, Are investors reluctant to realize their losses, *Journal of Finance* 53, 1775–1798.
- Odean, Terrance, 1999, Do investors trade too much? *American Economic Review* 89, 1279–1298.
- Oehler, Andreas, Klaus R. Heilmann, Volker Lager, and Michael Oberlander, 2002, Dying out or dying hard: Disposition investors in stock markets, Unpublished working paper, University of Bamberg.
- Redelmeier, Donald and Amos Tversky, 1992, On the framing of multiple prospects, *Psychological Science* 3: 191–193.
- Shapira, Zur and Itzhak Venezia, 2001, Patterns of behavior of professionally managed and independent investors, *Journal of Banking and Finance* 25, 1573–1587.
- Shefrin, Hershey, and Meir Statman, 1985, The disposition to sell winners too early and ride losers too long, *Journal of Finance* 40, 777–790.
- Statman, Meir, Steven Thorley, and Keith Vorkink, 2003, Investor overconfidence and trading volume, Unpublished working paper, University of Santa Clara.
- Thaler, Richard, 1999, Mental accounting matters, in Daniel Kahneman and Amos Tversky, ed.: *Choice, Values and Frames* (Cambridge, Russell Sage Foundation.).
- Thaler, Richard, and Eric Johnson, 1990, Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice, *Management Science* 36, 643–660.

Notes

¹On behavioral finance, Fama writes: “[G]iven the demonstrated ingenuity of the theory branch of finance, and given the long litany of apparent judgement biases unearthed by cognitive psychologists, it is safe to predict that we will soon see a menu of behavioral models that can be mixed and matched to explain specific anomalies.” (1998, p. 291)).

²See for example Lakonishok and Smidt (1986), Ferris, Haugen and Makhija (1988) and Heath, Huddart and Lang (1999). Genesove and Mayer (2001) find evidence of disposition effects in the housing market, and Oehler et al. (2002) find that it characterizes many world stock markets.

³We use the word “his” advisedly: Barber and Odean (2001) show that men trade more and earn lower returns than women, perhaps indicating that they are more susceptible to behavioral biases.

⁴See Harris (1988) for a discussion of the potential link between transactions costs and disposition to trade.

⁵(a) implies that agents care about wealth for its own sake, irrespective of what it implies for consumption, and (b) introduces a kink in the utility-of-wealth function at the level of current wealth.

⁶Grinblatt and Han (2002) and Coval and Shumway (2001) refer to (a), (b) and (c) collectively as “loss aversion.”

⁷Institutional holdings now equal 100 percent of GDP in G7 countries, and 200 percent in the U.S. and U.K. (Davis and Steil, 2001).

⁸We are exploring this more complex link in related work.

⁹The list of currencies is: Danish kroner, Norwegian kroner, Swedish kroner, Swiss franc, British pound, Australian dollar, Japanese yen, New Zealand dollar, Canadian dollar and euro. Prior to 1999 synthetic euro return and flow series are constructed by weighting across the euro member countries.

¹⁰Funds with fixed holdings in excess of equity holdings are defined as fixed income funds, and *vice versa*. Currency funds have no equity or fixed income positions.

¹¹99 percent of the trades value within one year of trade date, so trades with maturity greater than 265 trading days are ignored.

¹²The exponential decay rate used is 0.998, implying a half-life of decay for past observations of about 350 trading days.

¹³As mentioned earlier, the negative serial correlation evident at order two arises from the spot trades of fixed income and equity funds.

¹⁴In results not reported, we confirm that the changes in risk-taking arise from active trading rather than simply the passive changes in P&L.

¹⁵The results from the other levels of aggregation are similar.

¹⁶On the same point, Daniel, Hirshleifer and Subrahmanyam (1998) write that "...individuals too strongly attribute events that confirm the validity of their actions to high ability, and events that disconfirm the action to external noise or sabotage."

¹⁷To see this, note that a 10 percent loss is calibrated to pivot the left side of the utility function in Figure 6 from a slope of 2.25 to 2.55. This lowers the utility of a subsequent loss by an amount equal to $0.3 \times \text{loss}$. However a 10 percent gain is calibrated to raise the utility of a subsequent loss by $\min\{1.25, 1.25 \times \text{loss}\}$. This is a much bigger number for all but the most extreme negative realizations.

¹⁸The assumption here is that performance measurement periods correspond to calendar years.

Figures

Figure 1: Holdings by currency aggregated across all funds, 1/1/1995-1/1/2003

Figure 2: Risk exposure by currency aggregated across all funds, 1/1/1995-1/1/2003

Figure 3: Daily sample autocorrelation function for risk aggregates. The sample period is from 1/1/1995 to 1/1/2003. Sketched in lighter weight are 95% confidence bounds.

Figure 4: Daily sample autocorrelation function for P&L aggregates. The sample period is from 1/1/1995 to 1/1/2003. Sketched in lighter weight are 95% confidence bounds.

Figure 5: Impulse response functions for shocks to risk and P&L. The sample period is from 1/1/1995 to 1/1/2003. The sample includes ten currencies and 512 funds. The aggregated by fund and currency panel consists of a series of length T , where T is the length of the sample period. The aggregated by currency panel consists of 10 stacked series of length T . The aggregated by fund panel consists of 512 series of various lengths, depending on fund life. To generate the impulse response functions, panel VARs for risk and P&L are estimated. The lag length of each VAR is set at 13 weeks, the value selected by the Bayes-Schwartz Information Criterion for the aggregated by fund panel regression. The VARs allow for heteroskedasticity across currencies and funds. The vertical axes are scaled in standard deviation units of risk. Sketched in lighter weight around each function are 90% confidence interval bounds based on maximum likelihood standard errors.

Figure 6: Barberis, Huang and Santos (2001) utility function for gains and losses

Figure 7: Fund panel impulse response functions for gains and losses. The sample period is from 1/1/1995 to 1/1/2003. To generate the impulse response functions, risk is regressed on weekly lags of P&L conditional on gains and weekly lags of P&L conditional on losses for the fund panel. The regression allows for heteroskedasticity across funds and uses the same lag structure as in the VARs in Figure 5. Sketched in lighter weight around each function are 90% confidence interval bounds based on maximum likelihood standard errors.

Figure 8: Barberis, Huang and Santos (2001) utility functions modified to fit the evidence

Figure 9: Impact of P&L on risk-taking, conditional on past weekly and monthly P&L

Figure 10: Fund panel impulse response functions for first and second half of year. The sample period is from 1/1/1995 to 1/1/2003. The sample is split into the first half of the year (H1) and the second half of the year (H2). To generate the impulse response functions, risk is regressed on weekly lags of risk, weekly lags of P&L conditional on gains and weekly lags of P&L conditional on losses for the fund panel, for each subsample. The regression allows for heteroskedasticity across

funds and uses the same lag structure as in the VARs in Figure 5. Sketched in lighter weight around each function are 90% confidence interval bounds based on maximum likelihood standard errors.

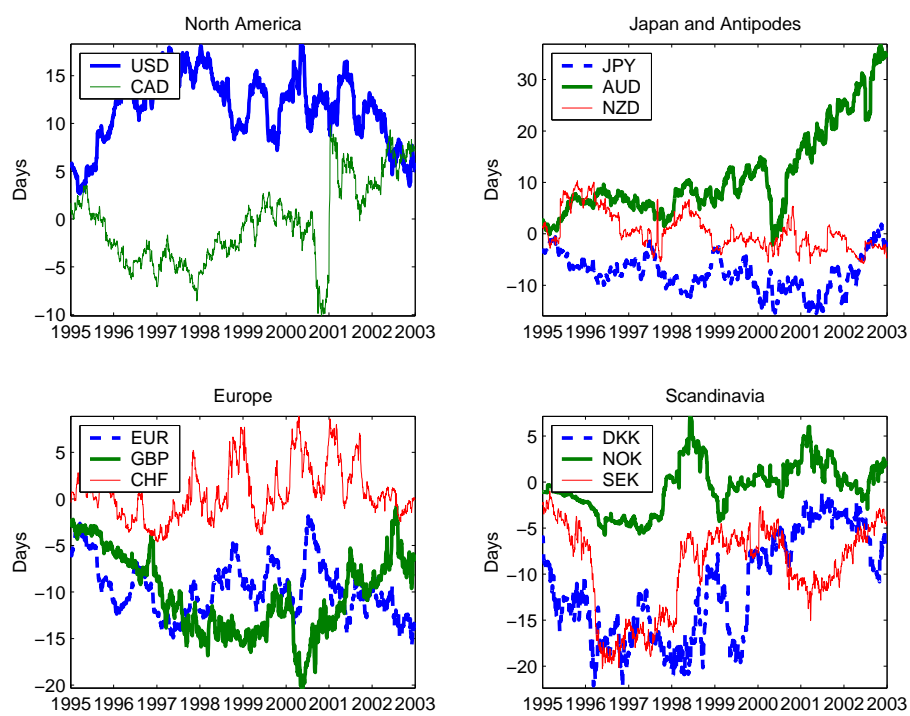


Figure 1: Holdings by currency aggregated across all funds, 1/1/1995-1/1/2003

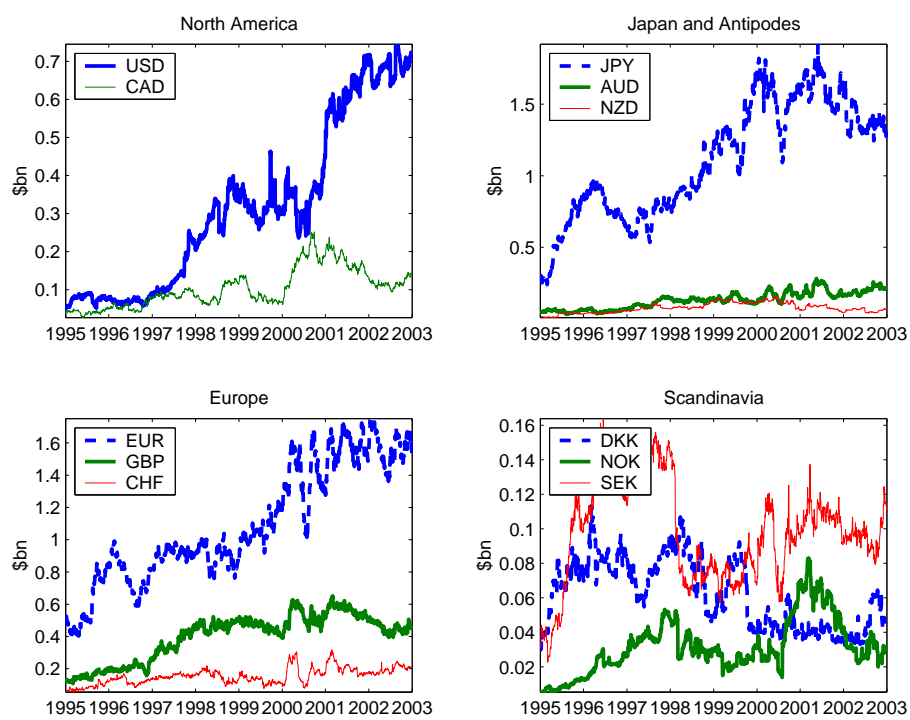


Figure 2: Risk exposure by currency aggregated across all funds, 1/1/1995-1/1/2003

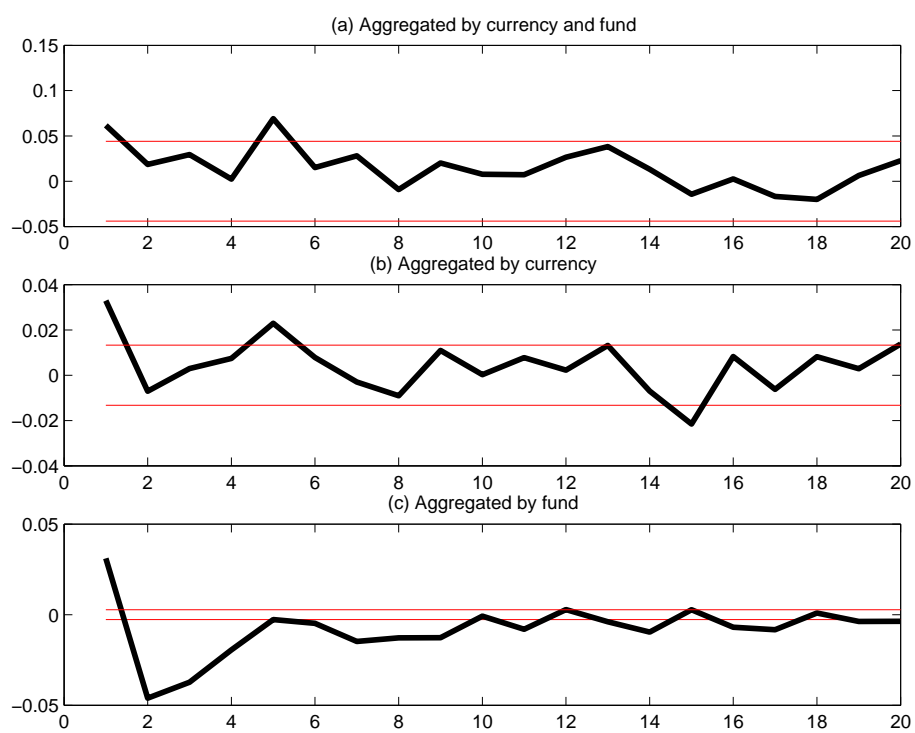


Figure 3: Daily sample autocorrelation function for risk aggregates. The sample period is from 1/1/1995 to 1/1/2003. Sketched in lighter weight are 95% confidence bounds.

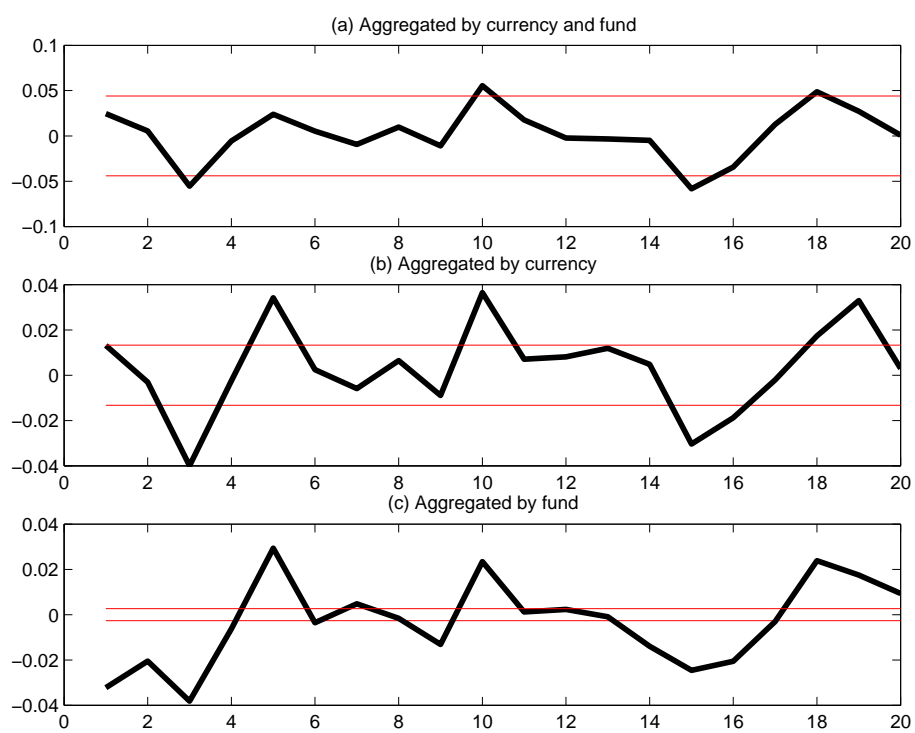


Figure 4: Daily sample autocorrelation function for P&L aggregates. The sample period is from 1/1/1995 to 1/1/2003. Sketched in lighter weight are 95% confidence bounds.

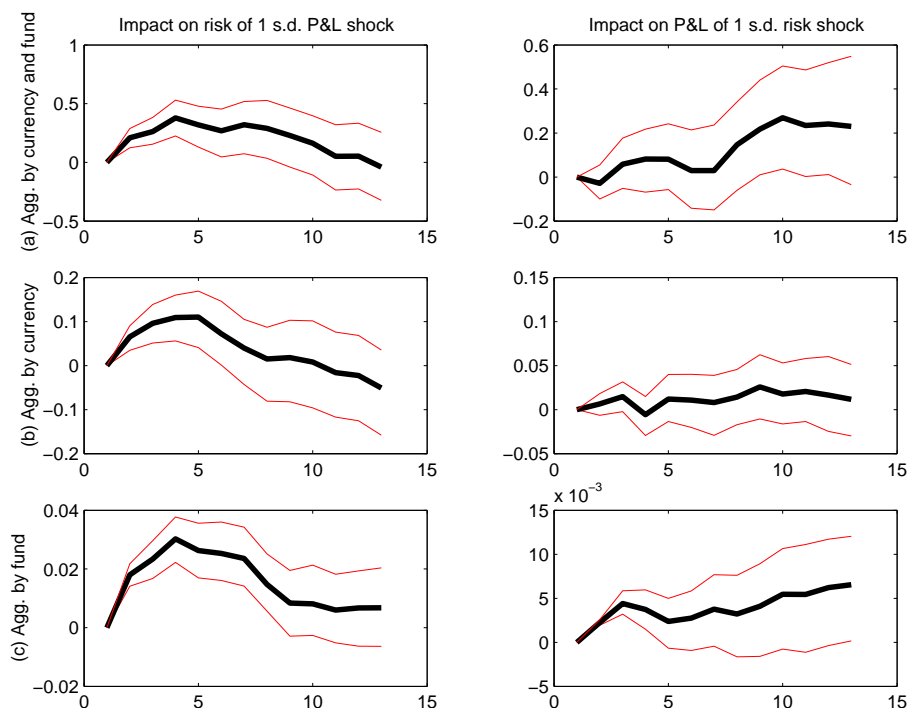


Figure 5: Impulse response functions for shocks to risk and P&L. The sample period is from 1/1/1995 to 1/1/2003. The sample includes ten currencies and 512 funds. The aggregated by fund and currency panel consists of a series of length T , where T is the length of the sample period. The aggregated by currency panel consists of 10 series of length T . The aggregated by fund panel consists of 512 series of various lengths, depending on fund life. To generate the impulse response functions, panel VARs for risk and P&L are estimated. The lag length of each VAR is set at 13 weeks, the value selected by the Bayes-Schwartz Information Criterion for the aggregated by fund panel regression. The VARs allow for heteroskedasticity across currencies and funds. The vertical axes are scaled in standard deviation units of risk. Sketched in lighter weight around each function are 90% confidence interval bounds based on maximum likelihood standard errors.

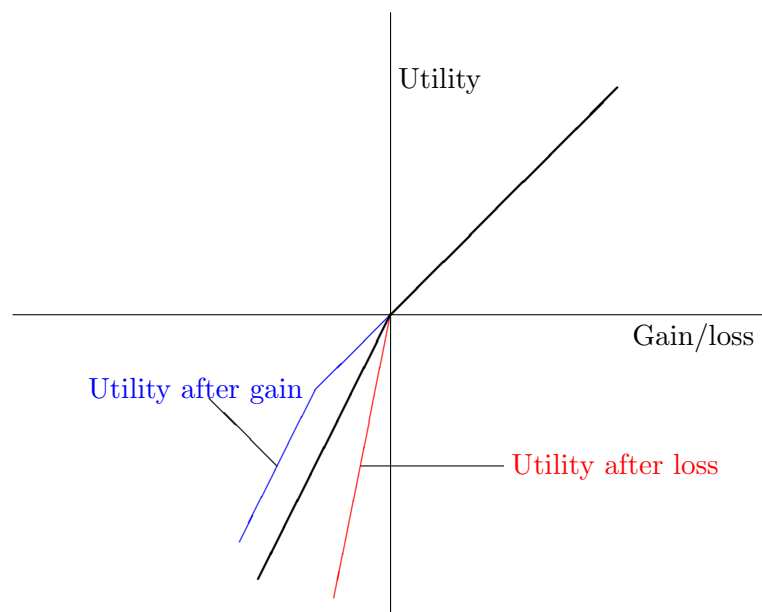


Figure 6: Barberis, Huang and Santos (2001) utility function for gains and losses

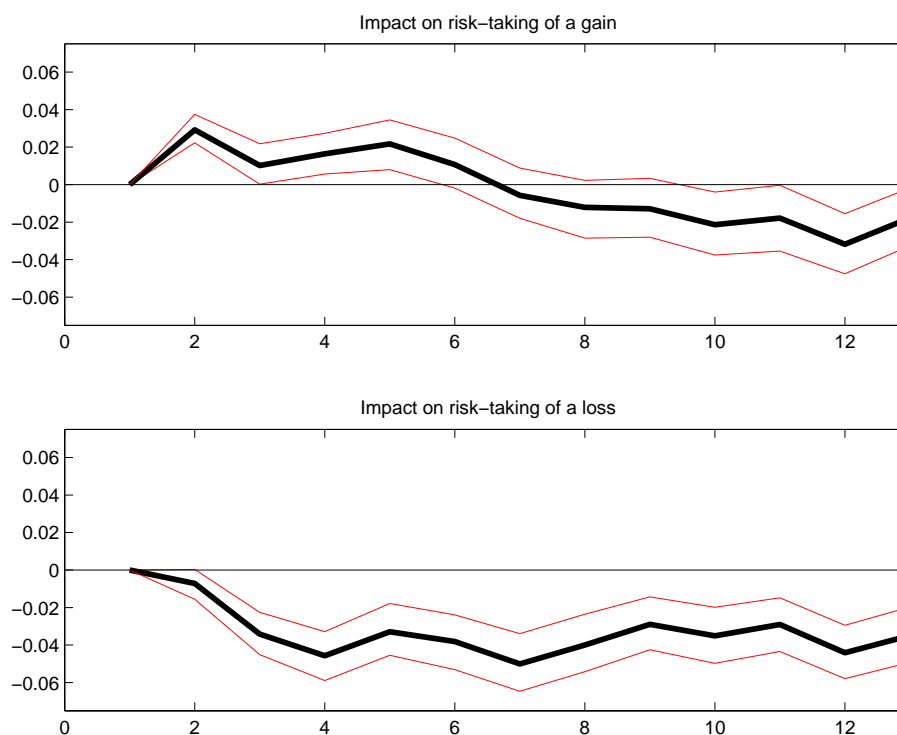


Figure 7: Fund panel impulse response functions for gains and losses. The sample period is from 1/1/1995 to 1/1/2003. To generate the impulse response functions, risk is regressed on weekly lags of P&L conditional on gains and weekly lags of P&L conditional on losses for the fund panel. The regression allows for heteroskedasticity across funds and uses the same lag structure as in the VARs in Figure 5. Sketched in lighter weight around each function are 90% confidence interval bounds based on maximum likelihood standard errors.

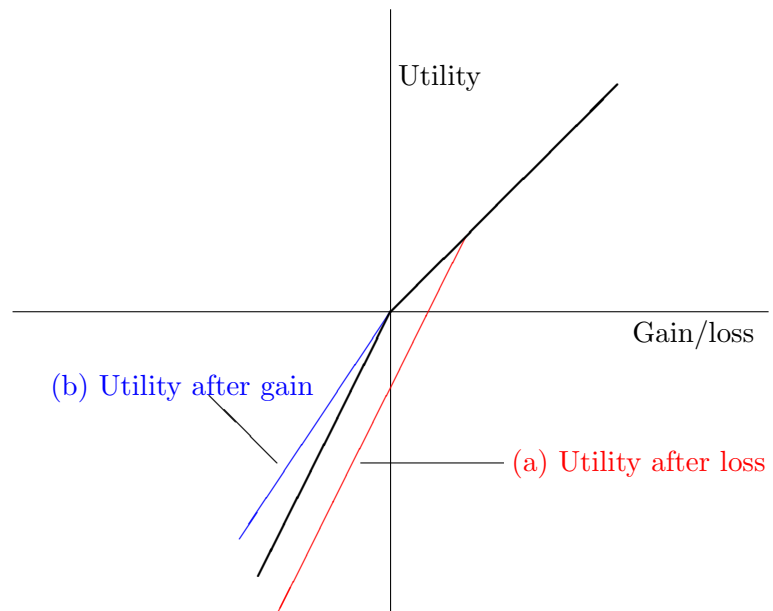


Figure 8: Barberis, Huang and Santos (2001) utility functions modified to fit the evidence

	Previous loss	Previous gain
Current loss	-32.2 (5.8) Weekly	-35.8 (5.6) Weekly
	-29.6 (5.9) Monthly	-37.5 (5.6) Monthly
Current gain	1.1 (5.8) Weekly	-4.1 (5.9) Weekly
	0.7 (5.9) Monthly	-3.3 (5.9) Monthly

Figure 9: Impact of P&L on risk-taking, conditional on past weekly and monthly P&L

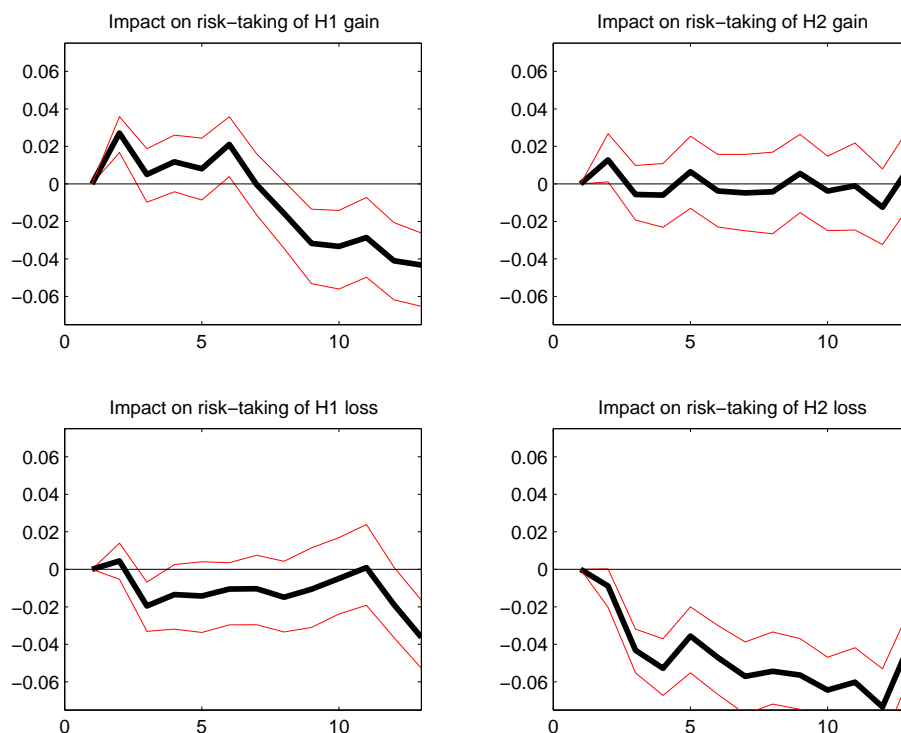


Figure 10: Fund panel impulse response functions for first and second half of year. The sample period is from 1/1/1995 to 1/1/2003. The sample is split into the first half of the year (H1) and the second half of the year (H2). To generate the impulse response functions, risk is regressed on weekly lags of risk, weekly lags of P&L conditional on gains and weekly lags of P&L conditional on losses for the fund panel, for each subsample. The regression allows for heteroskedasticity across funds and uses the same lag structure as in the VARs in Figure 5. Sketched in lighter weight around each function are 90% confidence interval bounds based on maximum likelihood standard errors.

Table I: Taxonomy of questions

Question	Rationale
<p>(a) Basic dynamics</p> <p>Is there performance dependence? What is its sign? Magnitude and duration?</p>	<p>Departure from (most of) rational finance paradigm</p> <p>Negative suggests disposition effects; positive supports overconfidence/dynamic loss aversion</p> <p>Economic relevance—is it large and long-lasting enough to explain major asset pricing anomalies?</p>
<p>(b) Conditional dynamics</p> <p>Does the impact of losses differ from impact of gains? Is effect conditioned on performance in prior periods? Are calendar effects present?</p>	<p>Disposition effect and dynamic loss aversion predict asymmetries</p> <p>Answers question of whether and how investors integrate outcomes over sequences of trades; speaks to theory of dynamic loss aversion (Barberis, Huang & Santos (2001))</p> <p>Incentives might alter managers' risk tolerance throughout the year</p>
<p>(c) Cross-sectional features</p> <p>Does effect vary across assets? Does effect vary across fund type? Other relevant fund characteristics? Cross-asset effects?</p>	<p>Is it pervasive?</p> <p>Equity, fixed income and currency funds may have different response functions</p> <p>Age and experience should temper overconfidence (Locke & Mann (2001))</p> <p>Narrow or broad framing—is it single asset or portfolio performance that matters</p>

Table II: Panel VAR estimates broken out by currency and fund-type

	P&L _{t-1}	P&L _{t-2}	P&L _{t-3}	P&L _{t-4}	P&L _{t-5}	P&L _{t-6}	P&L _{t-7}	P&L _{t-8}
(a) Currencies								
Denmark	-0.85 (1.86)	1.91 (1.85)	5.12 (1.85)	1.19 (1.85)	1.71 (1.89)	-2.31 (1.90)	-2.89 (1.88)	0.34 (1.88)
Norway	10.38 (8.67)	12.14 (8.67)	-0.79 (8.65)	3.91 (8.67)	-3.05 (8.70)	-0.58 (8.67)	7.61 (8.66)	24.94 (8.69)
Sweden	-6.27 (4.54)	-6.02 (4.52)	3.47 (4.51)	-10.15 (4.51)	0.24 (4.55)	5.18 (4.54)	-13.20 (4.52)	-2.34 (4.53)
Switzerland	9.40 (4.85)	2.99 (4.88)	12.58 (4.89)	10.01 (4.90)	4.84 (4.92)	4.13 (4.94)	5.62 (4.92)	3.00 (4.93)
U.K.	3.30 (4.67)	29.20 (4.67)	26.62 (4.66)	0.82 (4.46)	28.86 (4.44)	10.69 (4.44)	2.11 (4.45)	18.68 (4.41)
Australia	21.45 (5.59)	23.35 (5.58)	13.18 (5.55)	-12.75 (5.52)	-6.70 (5.49)	22.77 (5.48)	23.72 (5.41)	19.81 (5.47)
Japan	17.75 (4.24)	14.42 (4.24)	15.33 (4.25)	5.58 (4.24)	-2.01 (4.24)	-5.38 (4.25)	5.98 (4.24)	-7.62 (4.24)
New Zealand	-26.89 (5.68)	-6.90 (5.67)	-1.61 (5.71)	-1.48 (5.61)	-7.92 (5.62)	-20.91 (5.64)	-5.52 (5.60)	-5.63 (5.60)
Canada	10.64 (6.26)	13.87 (6.25)	10.32 (6.23)	1.78 (6.21)	-17.36 (6.21)	-8.61 (6.19)	-5.16 (6.19)	-5.70 (6.19)
Euro	19.64 (4.27)	10.80 (4.26)	8.23 (4.26)	-3.94 (4.27)	-5.32 (4.28)	1.45 (4.28)	-7.07 (4.26)	-12.67 (4.26)
All	16.61 (2.67)	7.20 (2.67)	9.87 (2.68)	-1.61 (2.68)	-1.99 (2.68)	-0.56 (2.68)	-8.20 (2.67)	-6.53 (2.67)
(b) Fund types								
FX funds	14.99 (4.73)	6.15 (4.72)	6.16 (4.73)	-6.63 (4.74)	-12.70 (4.74)	1.90 (4.74)	-12.95 (4.72)	-12.01 (4.73)
Stock funds	-8.18 (10.72)	2.32 (10.71)	-12.62 (10.69)	-19.77 (10.68)	-4.98 (10.67)	24.38 (10.67)	16.43 (10.66)	9.83 (10.66)
Bond funds	13.41 (3.96)	4.73 (3.96)	12.37 (3.98)	2.84 (3.98)	3.58 (3.98)	-5.19 (3.98)	-8.36 (3.97)	-6.34 (3.96)

This table shows the first eight coefficients on lagged P&L from a 13-lag bivariate panel VAR for risk and P&L. The sample period is from 1/1/1995-1/1/2003. The model allows for heteroscedasticity across funds. Estimation is carried out by maximum likelihood, stacking all of the funds in the sample. Standard errors are in parentheses. The estimates are shown by currency and by fund type.

Table III: The effects of age and experience on performance dependence

	$\beta_{t-1}^{P\&L}$ Total	$\beta_{t-1}^{P\&L}$ Gains	$\beta_{t-}^{P\&L}$ Losses
Age	-1.71 (0.91)	-5.72 (1.56)	2.49 (1.39)
Experience	-1.36 (0.65)	-3.47 (1.14)	0.67 (1.02)

This table illustrates the effect of age and experience on performance dependence. The sample period is from 1/1/1995 to 1/1/2003. The sample is split into two sub-periods. An evaluation period (the last three years of the sample, 1/1/1999-1/1/2003) and a formation period (the initial four years of the sample, 1/1/1995-12/31/1998). A fund's age is proxied by the length of time since the first day of trading in our sample, and experience proxied by the number of days trading. To test sensitivity to these two variables, we use a simple two-step procedure. In step one, the sensitivity of each fund to lagged P&L is measured across the evaluation period for all funds which exist in that period. Then in step two, the cross-section of coefficients is regressed on age and fund experience. Standard errors in parentheses.

Table IV: Conjugate P&L results by currency

	Conj. P&L _{t-1}	Conj. P&L _{t-2}	Conj. P&L _{t-3}	Conj. P&L _{t-4}	Conj. P&L _{t-5}	Conj. P&L _{t-6}	Conj. P&L _{t-7}	Conj. P&L _{t-8}
Denmark	0.03 (0.03)	-0.02 (0.03)	0.03 (0.03)	0.07 (0.03)	-0.02 (0.03)	0.05 (0.03)	0.03 (0.03)	0.04 (0.03)
Norway	-0.05 (0.04)	0.05 (0.04)	0.02 (0.04)	0.08 (0.04)	0.01 (0.04)	0.14 (0.04)	0.00 (0.04)	0.05 (0.04)
Sweden	0.15 (0.04)	-0.04 (0.04)	0.03 (0.04)	0.08 (0.04)	0.01 (0.04)	0.06 (0.04)	0.02 (0.04)	0.07 (0.04)
Switzerland	-0.01 (0.03)	-0.05 (0.03)	0.05 (0.03)	0.04 (0.03)	-0.03 (0.03)	0.01 (0.03)	0.03 (0.03)	0.06 (0.03)
U.K.	0.04 (0.09)	0.01 (0.09)	0.06 (0.09)	0.01 (0.09)	0.00 (0.09)	0.07 (0.09)	-0.03 (0.09)	-0.02 (0.09)
Australia	0.17 (0.10)	0.14 (0.10)	0.14 (0.10)	0.12 (0.11)	-0.01 (0.11)	0.32 (0.10)	0.12 (0.10)	0.01 (0.10)
Japan	0.79 (0.62)	-0.33 (0.62)	-0.89 (0.62)	0.66 (0.62)	-1.26 (0.62)	-0.46 (0.62)	-0.32 (0.61)	0.51 (0.62)
New Zealand	0.01 (0.02)	0.02 (0.02)	0.01 (0.02)	-0.01 (0.02)	-0.01 (0.02)	0.01 (0.02)	0.01 (0.02)	0.00 (0.02)
Canada	0.02 (0.07)	-0.02 (0.07)	0.10 (0.07)	0.10 (0.07)	-0.04 (0.07)	-0.01 (0.07)	-0.01 (0.07)	-0.01 (0.07)
Euro	-0.68 (0.22)	-0.91 (0.23)	-0.05 (0.27)	0.47 (0.24)	-0.81 (0.26)	1.17 (0.28)	0.34 (0.29)	-0.20 (0.27)

This table shows the first eight coefficients on lagged *conjugate* P&L from a regression that also lagged risk and lagged own P&L as regressors. The sample period is from 1/1/1995 to 1/1/2003. Full regression results are available from the authors on request. Conjugate P&L is defined as the profit or loss on all currencies except the regressand currency. Estimation is carried out by maximum likelihood, stacking all of the funds in the sample. Standard errors are in parentheses.