DYNAMIC RISK CONTROL FOR EQUITY PORTFOLIOS

WERNER E. KELLER, CFA
KELLER PARTNERS LLC

APRIL, 2008

Key Concepts

Too Much Risk. In the late afternoon of a twenty-year bull market in financial assets, most portfolios are significantly overexposed to equity risk, simply because the most widely-practiced portfolio allocation methods are, in the final analysis, slow-moving momentum strategies. Predictably, a "normal" allocation to equities will always be larger at the end of a bull market than at the beginning, and the best-performing asset class over the last five years will tend to have the greatest weight.

Given the compression that equity and bond returns have experienced over the last seven years, the retirement marketplace should be keenly interested in strategies that focus on risk control (drawdown control) and, when possible, enhancement of returns. A tall order, but certainly a goal worth pursuing. Managers wishing to implement such an active risk control process for their equity portfolios require: (1) access to an elegant, low-cost hedging mechanism, and (2) a robust set of decision rules to control the periodic activation of the hedge. This paper proposes approaches for both requirements.

An Elegant Hedging Mechanism. The recent launch of Leveraged, Inverse Exchange-Traded Index Funds is an important development. With these tools, portfolio managers can now implement highly efficient, low-cost hedges in portfolios of virtually any size. Since these specialized ETFs are both inverse and leveraged, even a modest portfolio allocation (15% - 20%) can significantly reduce the market exposure (beta) of a portfolio when desired. These ETFs represent a breakthrough in portfolio construction that is only beginning to be understood by the marketplace.

Decision Rules. A dynamic hedge is by definition sometimes active and sometimes not, ideally at all the right moments. While we have certainly not found that to be possible, we do believe that it is possible to identify intermediate trends in financial markets — not to predict them, but to identify them, based on weak and subtle — but statistically demonstrable — momentum effects (trending tendencies) that have existed, and persisted, in world equity markets for a long time. We believe it is possible to construct quantitative, rules-based, models that can keep the hedging decision on the right side of an intermediate trend perhaps 60%-65% of the time. This advantage can exert a major influence on the risk-adjusted performance of an otherwise unhedged portfolio.

This paper introduces such a trend-identification model whose output can control hedge decisions on portfolios correlated to the S&P 500, the Nasdaq 100, the Russell 2000, or a combination of the three. We also provide some hypothetical illustrations of how this model might have performed, making use of these new inverse ETFs (had they been available) for an active hedging regime over the last 11 years.

Dynamic Risk Control, like a balanced diet, is not a controversial concept. Even in favorable market environments, the possibility of providing a retirement portfolio with a smoother financial ride is an inviting idea, given the universal sensitivity of investors to drawdown, and the growing sense that stock and bond returns over the next twenty years will lag far behind those of the past twenty.

In the first three months of 2008, the key US equity indices experienced maximum drawdowns of 15% – 20%, well in excess of the 10% comfort zone that we believe to be almost universal among investors. Once again, portfolio managers are being challenged to become more active with their risk control strategies, as well as in their pursuit of alpha opportunities on the long side. Without such efforts, we expect these managers will have difficulty delivering return patterns that meet their clients' expectations.

This paper describes an active risk management process based on —

(1) An efficient hedging mechanism that is scalable and easily understood. The mechanism we propose is implemented with the innovative **Leveraged Inverse Exchange-Traded Index Funds** recently introduced by *ProFunds* and *Rydex.* We believe that they are tailor-made for these requirements, and that they represent a significant breakthrough for the replication of institutional hedging strategies in individual portfolios.

(2) A rules-based hedge management system that determines when a hedge should be applied. This dynamic hedging process should significantly reduce volatility — especially downside volatility (drawdown) — without penalizing performance. Secondly, a successful long/short market model should exhibit some added value versus its benchmark. In other words, beyond simply truncating the size of the negative returns and reducing the standard deviation of monthly returns, it would also be welcome if the hedging decisions contained enough "alpha" to enhance the investment performance of the unmanaged index.

LEVERAGED INVERSE INDEX ETFS

Inverse Funds have market movements that reflect the reverse of the movement of an underlying index (such as the Standard & Poor's 500). Thus, on a day when the index goes down, the fund rises by a similar percentage and vice versa. In Wall Street terms, these funds are a *de facto* short position on the index. When they are leveraged 2:1, as these new *ProShares* and *Rydex* ETFs are, the movement of the underlying index is magnified by a factor of two.

Unlike the similar inverse open-end mutual funds introduced by these two fund companies over ten years ago, these new ETFs trade continuously on a listed exchange.

The built-in leverage is especially important since it permits a small position in the fund to hedge a significant amount of market exposure in a portfolio.

ProFunds launched its initial offerings of leveraged and inverse index ETFs in June and July of 2006 under the *ProShares* ETF label, and has experienced dramatic market acceptance of their products, particularly the 2-beta short ETFs that (inversely) track the Nasdaq 100 (NDX) and the S&P 500.

Their 2-beta long counterparts have experienced asset growth also, but much less significantly, supporting our view that a serious constituency is now developing for these short ETFs as hedging vehicles. As the table below summarizes, the two dominant *ProShares* 2-beta short ETFs (SDS and QID) have accumulated very sizeable assets in the 22 months of their existence. On volatile days, their trading volume each exceeds $2 billion of daily market value ($4 billion in terms the underlying index), and bid/asked spreads have become very narrow. In recent weeks, the inside bid/asked spreads for the *ProShares* double inverse S&P ETF (SDS), as well as the corresponding inverse Nasdaq 100 ETF (QID) have narrowed to one cent (e.g., bid $45.13 / asked $45.14) on underlying share prices of roughly $60 and 45, respectively. Thus, spread slippage is almost trivial, especially in light of the significant daily percentage movements of these leveraged instruments.

Figure 1

2-Beta Inverse Exchange-Traded Index Funds

[$ millions]

	ProShares		Rydex	
Underlying Index	**Symbol**	**Assets**	**Symbol**	**Assets**
S & P 500 Composite	SDS	$ 2.920	RSW	$ 47
Nasdaq 100	QID	$ 1,610	n/a	$ -
S&P Midcap 400	MZZ	$ 175	RMS	$ 14
Russell 2000	TWM	$ 646	RRZ	$ 33

2-Beta Long Exchange-Traded Index Funds

	ProShares		Rydex	
Underlying Index	**Symbol**	**Assets**	**Symbol**	**Assets**
S & P 500 Composite	SSO	$ 802	RSU	$ 8
Nasdaq 100	QLD	$ 947	n/a	$ -
S&P Midcap 400	MVV	$ 64	RMM	$ 3
Russell 2000	UWM	$ 108	RRY	$ 5

Reported Assets in the Funds as of: March 14, 2008 ($ millions)

RAPID GROWTH IN ASSETS. In the brief period since their introduction in mid 2006, *ProShares*' Leveraged Inverse ETFs have established themselves among the most liquid ETFs in the marketplace. The Rydex counterparts were second-to-market late last year, and are just beginning to grow in assets. Given the ultimate size of the market for hedging strategies that we envision, we expect both to do well.

Rydex Investments launched a parallel group of ETFs products in November of 2007 which, given their short history, have yet to develop significant trading markets. Trading symbols for the most important of these 2-Beta index ETFs are provided in Figure 1.

The marketplace is beginning to recognize that using leveraged inverse ETFs for hedging long-only portfolios is significantly less complex and much more efficient than attempting similar strategies with options or futures, particularly for portfolios under $500,000.

CONSTRUCTING HEDGED PORTFOLIOS WITH 2-BETA INVERSE ETFS

Given the leverage built into 2-beta ETFs "at the factory" as it were, the investment portfolios being hedged can be kept largely undisturbed, since only a modest commitment to a 2-beta short ETF will significantly reduce the market sensitivity (beta) of the underlying portfolio.

In constructing a hedge policy, we must begin by measuring the real world CAPM statistics[1] of the underlying portfolio, which can be estimated in an Excel spreadsheet. As an illustration, we will use Fidelity's *Asset Manager 70% Fund* (FASGX), one of Fidelity's single fund portfolio "solutions." This long-established fund is diversified across all asset classes, and incorporates slow-moving adjustments in asset class weights over time.[2] For our purposes, the fund also offers the advantage of a long track record of daily portfolio values (fund NAVs) on which we can base our hedging evaluations, and it also (realistically) has embedded annual fees.

We begin by aligning the daily price changes for the S&P 500 (x-axis) with the daily price changes of FASGX (y-axis), plotting an X-Y graph and asking Excel to calculate a linear regression line and to display the formula, as illustrated in Figure 2. Since we use daily data pairs, the intercept value (0.0000815) needs to be annualized to arrive at the calculated alpha value (+2.075%).

Clearly, with an r-squared of 0.94, most of this fund's return is explained by the S&P 500, making it an excellent candidate for a dynamic hedge using an inverse S&P 500 ETF such as the *ProShares* 2-beta SDS. Since we have a large number of data pairs for the portfolio to be hedged and the benchmark, the hypothetical impact of the hedge can be calculated with considerable confidence.

---

[1] CAPM – The Capital Asset Pricing Model, one of the cornerstones of Modern Portfolio Theory

[2] Fidelity's description of its Asset Manager 70% Fund follows —

> *Fidelity allocates the fund's assets among stocks, bonds, and short-term and money market investments . . . . around a neutral mix of 70% stocks (can range from 50-100%), 25% bonds (can range from 0-50%), and 5% short-term/money market (can range from 0-50%). FMR regularly reviews the fund's allocations and makes changes gradually to favor investments that it believes will provide the most favorable outlook for achieving the fund's objective.*

Figure 2

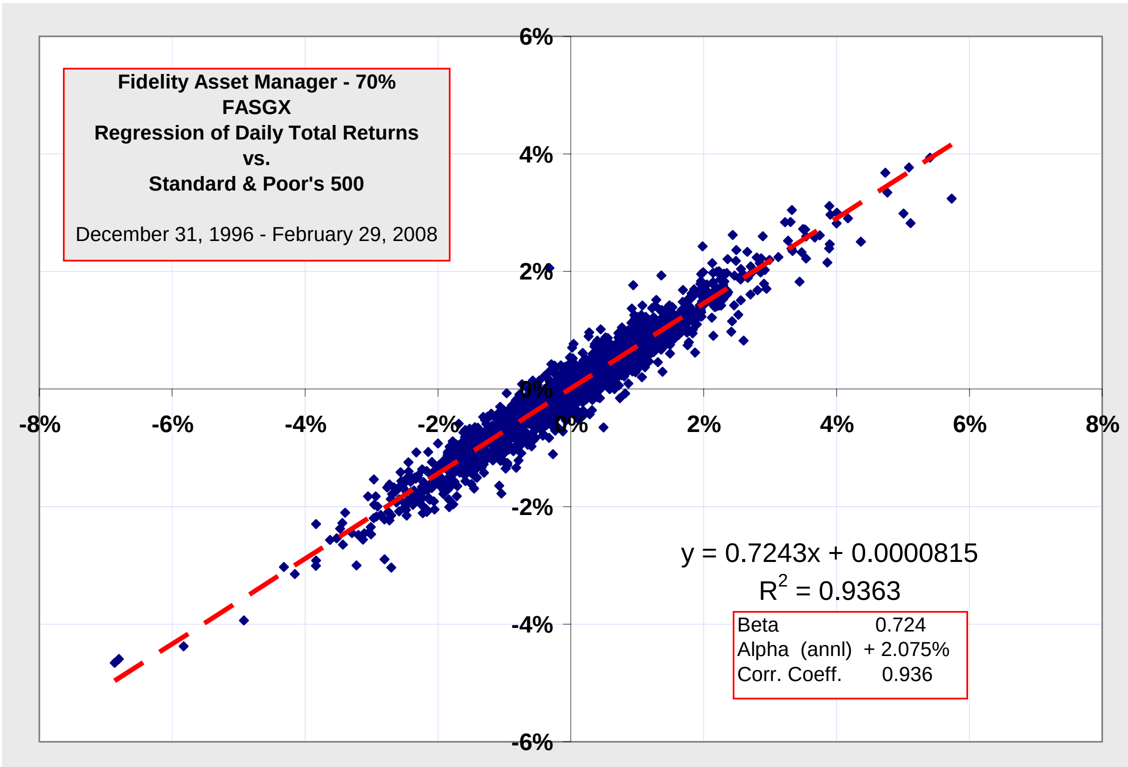

ESTIMATING BETA. Daily changes in the benchmark Standard & Poor's 500 Composite index on the horizontal axis, versus the daily total return of the *Fidelity Asset Manager 70% Fund*. As is the case with many of Fidelity's funds, the fit is quite close and statistically highly predictable.

When it comes to determining the degree to which the manager wishes to hedge the portfolio in question, we have suggested to most organizations that they consider under-hedging their portfolios, with a typical hedge allocation of 15% - 20% (using a 2-beta short ETF). This creates a meaningful, but not complete, reduction of systematic (market) exposure when the hedge is applied, keeping the portfolio at least a little net long all of the time.

In the case of FASGX, since its real-world beta was measured at 0.72, a bit of algebra tells us that a combination of 80% of the fund portfolio and 20% of a -2.00 beta ETF will blend the portfolio's market exposure to the S&P down to a beta of 0.18 whenever the hedge is invoked — this amounts to a very meaningful 75% reduction in the portfolio's exposure to systematic (non-diversifiable) risk. The mathematical relationship between the hedge commitment and the likely extent of systematic risk reduction is linear and is illustrated in Figure 3.

Figure 3

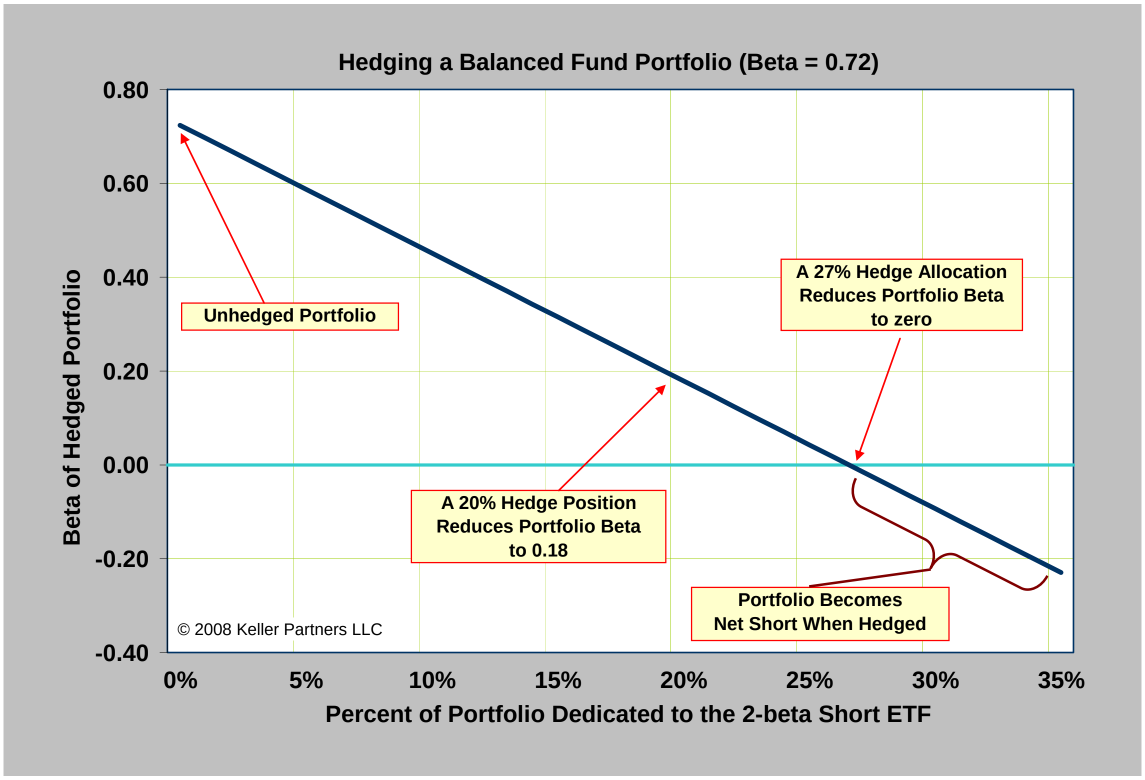

THE ALLOCATION TO THE HEDGE. When we introduce a small position of a security with a beta of negative 2.0 into a portfolio, the resulting reduction of the market exposure of that portfolio can be quite dramatic. Once the hedge represents 27% or more of this Fidelity fund portfolio, the hedged investment becomes net short, opening the door to replicating long/short strategies (with absolute return patterns) in portfolios of any size.

We should acknowledge that all of these calculations are really target range estimates, since, in real life, earnings announcements, economic news, the non-stationarity of betas, etc. all introduce statistical noise around these metrics. In general, however, the calculations are quite reliable over time.

Figure 4 illustrates our simulation of this fund's long-term total return performance with the addition of a dynamically-controlled hedge. Here we assume that the investment manager (perhaps Fidelity itself) has an effective way of reducing the portfolio positions by 20% *pro rata* whenever the hedge is called for, and purchasing one of the leveraged, inverse S&P ETFs with the proceeds. The dynamic hedge in this example is controlled by our firm's market trend model, whose characteristics we will discuss below.

Figure 4

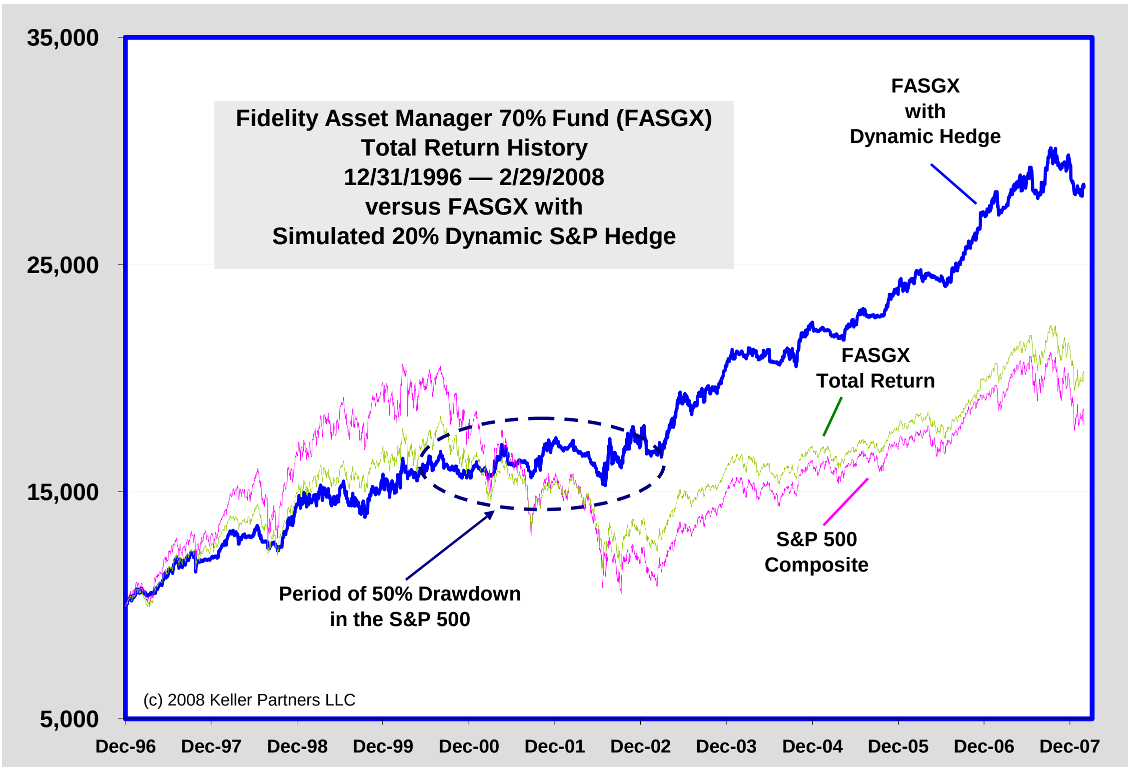

MUTUAL FUND PORTFOLIO WITH SIMULATED 20% DYNAMIC HEDGE. This chart examines the hypothetical impact of periodically shifting 20% of the portfolio's assets to an index ETF that inversely tracks the S&P 500 with a beta of 2.0. The timing of the hedge is controlled by our trend-following market model.

As Figure 4 suggests, considerable value could have been added in the 2000-2002 bear market, when the 11% maximum drawdown of the hedged fund investment (circled on the chart) was dramatically less than the 49% (!) MDD experienced over the same two-year period by the S&P 500 Composite[3]. On sunnier days in the markets, the hedge also needs to be nimble enough to get out of the way, as not to impede the performance of the underlying portfolio.

The central conclusion is that a market trend-sensing algorithm — even one that is slow-moving and only capable of identifying trends correctly "most of the time" — can still have a very meaningful impact on both portfolio drawdown and return over time.

Twenty years ago, there was considerable academic discussion about blending equity portfolios with managed futures (the optimal percentage allocation to futures was about 30%). In essence, these discussions recognized that combining two uncorrelated asset classes, each of which adds a little excess return (alpha) in its own way, can produce

[3] This dramatic reduction in drawdown comes from the combination of a relatively low-volatility fund (beta = 0.72) and a rather aggressive hedge allocation (20% in the inverse S&P ETF).

almost magical results for the investor — less volatility and higher IRRs. From this perspective, a dynamic inverse S&P position — really, an actively-managed index short — is the ultimate non-correlated asset class.

The market model that controls the hedge needs to be adequate to the task, but it does not need to perfect, since there is usually a second alpha generator working within the portfolio — the portfolio management process. The typical portfolio manager contributes alpha erratically, but his or her contribution pattern is very likely to be independent of the success pattern of the hedge algorithm.

Since even well-designed market models are erratic, especially over short time periods, they can be exceptionally frustrating if they are required to be the sole potential source of alpha[4]. This principle lies behind much of the traditional disenchantment with market timing-based managers. However, when this erratic process is paired with another, uncorrelated, value-added investment process — in this case, what we presume will be an intelligently-managed portfolio — the combination can become quite synergistic.

As a final note, from the chart in Figure 3 it is clear that one can also choose to hedge more aggressively, to the point where the portfolio becomes net short when the hedge is "on," thus crossing an important policy line, where its return pattern is more appropriately compared to a long/short fund and other absolute return strategies. We do expect that these kinds of portfolio structures will develop significant long-range appeal.[5]

IMPLEMENTATION PROTOCOLS

We will take a few moments to consider exactly how we would go about transforming a fully-invested, "normal" portfolio into a dynamically hedged investment: how the investment might be structured when the hedge is on and what it looks like when the hedge is off. There are several approaches, each with advantages and disadvantages.

I. Proportional Reallocation Approach. The first approach is the one we followed in the FASGX simulation above. It is perhaps the least elegant of the process alternatives due to its transactional intensity. When the hedge is called for, the portfolio is scaled back proportionately to raise funds for purchase of the inverse ETF position (in this example, 20%). This approach might be appropriate for retirement portfolios that are not in the position to borrow, are indifferent to tax events, and have access to low-cost transactions.

[4] This is why our friends at Theta Investment Research (www.thetaresearch.com), an organization that tracks a large number of active managers based on their ability to make timely shifts in index and sector positions, suggest that value-added active managers account for less than 10% of the total.

[5] Our firm is testing a logical extension of that concept with an ETF portfolio that becomes net short when the hedge is invoked, but shifts the hedge allocation to leveraged long position when the hedge is inactive (by rebalancing and switching the hedge to the corresponding long 2-beta Index ETF).

Unhedged Portfolio

Hedged Portfolio

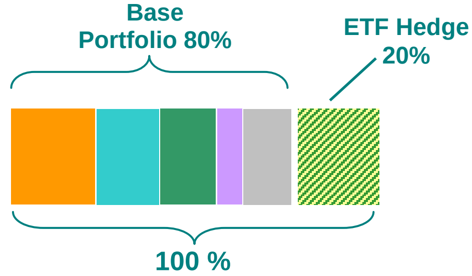

Pros

- The hedge is applied to the whole portfolio — the character of the underlying portfolio is unaltered.
- More appropriate to retirement portfolio constraints.
- No requirement for temporary borrowing.

Cons

- Reallocations create transaction costs.
- Reallocations create tax events in taxable portfolios.

II. Dedicated Hedge Allocation. The underlying "long" portfolio is constructed at the outset to include a fixed-income allocation that is equal to the policy allocation to the inverse ETF hedge. When the hedge is called for, only this fixed income portion (grey in the schematic below left) is converted to the hedge position. The other portfolio positions are untouched.

Unhedged Portfolio

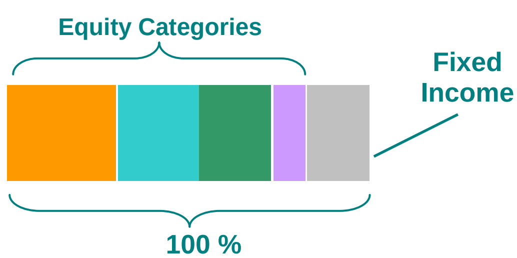

Hedged Portfolio

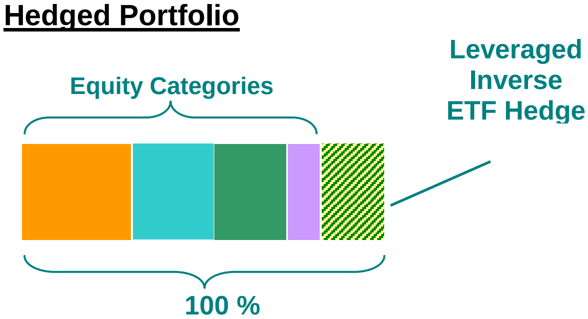

Pros

- Much more tax efficient, as the bulk of the allocation is left unaffected by the movement in the hedge.
- Most plausible structure for dynamically-hedged mutual funds, variable insurance trusts, or other collective investment vehicles.
- No requirement for borrowing.

Cons

- The basic asset allocation rules (investment policy statement) need to be altered unconventionally to facilitate efficient implementation of the hedge.

III. Margin. Perhaps the mechanically most elegant structure is one where the portfolio manager is able to borrow to purchase the hedge position. This allows the core portfolio to remain completely untouched while temporary borrowings (e.g., margin) are taken on to fund the hedge as needed. This structure would seem appropriate for separate accounts, and has especially favorable implications for taxable portfolios. It is also possible that it could be implemented in an open-end fund structure.

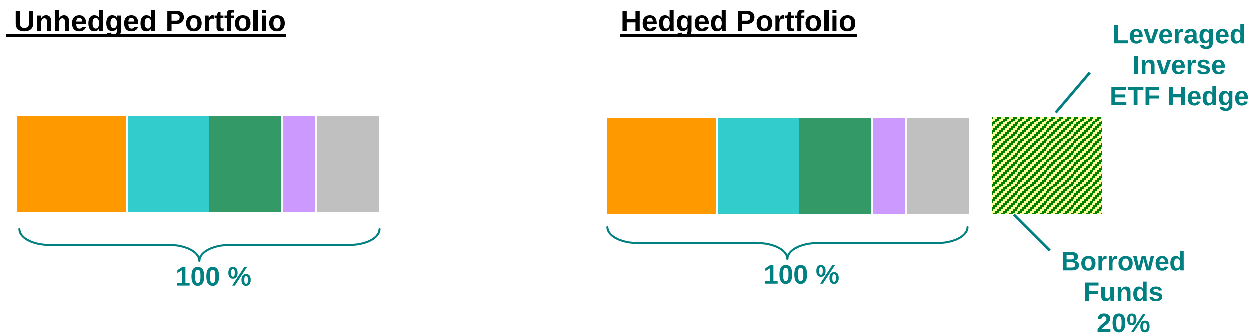

Pros

- Most tax efficient, as the entire underlying portfolio process remains unaffected by the insertion / removal of the hedge.
- Appropriate structure for separate accounts, or any entity that can borrow and repay seamlessly.
- Minimizes transaction costs.

Cons

- Cost of borrowing, although this is obviously only applied when the hedge is invoked, when, presumably, there is some work to be done.
- Might be complex to replicate structure in a mutual fund or variable trust.

THE KPLLC TREND ANALYSIS ALGORITHM

Our organization has an ongoing professional focus on developing proprietary market models, as well as a keen interest in the insights of others in this area. Over time, we have reached some conclusions about approaches that may add value and those whose periodic success, on closer evaluation, seems to be mostly random.[6] We have concluded that a small amount of useful trend information is indeed present in the daily flow of market data, and that it can be filtered from the vast amount of statistics thrown off by the financial markets every day. This information can be distilled and brought to

[6] Mathematician and speculator Nassim Nicholas Taleb wrote an excellent book a few years ago examining the difference between genius and long runs of good fortune, titled "*Fooled by Randomness.*"

productive use. As we shall discuss more fully, the kernels of information that are in fact present are subtle and easily overlooked.

General Design Criteria. A good hedging model should not trade too frequently, although it needs to be responsive enough to market movements to be on the right side of all the more significant price trend events in the market. The buy/sell decisions of a value-added model should, over time, make a measurable contribution to drawdown reduction, which is by far the most important risk metric for the individual investor. Ideally, there would be some improvement in performance (IRR) as well. In other words, the Sharpe ratio of the hedged portfolio should be materially higher than that of the unhedged portfolio.

Model Construction. The KPLLC model is grounded in our belief that internal, or secondary, market data contain more exploitable trend information than is contained in a sequence of historical prices. This is where, in our opinion, many academics have it wrong: security or index price changes considered on their do appear almost random[7]. Even when we shift our focus to the secondary data series that we consider more promising, we find (a) that the information is still embedded in considerable statistical noise, and (b) that one has to discover which of the secondary time series appear to contain the most useful information. Figure 5 maps some of the possible data inputs for a model, and illustrates the tradeoffs we have observed between higher predictive content and greater (tighter) connection to portfolio prices.

Optimization. The construction of any market trend model necessarily requires the builder to test theories and smoothing parameters on historical data in an attempt to discover the underlying personality, or "signature" of a particular market, assuming of course that there is one to be discovered. One always begins with a conceptual framework — an idea of how the trends of a particular financial market might be analyzed and what factors and proxies one might incorporate, followed by rigorous day-by-day bench-testing with historical market data to investigate how the model would have performed in the past. The obvious danger is that the investigator gets carried away and over-optimizes this stage of the process to the point where the model explains the historical market swings perfectly but, not surprisingly, loses all of its (already fragile) predictive powers.

There are several defenses against this data-fitting error: (1) independent time periods — it is always important to test any algorithm against data other than the data against which it was calibrated; (2) test the robustness of the conceptual design by mapping the sensitivity of the IRR of the model to changes in its key parameters — it should be relatively insensitive to small changes; and (3) keep the construction of the model simple and have it make intuitive sense to begin with — a model based on the weekly rainfall in Minneapolis should not be expected to add value in the long run.

[7] A recent study published by the London School of Economics has documented significant long-term momentum effects in global equity prices. A summary of that paper is available for download at: http://www.london.edu/assets/documents/786_GIRY2008_synopsis.pdf. While we would still categorize price momentum among the weaker data series with less predictive value, these studies are encouraging in terms of what they suggest about the potential of statistical models in general.

Figure 5

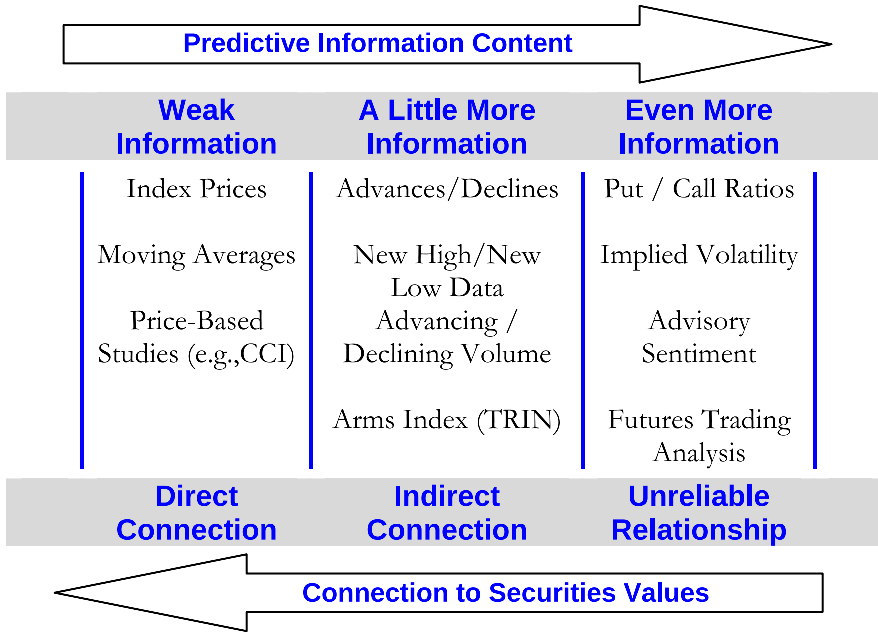

KNOWING WHERE TO LOOK. If our task is to identify the trend of a particular financial market, not to predict it, but identify it, we can certainly study historical prices, as many analytical techniques do (with moving averages, for example). A more evolved approach might also consider other forms of market data for possible trend information. The graphic above summarizes our conviction that a robust market model needs to incorporate data series beyond historical prices.

Design Concepts underlying the KPLLC Equity Trend Model

- Price tops have a different structure than price bottoms, and thus need to be analyzed differently by a competent model.
- As discussed above, the greatest information content for forecasting future prices is often not found in prices, but in data series that are once removed, such as advances and declines, or even further removed such as put/call ratios or implied options volatility.
- The securities markets are fractal, but not completely so. The presence of very large players has inhibited the fractal structure of intra-day movements and made them more random.
- Contrary to common sense, no one has ever demonstrated an exploitable connection between economic data and stock prices. We incorporate no economic inputs.
- Any good algorithm needs to have adaptive design features, since the internal rhythms and characteristics (the typical “signature”) of financial markets change over time. This has been especially the case recently.
- Less is more. Simple formulas with simple parameters have a far greater chance of adding value over time than complex, multi-factor models.
- An underlying characteristic of any value-added model is robustness: where more value comes from the design of the decision model than from the precise value of any single formula parameter. Enhanced performance over time comes from many small statistically-favorable decisions, not from getting an isolated big event right.

The historical shift points of the KPLLC model are presented graphically in Figure 10 on the last page, and summarized statistically on the next page. A few observations on our model:

- The model covers over 11 years of market history, including both sides of the tech stock mania. While the period is certainly interesting, it may not be typical. The model issued 56 changes of direction in that period, roughly five per year. The optimization period covered five years — thus there are six statistically independent years in the period under study.

- The model periodically arrives at a state of indecision where it reverses several times over a few weeks. Inevitably, these short holding period trades will generate losses — with hindsight, it would have been more productive to do nothing. However, it is critical in any regime of this sort that all signals be followed religiously, and that false signals be embraced as the cost that managers and their clients are required to pay in order to be on the right side of major market trends, most of the time.

- The intervals where the model earns its keep are those when the benchmark (S&P 500) declines and the model (correctly) hedges all or much of that decline. While our model does accomplish that over time, it does not accomplish it every time, and while the months where it doesn't quite "get it right" are, as we shall see, statistically very acceptable, they can seem interminable when one is actually living through them day by day. For this reason, having such a model play a collaborative support role in the portfolio construction process is well-advised. As we've already mentioned, this usually means selecting a hedge ratio that keeps the portfolio slightly net long, even when the hedge is active.

Bench Test. We begin our evaluation by asking how the return patterns of a long position in the S&P 500 index might have been altered / improved by the addition of a dynamic hedge steered by this particular model. How effective an alpha generator is the model on its own?

Figure 6 summarizes the results of this exercise: the first column describes the statistical behavior of the benchmark S&P 500 as a buy/hold investment, and the second column of data represents a test strategy that is always either long or short the S&P 500, depending on the position of the hedge. Both trials assume the investment at a beta of 1.0 — we are simply examining the statistical characteristics of the model's signals on its own.

Figure 6

Keller Partners Equity Trend Model Dec 31, 1996 — February 29, 2008	SPX Buy / Hold	SPX Long / Short
Standard Deviation of Monthly Returns	4.32%	4.16%
Number of Trades	0	50
Average Month Ret	0.53%	0.81%
/ IRR	5.40%	9.03%
Max Drawdown (1)	-27.9%	-24.8%
No. of Positive Mos	80	79
/ Avg Gain Pos. Mos.	3.33%	3.49%
No. of Negative Mos	54	55
/ Avg Neg Mos	-3.57%	-3.07%
Sharpe Ratio (annl)	**0.23**	**0.47**
Beta	1.00	0.34
Alpha	0.0%	9.8%
R-squared	1.0000	0.0013

(1) Rolling three-month drawdown.

THE HEDGE ON ITS OWN. If we examine the return patterns generated by the hedging model's signals, trading the S&P 500 long and short in line with the model's output, we note that (a) it is inherently noisy, (b) that by itself it does not reduce drawdown or volatility very much, but (c) the IRR and, accordingly, the Sharpe Ratio, is significantly improved over that of the unmanaged S&P 500.

The short hedge so dramatically alters the return pattern of the hedged investment that normal MPT metrics are no longer helpful. Note that the R-squared of the investment in the second column is effectively zero, rendering the measured alpha and beta meaningless.

However, the annualized Sharpe Ratio (a simple, but widely-accepted measure of risk-adjusted returns which is not dependent on benchmark correlation) appears to meet our basic requirement of significantly improving the risk-adjusted return of the underlying investment.

Although most of the metrics in the table did move in the hoped-for direction, nothing is very spectacular except for the change in IRR and the significantly-higher Sharpe Ratio. So, in a conceptual sense, alpha is most likely present, but it is erratic. We conclude that our model captures weak, but real, trend direction information for equities which, over time, can translate into a significant advantage to the final investor.

Figure 8

KPLLC Equity Trend Model		
Dec 31, 1996 — February 29, 2008		
Number of Transactions	56	
Average Time Between Changes	73	days
Average Trades / Year	5.01	
Successful Signals (both directions)	32	
Average Gain	5.19%	
Unsuccessful Signals	23	
Average Loss	-3.86%	

A Dynamically-Hedged ETF Portfolio

Simulated portfolio results never carry the weight of actual performance, but they do allow us to investigate whether our mechanical structures are likely to work as we might expect, whether there are any unintended consequences, and — most important — they generate "what-if" exercises covering market periods that are still very vivid in our collective memories, such as the 2000-2002 tech meltdown.

We begin with a straightforward strategically-allocated ETF portfolio, based on the real-time historical asset class allocations of an investment manager colleague[8]. We took the exercise back to May of 2000 when the first wave of style-box Russell ETFs were introduced by i-Shares (Barclay's). They included ETFs for *Large-Cap Growth*, *Large Cap Value*, and *Small Cap* and, soon thereafter, *EAFE*. Today, of course, there are ETFs for virtually every imaginable asset class.

Beginning the exercise in mid-2000 provided two advantages:

1. It allowed us to use historical ETF market values for the base portfolio, including their real fees, real tracking errors, and so on. We had no choice but to simulate the behavior of the inverse index ETFs from the daily price changes of the S&P 500, less imputed expenses, but we think that this is a defensible assumption, given how well they track today.

[8] These historical semi-annual portfolio allocations were graciously provided to us by Tom Ashton of Vantage Advisors. LLC of Raleigh, North Carolina.

2. Mid-2000 was also the approximate peak of the Nasdaq bubble, providing us with a very relevant test period during which the buy-and-hold performance of even large-cap portfolios that mirrored the S&P 500 were subjected to unacceptable levels of drawdown (-49% for the index) between August 2000 to October 2002.

The performance metrics for the unhedged base portfolio, together with two experimental percentage allocations to the hedge (15% and 20%) are presented in Figure 7.

Figure 7

Portfolio	IRR	MDD	StdDev	Beta	Alpha	R-Sq.
Benchmark (S&P 500 with divs)	3.13%	-28.87%	1.09%	1.00	0%	1.00
Aggressive ETF Portfolio	4.39%	-24.34%	0.90%	0.81	2.59%	0.94
/ with 15% Hedge Allocation	5.68%	-14.98%	0.69%	0.56	3.69%	0.80
/ with 20% Hedge Allocation	5.98%	-14.10%	0.65%	0.53	3.92%	0.78

- *Time period of simulation: May 31, 2000 — December 31, 2007.*
- *MPT Statistics are calculated against S&P 500 with dividends.*
- *All drawdown calculations in this table are based on rolling 90-day lookbacks.*

SIMULATED ETF PORTFOLIO. This exercise illustrates the potential benefit of combining two asset classes with very low correlation, but each with positive, albeit erratic, alphas. The two components tend to mutually reduce volatility, while the value-added contributed by each is allowed (for the most part) to make its way to the bottom line. Here, in the interaction between the two processes, we see significant reductions in drawdown and some improvement in IRR as well.

This exercise utilizes the second implementation scheme discussed above, that is, the 15%-20% fixed income allocation of the portfolio is converted to cash and used to fund the hedge when it is called for. At other times, the returns include the 15%- 20% short-term fixed income component. The graphic in Figure 8 on the next page displays the 15% hedged portfolio, the unhedged buy-hold portfolio, and the benchmark

Several comments on this simulation:

- The middle plot (dark blue line) is the unhedged ETF portfolio[9] using actual ETF market prices. It exhibits a classic "efficient frontier effect" of a small positive alpha against the S&P 500 total return benchmark. The S&P is the lowest, magenta line on the graph.

[9] The ETF portfolio is based on an active allocation model maintained by Vantage Advisors, LLC which rebalances quarterly. The current allocation is: Large Cap Growth (IWF) 19.6%, Large Cap Value (IXX) 23.8%, Small Cap. (IWM) 15.3%, International (EFA) 26.4%, and Fixed Income, 15.0%

- The 15% hedged portfolio simulation (top, red line), on the other hand, proves much more defensive. While it struggles to stay even during the bear market, it does a much more credible job of maintaining portfolio value with a maximum 3-month drawdown of 15%, versus the unmanaged index at 24%.

- Both hedged portfolio variations generated a worthwhile improvement in metrics. With the 20% hedge allocation, volatility is reduced a bit more and the IRR rises a little more versus the portfolio with the 15% hedge allocation.

Figure 8

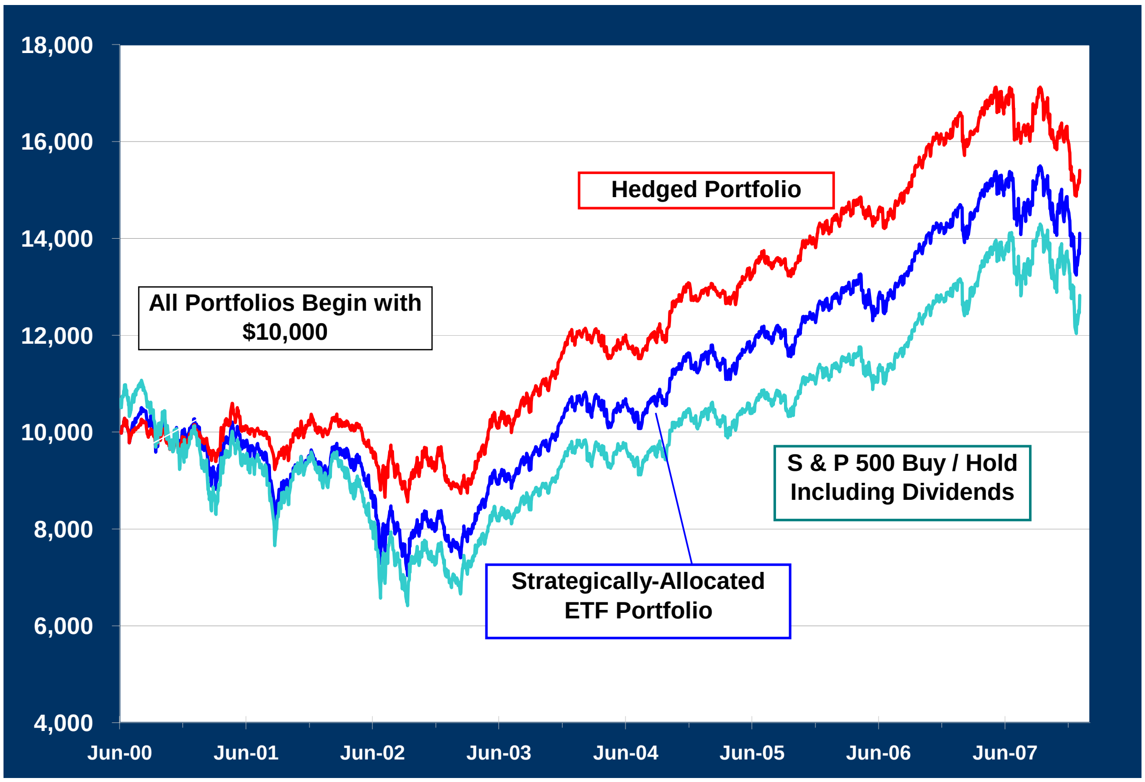

NAVIGATING TROUBLED WATERS. It appears that the asset class-diversified ETF portfolio has some "efficient frontier" characteristics that improve both its risk and return metrics versus the S&P 500. However, diversification is not a strong defense when prices go down, since correlations inevitably rise when markets decline. The lowest line is the performance of the S&P 500, the middle line is the unhedged ETF portfolio, and the top line is the simulated hedging exercise with a dynamic 15% hedge allocation.

SUMMARY: SUPPORTING THE MISSION OF THE PROFESSIONAL MANAGER

In the final analysis, it is the investment manager who is able to deliver a bend in the Capital Market Line that deserves the largest management fee[10]. This bend — or non-

[10] A hedge fund analyst and author, Alexander Ineichen, has recently published a fine book titled *Asymmetric Returns: The future of Active Investment Management* which underlines the importance of non-linear return patterns: equity-like returns on the upside and bond market-like volatility on the downside.

linear return pattern — is illustrated in Figure 9 below. Warren Buffett achieves non-linear returns with a value approach where individual portfolio selections have asymmetrical risk/reward profiles. For the rest of us, we feel that rules-based dynamic hedging has considerable potential for transforming conventional beta-dependent portfolios, with their inherently linear gain/loss potential, into an investment service that the client truly yearns for: equity returns on the upside with controlled drawdown risk — the dashed line in the slide below. The potential benefits to client comfort, as well as to their long-term returns are substantial, especially in markets where beta-based returns can no longer be relied on.

Figure 9

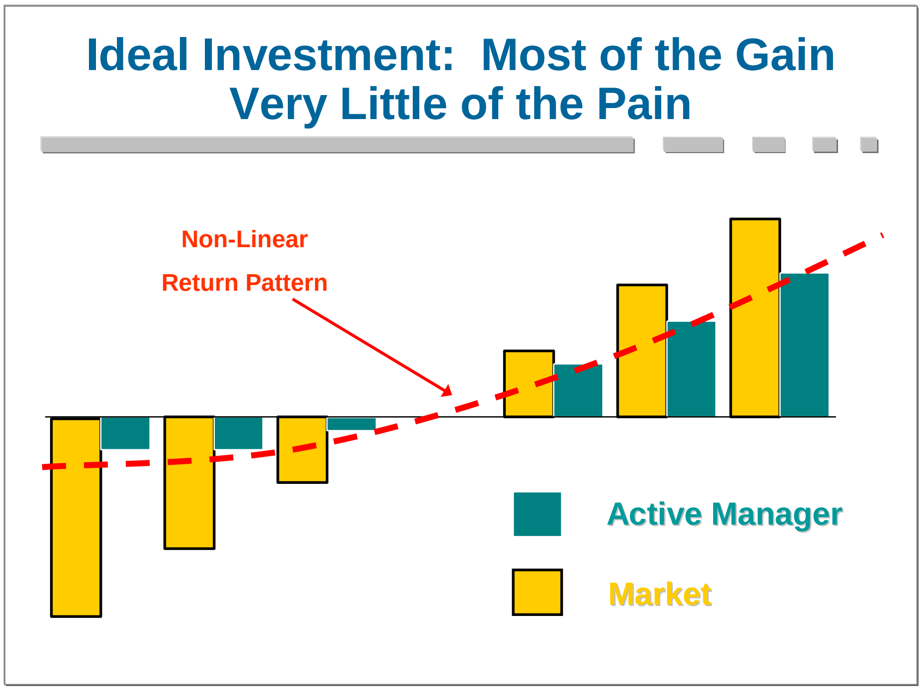

ASYMMETRIC RETURNS. This slide comes from a 1996 client presentation of our predecessor firm. It illustrates our portfolio return tradeoff proposition at that time: two-thirds of the gains, with control of large negative returns. This was not always easy to deliver. Leveraged, inverse index ETFs probably would have made our work much easier!

The recent arrival of liquid, listed, Exchange Traded Funds that deliver leveraged, inverse index returns has presented the financial professional with an exceptionally-convenient low-cost mechanism to implement such hedging strategies.

To manage the hedge decision, we have found it possible to build market models that can identify the underlying trend of a particular financial market (in this case, the US stock market) with enough certainty to permit a statistically meaningful improvement over simply holding the unmanaged index. The value added by any such a model is admittedly irregular and unpredictable in the short-term. But, when integrated into a well-constructed, diversified portfolio, the combination of the two components can be financially quite powerful.

Werner E. Keller, CFA is Principal of Keller Partners LLC, a registered investment adviser focused on the development and implementation of quantitative investment management tools. Previously, he was the founder of Centurion Capital Management, an active investment management firm, which was acquired by General Electric Capital (now Genworth) in 2002.

He can be reached at wkeller@kellerpartnersllc.com or at 888-255-9000.

Figure 10

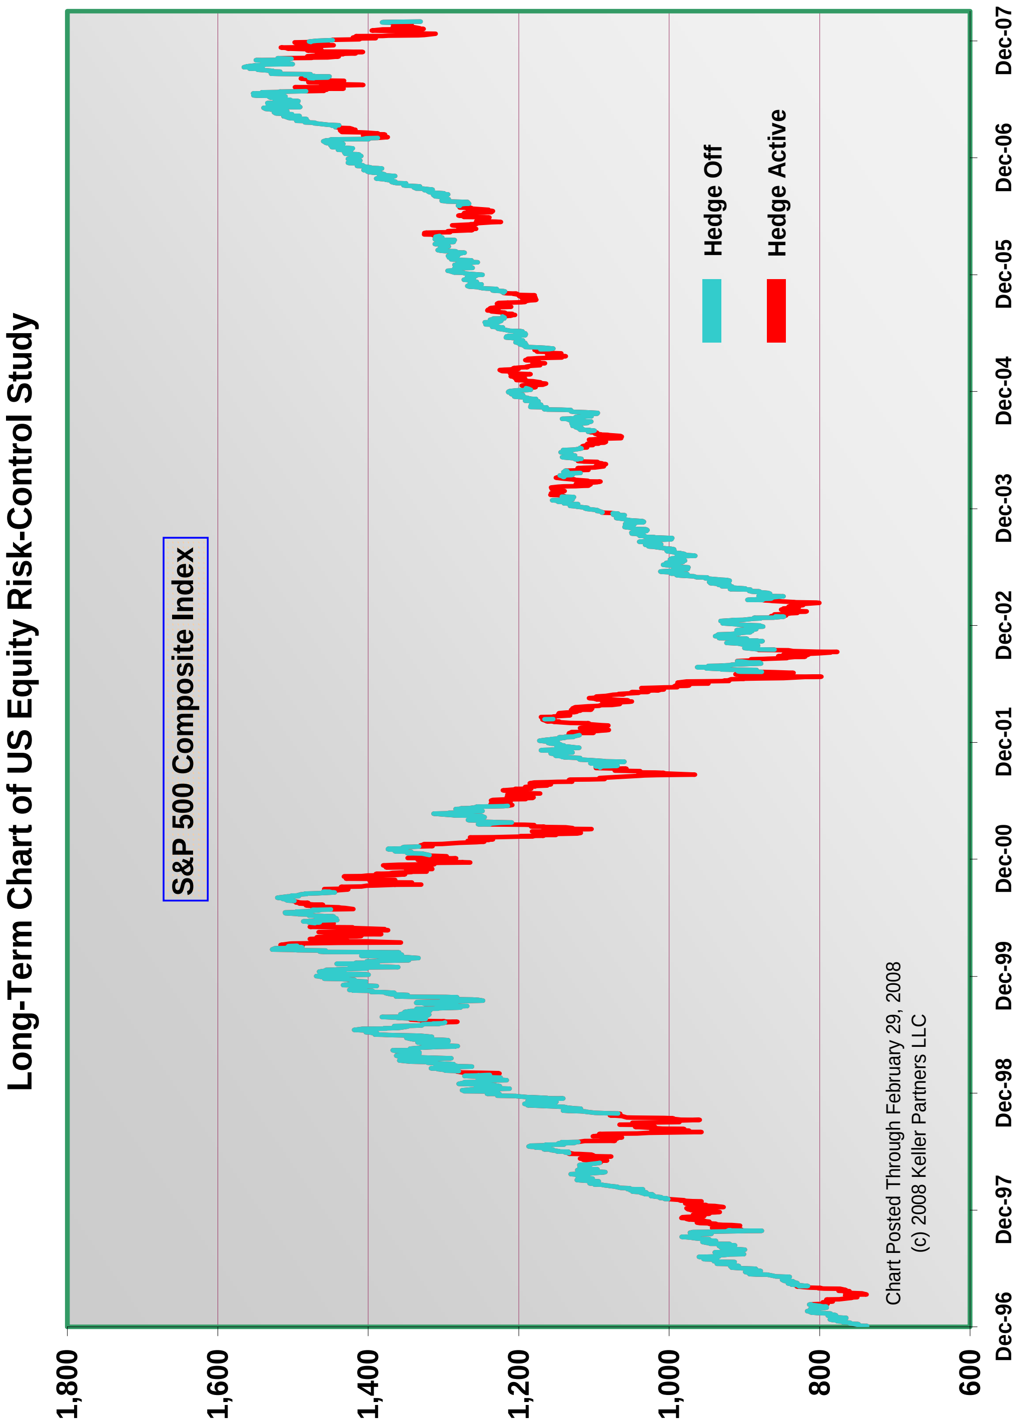

THE RIGHT SIDE OF THE TREND, MOST OF THE TIME. Since the KPLLC trend model includes internal (non-price) data, some of the buy/sell points do not occur at completely logical places — buys can occur when prices are still declining and sells can occur when they are rising. The most important design criterion is that, for the comfort of clients and their advisers, the hedge needs to have been invoked for several weeks prior to any major market disruption — the sharp downward spikes that tend to generate high client anxiety.