Individual CTA Replication: Theory and Application to Manager X

Conquest Capital Group LLC

January 2007

CONQUEST
CAPITAL GROUP

Individual CTA Replication: Theory and Application to Manager X

A Conquest Capital Group LLC Study

In this report, we describe a general method of replicating individual CTA managers based on the Conquest MFS method. We then apply it to replicating the performance a sample manager, heretofore referred to as Manager X, whose returns were provided by a third party.

The MFS method consists of using a system-market portfolio as widely diversified in both markets and time frames. The same approach is applicable to an individual CTA. The allocations to the six sectors (FX, fixed income, etc.) and among time frames have to vary significantly to match the reality of actual varying allocations of a specific money manager. Yet, we still take the approach of minimizing optimization. For example, we try to avoid as much as possible varying individual market weights within each sector. For example, allocating 5 times as much to currencies as to commodities may be a reasonable strategy, which our clone has to match by doing the same. But allocating 5 times as much to the euro as to the yen is not reasonable, usually not done in reality, and we explicitly avoid it.

We start this report with a general overview of the MFS method as applied to the CTA sector, which should be helpful in understanding the logic behind the whole procedure. In Section 2, we discuss other replication methods and their inapplicability to CTA replication. Section 3 deals with the peculiarities of replicating an individual CTA vs. a CTA index. Section 4 discusses two main issues of correlations. First, the issue of the correlation stability in the process of creating a replication portfolio. Second, the level of the correlation between a replication portfolio and the CTA that can be reasonably expected in an individual CTA replication. In Section 5, we discuss the principles of selecting systems and markets. Having developed the guiding principles, in Section 6 we describe the process of building the replication portfolio step by step for Manager X. Finally, in Section 7, we discuss the results. Appendix A discusses the intermediate results obtained during the replication process itself in detail. Replication of Manager Y will be considered in our next report.

1. MFS Method Overview

The MFS approach to replicating CTA strategies is based on the simple well-known fact that profitable opportunities in financial markets are due to the existence of trends, which permeate all markets in all time frames. MFS uses simple trend-following systems that cover the whole spectrum of CTA time frames with equally high density on all scales.

The approach works for the following reason. Any CTA strategy is a combination of entry signals, filters, and exits. Each of these components is based on capturing trends in its specific time frame, and thus the strategy as a whole can be presented as a portfolio of trend-following strategies. As a simple example, consider a strategy that takes a long position based on the price crossing above a 20-day moving average, but only if the 100-day MA is rising, and exits on a stop set at the one-week low. Obviously, it is a combination of exposures to 20-day (signal), 100-day (filter) and 5-day (exit) trends. When CTA strategies are combined into a portfolio, the trend exposures of each strategy overlap and combine into a coherent overall time frame exposure. Such exposure can, for example, be discovered by calculating the CTA correlations to the MFS breakout strategies.

This method is not confined only to trend-following CTAs. As is well-known, a plain-vanilla counter-trend approach, i.e. just taking a position against a trend, doesn't work. A proof of it is the fact that, ignoring slippage expenses, trend-following strategies are profitable in all time frames. A CTA can have a strategy that takes trades against trends in a specific time frame. Yet, to make it profitable, it will have to employ multiple filters, which will be trend-following in some other time frames. The correlation chart to MFS systems in such a case will show the CTA having long exposures in all time frames except for the one in which it is counter-trend. Interestingly, in studying real CTAs it is very difficult to find negative correlation in even a single time frame, except for a couple of the shortest-term strategies. Since CTAs have many strategies in their portfolios, the explicit single counter-trend exposure in one strategy is easily offset with multiple residual trend-following exposures from the other ones.

Therefore, while a CTA views its portfolio as a combination of strategies, the MFS approach views the same portfolio as a combination of time frame exposures.

2. Applicability of Other Replication Methods

While several other replication methods have been discussed in research papers on the subject and applied to the overall hedge fund space or to strategies like long/short and merger arbitrage with some level of success, we believe that they are inadequate for CTA replication.

Several replication methods are based on using exotic options, such as lookback straddles. While very interesting in theory, they are impractical in reality due to prohibitive trading expenses. Exotic option positions suggested by such models either have to be created synthetically, leading to large trading turnover in rather illiquid markets and thus to excessive slippage, or would have to be negotiated with a bank as a counterparty, which would result in excessive fees.

On the other extreme of the complexity spectrum lie simple factor models. They replicate strategies by taking positions in several (usually less than 10) broad market indices, replicating exposure to equities, bonds, US dollar, emerging markets and commodities. The exposures to such indices are either fixed or vary monthly based on monthly data over several previous years. This method gives good results for strategies that do have stable exposure to such markets. For example, a long/short fund can have a significant long exposure to the stock market. Or a fixed income arbitrage fund can be replicated by buying high yield or convertible bonds. Unfortunately, simple factor models don't work for CTA strategies due to the nature of such strategies. Significant portion of CTA trading is done in relatively short time frames based on data spanning several days to several weeks. Even monthly rebalancing based on several years of price data is far too slow to capture such behavior, not to mention market weights that remain fixed for decades. Actually, using years of data in changing factor weights or even keeping them the same means that the shortest trends that can be captured with such approach must last for years. This is way longer than the longest trends which concern CTAs. The problem has been recognized by authors of factor models, who usually point out very low correlations between the replicated portfolios and CTAs.

The only viable approaches to CTA replication that are discussed in literature consist of ignoring generic asset classes and using portfolios of basic trend-following strategies. Yet, for some reason, each of them is limited in some respect. Some use very few time frames (e.g., the 12-month moving average approach of MLM), some use long-only exposures (e.g. CRB or GSCI commodity indices), while practically all are limited in their selection of markets.

MFS has taken this idea one step further. Its main purpose is maximum diversification both in terms of markets and systems. As in any industry, an increasing level of competition causes real CTAs to move away from the beaten path and look for new opportunities both in terms of new markets and new time frames. Since the number of liquid markets and time frames is very limited, we can thus be sure that the CTA industry as a whole has a more or less balanced exposure to all of them. And having an exposure to most markets and all time frames is exactly the goal and the advantage of MFS.

While this market and time frame diversification is obviously relevant for CTA indices, it is true for individual CTAs as well. Diversification is a key to improved risk-adjusted performance. For example, a short-term CTA can focus on the short-term time frame as a matter of its trading style, but may notice that adding long-term filters improves its risk-adjusted returns. Such filters, as we have seen above, can be viewed as time-frame diversification. Therefore, even an individual CTA has to be approached in the diversified style of the MFS method.

3. Individual CTA vs. Index: Replication Differences

While the currently traded MFS program is meant to replicate CTA indices, its application to individual CTAs is quite similar. Each CTA in a sense is an "index" consisting of a portfolio of its strategies. On the other hand, each real index is in a sense a "fund", whose "strategy" is the selection approach of its creators.

There are actually good reasons for CTAs becoming as diversified as indices and for indices becoming as specific as CTAs. As most investors look at risk-adjusted returns, CTAs respond with building more diversified portfolios. As more indices appear in the market, index creators try to differentiate themselves from the rest by designing their own creative approach to individual manager selection. More diversification of CTAs and more creativity of indices will ultimately drive them closer together.

The success of CTA replication depends on two factors. First, on the stability of the CTA trading strategy. Significant abrupt shifts in market (sector) allocations or time frame exposures cannot be captured immediately by anyone, unless they are explicitly announced by the manager. Return correlations will show the effect of such changes, but only after many months of new data are collected (please see the next Section for details). Fortunately, large successful managers usually don't overhaul their strategies overnight or at least announce such major overhauls to investors.

The second success factor is the degree of CTA diversification. As an extreme example, a small CTA running a single highly unconventional (quirky) strategy will be very difficult to replicate. Again, fortunately, large successful CTAs diversify their strategy portfolios, smoothing away all individual strategy irregularities. Also, for any CTA the degree of such irregularities can be easily seen in the first stages of a replication procedure, since any such noise drives down correlations with CTA indices and MFS.

To summarize this Section, we believe that for large successful CTAs, i.e. those CTAs that are really worth replicating, the difficulty of replication should not be much higher than that for a CTA index.

4. Individual CTA Replication – Expected Correlation

a) Correlation Issues in Replication Portfolio Construction

Here, let us summarize what correlations we will be looking at in calculating replication portfolio factors. The most obvious parameter is the overall correlation between the CTA and the portfolio over the full 15-year period. But it does not have to be the only one. There are two problems with the 15-year correlation.

First, a lot could have changed in the fund's system-market portfolio over that period, lowering the correlation and making portfolio weight adjustments to maximize it less meaningful. Second, any correlation can be partially spurious – we need some way to check its stability.

We therefore split the total 15-year time period into 3 equal 5-year periods. We can now also look at 3 independent correlations over the three non-overlapping time periods. The most important of these is the correlation over the most recent 5 years. Things might have changed in the past 15 years, and the most recent time period definitely has the highest significance.

Speaking about recent performance, at first glance, it seems that we must consider much shorter time frames to infer the current CTA structure with total accuracy. Fortunately, this is not absolutely necessary. As was mentioned before, large CTAs don't overhaul their portfolios dramatically on a whim. Also, desire for diversification and avoidance of optimization usually causes their market weights to be relatively similar within the same sector. Also, a highly important recent development, such as an addition of a very different new market one year ago, will most likely still show in the last 5 years correlation. The effect will be weaker if it comes from only part of the data stream, but if it is strong in the last year, it should be at least discernible in the last five.

Going to very short time frames is mostly not advisable due to issues of accuracy. The simple rule to estimate the significance of the hypothesis that an observed correlation is at least different from zero consists in calculating its Student's t-ratio by the formula $t_{N-2} = r \sqrt{((N-2)/(1-r^2))}$, where N is the number of observations and r is the observed correlation. Therefore, even for a 5-year period (N=60 months), an observed correlation of within ±25% is still consistent with the real correlation being equal to zero (at the 5% level of significance). For 15 years of data, the significantly non-zero level of correlation drops to ±15%. For three years, it rises to ±33%, and for one year to ±58%. Thus, there is very little that we can reliably infer even from very high absolute levels of one- or even three-year correlations – in most cases, we cannot be even sure of their difference from zero.

b) Expected Correlation between a CTA and its Replication Portfolio

As described in the following sections in more detail, we calculate the replication portfolio weights based on maximizing correlations. While, in principle, any target values of correlations (or any other parameters) are achievable if a large enough number of parameters is allowed to vary in an optimization procedure, we must take special precautions to avoid over-optimization. One such precaution consists in estimating the reasonable levels of parameters that we want to achieve. For returns and volatilities these are obvious – they must be similar to those of the replicated manager. Here we consider correlations.

As is known from the multiple regression theory, the natural multi-dimensional extension of the notion of correlation is the coefficient of determination. In one dimension, it is simply the square of the correlation coefficient. On the other hand, the coefficient of determination has a clear meaning – its difference from 1 is the relative level of noise (or the sum of squares of errors) in the signal vs. the benchmark. This provides a clear intuitive notion that adding noise to a system is bound to decrease the correlation between its returns and those of any model. As CTA trading decisions can be very creative, their level of "noise" to benchmarks is quite high, especially in individual manager returns.

Since individual CTA returns involve more non-diversified high-impact decisions by their managers than average CTA industry indices, they are noisier relative to any benchmarks and their correlations must be lower than those of indices. It thus looks reasonable to use CTA index correlations as the upper target for correlations in individual CTA replication. MFS correlations to both Calyon Barclay and S&P Managed Futures CTA indices are shown below:

Correlations of CTA Indices			
	MFS	S&P Managed Futures	Calyon Barclay
MFS		86.9%	86.8%
S&P MF	86.9%		91.8%
Calyon	86.8%	91.8%	

Since each index is in a sense the representative of the CTA industry, correlations of indices to each other can be viewed as an ideal limit of the CTA industry replication, in which two CTA industry representations "replicate" each other. Even in this ideal case, the correlation is only 92%. For a real replication, e.g. using MFS, the correlation drops to around 87%.

As mentioned elsewhere, the replication correlations calculated in research papers fall significantly below this level. For example, L. Jaeger and C. Wagner in *Factor Modeling and Benchmarking of Hedge Funds* obtain the coefficient of determination for manage futures investable index replication, which corresponds to a 58% correlation. A. Lo & J. Hasanhodzic in *Can Hedge-Fund Returns be Replicated* study individual fund replication. Again, converting their R2 to correlation coefficients, for managed futures they on average get a 39.1% correlation. We see that MFS correlation being so close to raw index-to-index results is quite unusual.

Therefore, in our case of a one-manager replication, we can reasonably target the correlation to be somewhere in the range of 80-85%.

Also, as mentioned above, the correlations in the most recent years must be higher. We apply the same system-market portfolio weights to the whole time period, while in reality they have changed significantly. Since our goal is to match the current CTA portfolio composition, we must see the effect of a high recent correlation trailing off going back in time.

As will be shown below, for Manager X we can obtain a correlation of over 85% for last 5 years and over 80% over the whole 15-year period without any optimization of individual market weights within sectors.

5. Principles of Replication Procedures

a) System Selection

Now we are ready to discuss the particular principles of the replication process for an individual CTA, starting with principles of selecting systems for the replicating portfolio.

MFS systems, which we use in this replication, are simple trend-following strategies based on N-day price breakouts, with N ranging from 5 to 200 days. There are 20 systems in the MFS portfolio. The range of N's is wide enough to cover the whole CTA time frame spectrum. For N < 5 slippage becomes excessive, for N > 200 the average trade length exceeds 1.5 years and becomes too long for practically all CTAs.

The specific values of the time periods N for the systems were chosen to make the ratios of N's of each two consecutive systems the same, thus evenly spacing them on a logarithmic scale[1]. As the table below

[1] A detailed report is available on the selection of the MFS time frames N. Here we summarize its conclusions.

shows, this approach results in very similar correlations between each two consecutive systems (and between any two of them separated by the same number of systems):

Correlations of Returns of 20 MFS Systems																				
	B005	**B006**	**B007**	**B009**	**B011**	**B013**	**B016**	**B019**	**B024**	**B029**	**B035**	**B042**	**B051**	**B062**	**B076**	**B092**	**B112**	**B136**	**B165**	**B200**
B005		94%	87%	77%	68%	61%	54%	49%	44%	40%	33%	28%	24%	22%	21%	16%	14%	12%	10%	7%
B006	94%		95%	86%	77%	70%	62%	57%	51%	46%	38%	33%	29%	27%	25%	19%	17%	14%	13%	9%
B007	87%	95%		92%	84%	77%	69%	64%	57%	52%	44%	39%	34%	31%	29%	23%	20%	17%	15%	11%
B009	77%	86%	92%		94%	87%	79%	73%	66%	61%	52%	47%	42%	39%	35%	29%	25%	22%	19%	16%
B011	68%	77%	84%	94%		95%	88%	82%	74%	69%	59%	54%	48%	45%	40%	34%	31%	27%	24%	19%
B013	61%	70%	77%	87%	95%		94%	89%	81%	76%	65%	60%	53%	48%	44%	38%	34%	30%	27%	23%
B016	54%	62%	69%	79%	88%	94%		96%	88%	82%	71%	65%	58%	53%	48%	41%	38%	33%	30%	26%
B019	49%	57%	64%	73%	82%	89%	96%		94%	88%	76%	70%	62%	56%	51%	45%	41%	36%	32%	28%
B024	44%	51%	57%	66%	74%	81%	88%	94%		95%	84%	77%	69%	62%	56%	50%	45%	39%	36%	31%
B029	40%	46%	52%	61%	69%	76%	82%	88%	95%		89%	83%	75%	68%	61%	55%	50%	43%	40%	34%
B035	33%	38%	44%	52%	59%	65%	71%	76%	84%	89%		95%	88%	80%	72%	67%	61%	54%	49%	43%
B042	28%	33%	39%	47%	54%	60%	65%	70%	77%	83%	95%		94%	87%	78%	73%	67%	60%	55%	48%
B051	24%	29%	34%	42%	48%	53%	58%	62%	69%	75%	88%	94%		94%	86%	80%	73%	65%	60%	53%
B062	22%	27%	31%	39%	45%	48%	53%	56%	62%	68%	80%	87%	94%		94%	88%	81%	74%	68%	60%
B076	21%	25%	29%	35%	40%	44%	48%	51%	56%	61%	72%	78%	86%	94%		92%	86%	80%	73%	63%
B092	16%	19%	23%	29%	34%	38%	41%	45%	50%	55%	67%	73%	80%	88%	92%		95%	87%	80%	72%
B112	14%	17%	20%	25%	31%	34%	38%	41%	45%	50%	61%	67%	73%	81%	86%	95%		92%	86%	77%
B136	12%	14%	17%	22%	27%	30%	33%	36%	39%	43%	54%	60%	65%	74%	80%	87%	92%		94%	86%
B165	10%	13%	15%	19%	24%	27%	30%	32%	36%	40%	49%	55%	60%	68%	73%	80%	86%	94%		93%
B200	7%	9%	11%	16%	19%	23%	26%	28%	31%	34%	43%	48%	53%	60%	63%	72%	77%	86%	93%	

We see that the 20 systems cover the whole time frame spectrum quite "densely". Even if we left out every second system, each two new consecutive systems would still be 83-89% correlated. This means that there is nothing unique about the choice of the specific time frames N. There are enough systems to cover all time frames with a significant overlap.

The number of the systems does not matter either, as long as they are close to each other enough to cover the whole spectrum without major gaps. The overall number of systems is based on a balance between two opposite considerations. On one hand, it must be as high as possible to cover all time frames. On the other hand, it must be small enough due to the finite size of market contracts. The required number of contracts in each position is calculated as a real number, which is then rounded to an integer. A lower number of systems results in higher individual system weights and thus higher positions. The relative rounding error is minimal when the position size is as high as possible. In the worst case, when the real-number position size falls below 0.5 lots, the position is missed completely. For our individual CTA replication, for which market allocations must be more specific and thus vary more among systems and markets, this issue becomes more significant than for MFS. Also, for decreasing the number of systems naturally decreases the amount of optimization in allocating system weights.

Therefore, for the individual CTA replication, we have decreased the number of systems from twenty to ten. Their highest and lowest N's are still 5 and 200. All other N's are again chosen by the same method as that of MFS: so that the ratio of N's remains constant for each two consecutive systems. A straightforward calculation based on these rules generates the following new system lengths, shown here together with the old 20 MFS lengths:

20 Systems (MFS)	5	6	7	9	11	13	16	19	24	29	35	42	51	62	76	92	112	136	165	200
10 Systems (Single-CTA)	**5**		**8**		**11**		**17**		**26**			**39**		**58**		**88**		**133**		**200**

Since all but three of the new systems have N's that are different from those of MFS, this will have an added advantage of diversifying trade entries between this replication program and the old MFS portfolio, which will lower their trade entry interference with each other and thus slippage expenses.

The numbers in the system names after the letter "B" (for "breakout") indicate their time frames in days. E.g., B200 is a 200-day and B007 is a 7-day breakout systems.

As the new cross-correlation table below shows, the systems are still close enough to each other, with adjacent systems having 80-89% correlations to each other:

Correlations of Returns of 10 Individual CTA Replication Systems										
	B005	**B008**	**B011**	**B017**	**B026**	**B039**	**B058**	**B088**	**B133**	**B200**
B005		83.6%	62.6%	37.2%	29.2%	15.3%	15.6%	20.8%	21.7%	22.0%
B008	83.6%		86.7%	60.7%	46.6%	25.1%	24.1%	20.1%	21.2%	28.2%
B011	62.6%	86.7%		79.5%	65.0%	38.9%	34.2%	25.0%	21.9%	22.9%
B017	37.2%	60.7%	79.5%		80.6%	55.0%	47.1%	31.5%	23.4%	22.0%
B026	29.2%	46.6%	65.0%	80.6%		80.9%	64.4%	45.9%	36.2%	34.9%
B039	15.3%	25.1%	38.9%	55.0%	80.9%		85.9%	65.4%	56.6%	51.6%
B058	15.6%	24.1%	34.2%	47.1%	64.4%	85.9%		87.9%	77.0%	67.1%
B088	20.8%	20.1%	25.0%	31.5%	45.9%	65.4%	87.9%		89.3%	71.8%
B133	21.7%	21.2%	21.9%	23.4%	36.2%	56.6%	77.0%	89.3%		84.6%
B200	22.0%	28.2%	22.9%	22.0%	34.9%	51.6%	67.1%	71.8%	84.6%	

Exposures of a CTA to different time frames should be replicated by varying system weights, not by discarding individual systems or changing their time frames N. The systems themselves, if we may use an analogy, present an even fabric, on which we can "paint" our time frame exposure. The fabric should be as smooth as possible, we don't want to tear it (delete systems) or to move its individual fibers.

While the evenness of the time frame spacing is no longer used to ensure the even exposure to all time frames, as the weights are now different, it is still useful in making all external weights transparent. For example, consider the same 10 systems, only with B026 replaced with B018. As it is now very close to B017, equal-system allocations now would result in an almost double exposure to 17-18 day trends and much lower exposure to, say 25-day trends. Equal weights no longer mean equal time frame exposures.

Each complex CTA strategy has exposures to many time frames simultaneously. Charts of correlations of individual CTAs to MFS strategies show that such exposures vary gradually with N. The reason for this, of course, is that the complexity of a CTA strategy can be represented as adding exposures to trends in a multitude of other time frames. Each time frame, in turn, is correlated to adjacent frames, which further "smears" the exposure among the time frames.

Unfortunately, any optimization procedures used in selection of system weights (regression, maximization of correlations, etc.) may lead to abrupt discontinuous results. A slight advantage of one of the two highly correlated systems over the other may result in the former being chosen and the latter discarded. This problem is due to the high cross-correlation between systems and has nothing to do with real exposures. To avoid erroneous over-optimized results, we need to correct for it "by hand" by explicitly watching for the gradual nature of system weight variation.

To summarize, we use the following principles of system selection:

1) We use all of the 10 MFS-type breakout systems with time frames evenly distributed between 5 and 200 days, no system is discarded or added;

2) System weights are adjusted to match the CTA time frame exposure so as to achieve the highest correlation over the last 5 years and over the total 15-year period;

3) System weights must be gradual functions of their time frame N. To ensure this, we require that the highest and the lowest weights have the ratio of no more than 10, and check the results to monitor their adherence to this principle.

b) Market Selection

The issues in determining the best market (sector) exposure are different. As discussed above, adjacent system time frames are highly correlated, causing the time frame exposure to smoothly flow from one time frame to the next. Markets, on the contrary, are separate discretionary entities.

Technically, the perfect goal of a replicating portfolio is to figure out and include those and only those markets that are used by the replicated CTA. Yet, the exact completion of this task is, on one hand, impossible, but on the other hand, unnecessary. The reason for both is simple: although they are different from each other, markets can be split into highly correlated groups. The markets belonging to each group are easily interchangeable. For example, a CTA may trade US 10-year bonds, while the replication portfolio may use US 5-year and US 30-year bond combination with very similar results.

The largest groups of correlated markets are, naturally, market sectors. We consider six sectors: FX, equities, bonds, energy, metals and commodities. Sector weights are chosen by each investment manager in a discretionary way. Since there are only 6 sectors with relatively low cross-correlations, over-optimization becomes much less of an issue. We can calculate sector weights using simple optimization procedures based on the correlation of the replication portfolio to the replicated CTA. As described below, to guard against any remaining over-optimization threats, we base such optimization on several parameters. This part of market weight selection is the "top-down" part of the approach.

The other part can be termed "bottom up". The degree of mutual similarity due to which markets belong to groups vary greatly, and some markets definitely can and have to be excluded if they obviously don't belong to the replicated CTA's portfolio. The fact that they don't belong to the portfolio can be seen from their negatively influencing various correlation parameters. Practically speaking, by far in most cases a market that significantly decreases the overall correlation and hurts recent performance, will also decrease the correlation over the last 5 years and at least one, or often both, of the other 5-year correlations. Such obvious "misfits" should be simply discarded. While it may look like an optimization procedure, it is rooted in reality, as, obviously, no CTA trades all markets, and the markets that are irrelevant to the replicated CTA shouldn't be in the replication portfolio.

The "middle scale" weights, i.e. allocations to the remaining specific markets in each sector, should be made as even as possible. The reason for it is avoiding portfolio optimization and the risk of low future performance. This is especially true for very similar markets. In fact, in our procedure we choose all market weights in the same sector to be the same with no more than a couple exceptions, driven by fundamental reasons, not by optimization results.

We thus can summarize the following principles of market selection:

1) Top-down: sector weights are optimized to match the replicated CTA returns;
2) Bottom-up: the individual markets that hurt correlations between the CTA and the replicated portfolio, should be discarded;
3) Middle scale: market weights within a sector should are chosen to be the same.

6. Replication Procedure for Manager X

Having discussed all the necessary theory, we are now ready to start the practical replication procedure for Manager X.

a) Gross Returns

For the purposes of this study, we were told to assume that Manager X uses a 2/25 fee (2% management, 25% incentive) structure. Using these inputs, we recalculate its returns back to gross values. This allows us to compare and match gross returns of Manager X and the replication portfolio (the "Portfolio") directly, without the influence of smoothing and fee giveback provided by incentive fees. As management and incentive fees together influence statistical characteristics of performance, the replication of gross returns should yield more accurate results.[2]

Regarding the fees, as the volatility of Manager X gross and net returns (29.5% and 25.1%) is, respectively, 2.2 and 1.8 times higher than that of MFS (13.7%) over the past 15 years, we believe that its management fee should be set at 2%, to be consistent with our MFS fee policies. Similarly to MFS, there will be no incentive fee.

Once we match the gross returns of Manager X to those of the Portfolio, we believe that our lower fees will allow the Portfolio to outperform Manager X.

b) Manager X vs. Actual MFS

As we have discussed in Section 3, a CTA is much easier to replicate if its strategy is diversified, "index-like", so to speak. The most straightforward way to check whether a CTA is easy or not to replicate by our method is to compare it to the actual MFS portfolio, without any changes made to the latter.

The table below shows the results. All returns are gross. MFS returns are actual from June 2004 and pro forma prior to that. *Relevered MFS* are the MFS returns with the amount of leverage such as to match the 29.5% gross annual volatility of Manager X over the last 15-year period, assuming full investment into MFS.

MFS versus Manager X (1992-2006)

	Gross Manager X		Relevered MFS	Gross MFS
Annual Return	**29.3%**		**37.3%**	19.3%
Annual Volatility	**29.5%**		**29.5%**	13.7%
Sharpe Ratio (RFR = 0)	**0.99**		**1.26**	1.41
Correlation - Full 15 Years		**70.4%**		
Correlation - Last 5 Years		79.5%		
Correlation - Middle 5 Years		69.8%		
Correlation - First 5 Years		69.0%		

As one can see from the table, even without any adjustments to the MFS system-market portfolio, the returns are good and the correlation is quite high for a single manager (70.4%). It thus should be easy to replicate. In a sense, it is already replicated, all we need to do is improve on it. As we've mentioned in Section 4b, best results from published research articles on replicating managed futures show correlations of 40-60% – and those are based on direct optimization of the replication portfolio weights (e.g. by

[2] Although, in general, having information about the fees and, therefore, gross returns of a CTA are by no means necessary for a successful replication procedure. For example, we will use only the provided net returns in replicating Manager Y in our next report.

multiple regression). Here we have a 70% correlation with absolutely no optimization or even an attempt at replication.

As we've discussed above, one way to check that the correlation is stable and consistent is to split the full time period into three non-overlapping 5-year segments. The lowest correlation (69.0%) does not fall much below the average, which is very good. The last-5-year correlation, 79.5%, is very high, exceeding the average by a wide margin. Thus, all aspects of the stability study show good results.

As the CTA space becomes more competitive and CTA assets under management grow, funds become more diversified and thus more index-like in their structure. This, naturally, should lead to higher correlations to diversified programs like MFS.

From this comparison we can conclude that the most reasonable approach is to start with the MFS portfolio and gradually change its composition to make it closer to that of Manager X.

c) *Actual MFS Test*

Since we already have a back-test of the actual MFS portfolio, which we have studied above, let us start with adjusting its system and sector weights.

The following chart shows correlations of the MFS and Manager X performance to the individual MFS markets:

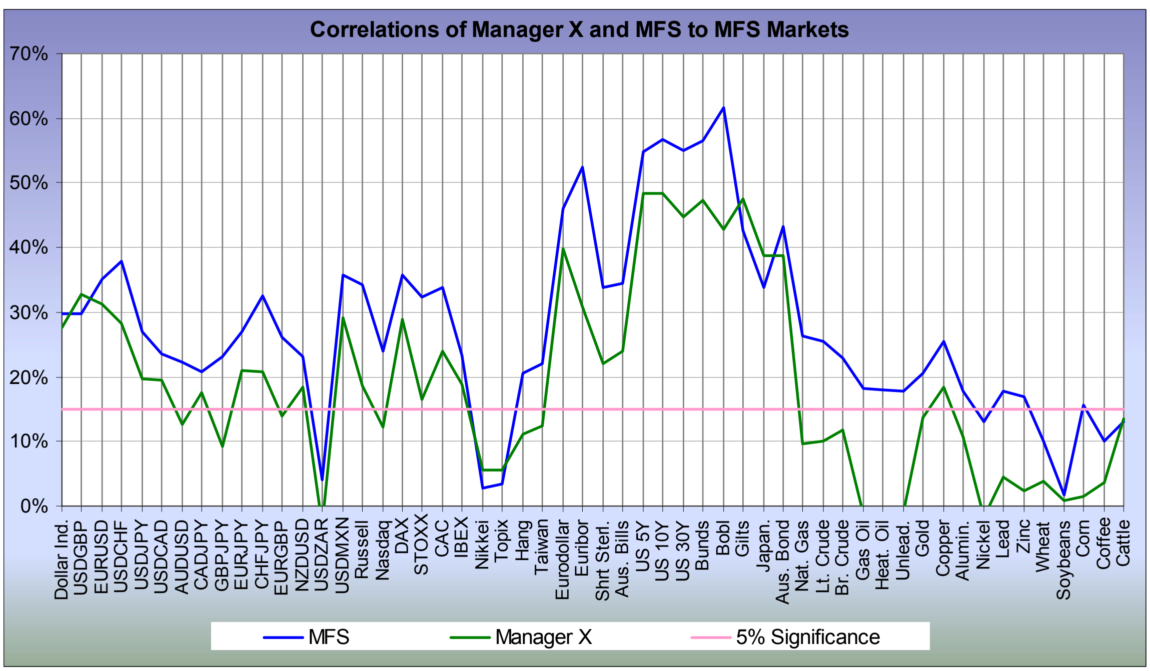

The red line at 15% shows the level of correlations at which they are large enough to be significantly different from zero.

The reason we look not only at Manager X correlations to MFS markets, but also at those of MFS itself, is that different markets have different "natural" correlations to a balanced portfolio. For example, US/European fixed income markets are highly sensitive (and thus highly correlated) to general financial trends. These markets are often correlated even to portfolios that don't trade them. On the other hand, a market like the rand is unusual enough to be non-correlated even to a portfolio that has it. To see whether a market belongs in the Portfolio, we need to view its correlation to the replicated CTA on its own scale, irrelevant to other markets. The measuring benchmark of such a scale is its correlation to MFS. This also explains why, in building a replicating portfolio, we decide whether to exclude a specific market based not on the market's correlation to the CTA, but on how the inclusion of the market changes the total portfolio correlation to its replicated CTA. However strange it sounds, the value of the market's correlation may be irrelevant to the fact of that market's presence in the CTA portfolio.

The markets with the highest difference between the correlations on the chart above are, naturally, the first candidates for being excluded from the portfolio. But we won't do it this way – we'll use an automated procedure, in order to minimize optimization. At this stage, although, we won't drop any markets, only adjust sector weights to improve the total correlations.

The next chart shows correlations of the MFS and the CTA performance to the individual MFS systems:

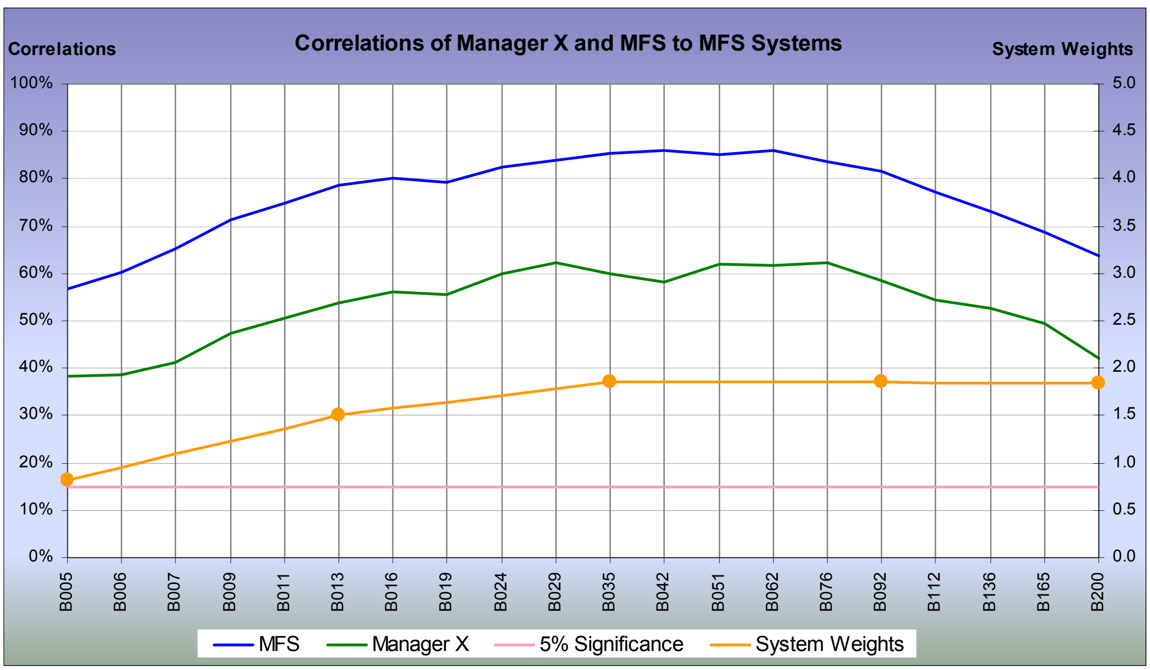

We can see that, as discussed in Section 5a, correlations to systems are distributed very smoothly. If we run an optimization procedure (to figure out which systems to use) based on the 20 system weights directly, we will get meaningless results. Namely, almost half of the system weights will turn out to be zero. The problem with optimization is that it works best only if the input data streams are not correlated to each other. If this is not true, it will give an unfair advantage to every data stream that is just a little better than the other ones, and give those other ones zero weights.

Such problems are dealt with later by decreasing the number of systems to 10 and by limiting the ratio of the highest to the lowest weight. At this stage, we get a smooth system weight curve (shown in orange in the chart) by simply limiting the number of optimized points. We split all systems into 4 segments separated by the 5-, 13-, 35-, 92- and 200-day systems and only vary the weights of those five systems. The weights of other systems are set using linear interpolations. Overall, the curve is rising, implying lower allocations to the short-term systems.

d) *Replication Procedure: The Main Iterations*

Finally, as discussed above, we take the 10 systems, the 55 MFS markets, add 6 more markets that we consider interesting, and start with the system/sector weights that we have obtained from the procedure discussed in the previous subsection.

To calculate replication weights, we need to run several iterations of back-testing. Sector and system weights are obviously interrelated in a sense that the correlation of the whole portfolio to any sector depends on system weights, and vice versa. We thus try to change either only system weights or only sector weights at each stage of the iteration process.

The optimization approach in selecting the weights remains the same. We target the overall 15-year correlation while watching the last 5-year correlation. Fortunately, both other 5-year correlations improve with this procedure significantly. Although we do not specifically target returns, we do make sure that they don't fall below preset levels, especially when discarding markets. Market discarding procedure is run automatically based on maximizing correlations (checking whether market removal would improve the overall correlation), while keeping returns above specific levels. After the procedure is run, each market is checked by hand. As each discarded market is taken out of the future procedures completely and decreases the portfolio diversification, cases of doubt are usually resolved in favor of keeping the market. Also, only markets picked by the automatic procedure are viewed as candidates for being discarded. Again, this is meant to decrease the chance of getting rid of a "good" market.

We thus construct the replication portfolio in three simple stages:

1) Fix system weights. Optimize sector weights. Run the automatic procedure to discard markets. (At this stage, 14 markets were discarded.) Optimize sector weights again with the remaining markets.

2) Fix sector weights, optimize system weights.

3) Same as step one. This time, 7 markets were discarded. This was the final market removal; the remaining 40 markets were kept. The procedure required removal of one or two of three similar metal markets – nickel, zinc or lead – but without a clear preference for any one of them. Thus, in the only individual market weight adjustment case, all three were kept, but their weights were reduced by half.

In total, including the construction of the original portfolio before running the first stage above, we ran three sector weight optimizations, two system weight optimization and two market elimination procedures. Two of the three sector weight optimization were complex two-part ones, before and after market elimination. We believe this is a very modest amount of data manipulation.

To see that the results are not overly optimized, we consider, first, the new sector weights:[3]

[3] Sector, market and system weights are used in determining every trade/position size. When a system X takes a position in a market Y, the position is determined by the product of the respective system and market weights. Market weights usually equal sector weights (in our case, for all markets except for the 3 metals).

	FX	Stocks	Bonds	Energy	Metals*	Commodities
Sector Weights:						
New CTA Replication	3.90	3.97	6.35	1.55	5.75	3.23
Actual MFS	3.75	3.00	5.00	4.75	3.25	5.00
Number of Markets:						
New CTA Replication	12	6	9	4	3.5	4
Actual MFS	15	10	12	6	6	6
Share of the Total Sector Weights:						
New CTA Replication	28%	15%	37%	4%	11%	5%
Actual MFS	27%	15%	31%	12%	7%	8%

*Lead, Nickel, and Zinc are each half-weighted in the final construction

With the exception of the drastic decrease in energy weight, we can see that sector weights are quite similar. Not only are the weights that are used to determine position sizes for each market similar, but the total weights of each sector are also reasonably close.

Second, consider the new system weights (the numbers after the letter "B" in the system titles equal the time frame length in days, e.g. B005 = 5-day system):

	B005	B008	B011	B017	B026	B039	B058	B088	B133	B200
New CTA Replication	0.22	0.34	1.01	0.94	1.62	0.22	2.21	2.10	0.61	0.73
MFS-Like Strategy*	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00

**As discussed above, MFS itself has 20 systems, also evenly spread on a logarithmic scale.*

The distribution is in line with what we've seen as a first approximation. Allocations to long-term systems are higher than to short-term ones. For very long-term systems, they become lower again. All this looks natural for a large trend-following CTA. Again, as explained in Section 1, the inclusion of extreme time frames, like 5- or 200-day ones, does not mean that the CTA actually has to deliberately trade 5- or 200-day systems. It may mean that the systems that it trades include a high enough number of secondary components, like exits or filters, in the respective time frames.

7. Replication Results

As a result of the replication procedure, we have obtained a portfolio with an 86.8% gross return correlation to Manager X over the last 5 years and 80.4% correlation over the whole 15-year period. The table below also shows the changes of portfolio correlations during the replication process. Here, *First Run* means the portfolio after the first MFS optimization but before the first stage (as defined in Section 6d above). *Mid Point* is the result of the first stage, whose impact (mostly due to discarding the 14 markets) was comparable to the two remaining stages. *Final Portfolio*, naturally, shows the final results:

Correlations between Manager X and the Replication Portfolio

	Final Portfolio	Mid-Point	First Run
All 15 Years	**80.37%**	76.64%	70.70%
Last 5 Years	**86.80%**	82.08%	80.20%
Middle 5 Years	78.96%	77.46%	68.86%
First 5 Years	80.54%	75.32%	70.59%
Last Year	85.96%	78.14%	78.62%

We can see that not only the highlighted correlations, but also the correlations in all three 5-year periods and in the last year have significantly and similarly improved. This gives us additional comfort that the Portfolio is similar to Manager X.

Next, let us look at performance statistics. The leverage of the Portfolio can be easily changed. We adjust it so that the 15-year volatility of the Portfolio matches that of Manager X. The actual leverage of the Portfolio can be easily changed at the time of investing taking into consideration preferences of the investor. As we can see, the returns over the last 5-, 10- and 15-year periods are very similar:

Gross Performance Comparisons(1)

	Manager X	Replication Portfolio
Last 15 Years		
Annual Return	29.3%	29.7%
Volatility	29.5%	29.5%
Last 10 Years		
Annual Return	18.6%	23.8%
Volatility	25.6%	29.9%
Last 5 Years		
Annual Return	20.0%	20.0%
Volatility	24.0%	30.0%
Last Year	26.0%	28.7%

(1) *Volatilities are annualized, returns are non-compounded averages.*

Finally, we present a year-by-year comparison between the non-compounded annual returns of Manager X and those of the Portfolio. In our opinion, the results are consistent with the 80-85% level of correlation. Only in very few years do returns differ significantly, and even in those years the Portfolio does well compared to Manager X. In three such years, 2000, 1999 and 1997, the Portfolio outperformed Manager X by 21%, 11% and 24%, respectively. 1997 was the only significant underperformance year, although both Manager X and the Portfolio did well. While in 2005 Manager X also outperformed the Portfolio, it followed a much bigger loss in 2004. The total effect of 2004 and 2005 added together was that both Manager X and the portfolio were close to flat, with Manager X showing a somewhat higher return, while the portfolio had a much lower year-to-year volatility. The remaining years are very similar:

CONQUEST
CAPITAL GROUP

Gross Annual Returns of Manager X and the Replication Portfolio*		
	Manager X	Replication Portfolio
2006	26.0%	28.7%
2005	23.5%	-1.3%
2004	-18.4%	-5.1%
2003	31.2%	33.7%
2002	37.9%	44.2%
2001	13.4%	15.8%
2000	-4.0%	17.3%
1999	-1.0%	10.1%
1998	51.4%	55.5%
1997	25.8%	50.2%
1996	36.5%	42.2%
1995	75.4%	71.4%
1994	34.7%	25.8%
1993	62.7%	68.1%
1992	75.6%	90.0%

**The Replication Portfolio's leverage is chosen so as to achieve the same volatility over the full 15-year period as that of Manager X.*

We thus believe, based on the presented back-test performance, that we have successfully replicated Manager X.

Appendix A: Evolution of System and Market Allocation from MFS to Manager X

In this Appendix, we discuss in detail the evolution of the portfolio composition (systems, sectors, markets) from MFS with 6 additional markets to the 40-market replication portfolio for Manager X. In the following text and tables, the word "Stage" refers to the three stages of the replication portfolio creation as described in Section 6d.

a) Systems

As discussed above, system optimization needs to be approached very carefully. First, trend-following systems are highly correlated to each other. Even the 5-day and 200-day systems have a 22% correlation. This makes optimization difficult in principle, as the result becomes indifferent to a specific allocation. Roughly speaking, if data streams are almost the same, you can allocate to any of them with almost the same outcome. On the other hand, since all that we care about is the result, this is not that much of an issue. If equal allocations to, say, the 58- and 133-day systems give an effect very similar to allocating to the 88-day system, then we also should not be extremely concerned about choosing between these two allocations.

Second, the first problem would be a non-issue if we could claim that the optimization can be avoided altogether. For example, we avoid over-optimization of markets by allocating same weights to each of them. Yet, for markets we did it because CTAs do usually make their market allocations very balanced, if not the same, precisely in order to avoid over-optimization. For time frames, this is not true. Every money manager has its own time frame profile. Actually, each CTA system has its own profile. Those add together to make up the profile of the CTA. It can be unusual and non-even to any degree imaginable. Thus, the profile that we have to come up has no reason to look very smooth.

To balance these two issues, we cut the number of systems to ten, cut the number of optimizations to 2, and take whatever results the optimization gives. We also use an arbitrary limitation that no two system weights should have the ratio to each other exceeding 10. This is to make sure that our allocations do not become ridiculously small, which is meaningless practically. In our case, the lowest-allocation system was the 5-day one. We then manually set its weight to zero to see if it should be removed altogether. This did not produce any noticeable effect on the portfolio, and thus we have kept it for diversification. As we explained above, technically, each time frame exposure is present in any balanced portfolio due to the multitude of its entry, exit and filter exposures and cross-correlations among the time frames themselves.

The first optimization is done before Stage I, and the resulting allocation is kept until Stage III, when the second and final allocation is obtained. The table and the chart below show the results of the two optimization runs:

System Weight Evolution

	B005	B008	B011	B017	B026	B039	B058	B088	B133	B200
Actual MFS*	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Stage I to Stage III	0.52	0.69	0.87	0.99	1.08	1.17	1.17	1.17	1.17	1.16
Final Replication Portfolio	0.22	0.34	1.01	0.94	1.62	0.22	2.21	2.10	0.61	0.73

**MFS is based on 20 systems with different time frames and equal weights.*

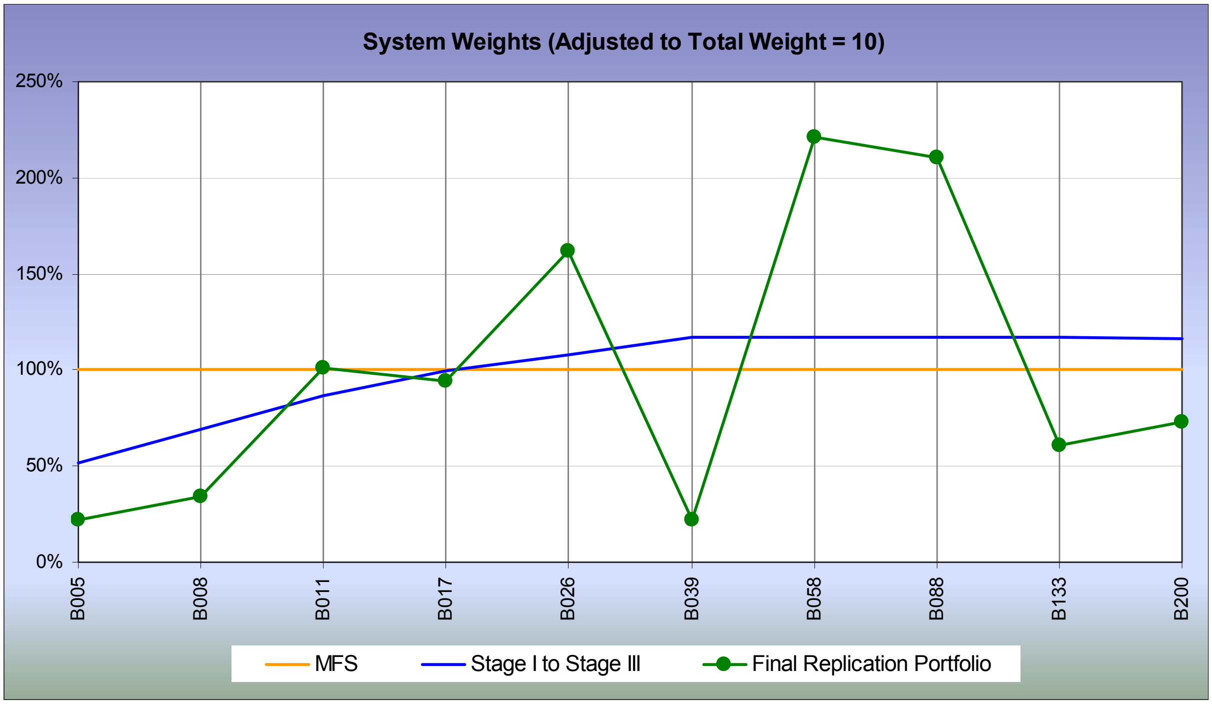

The short- and long-term parts of the curve behave very reasonably for a CTA balanced across time frames. Both drop off when getting to very short- and long-term values. The highest allocation is to mid- to long-term frames. The drop off at B039 is hard to explain, except that it does increase correlations. In any case, the allocation to any one system, taken out of context, is not important. For example, B039 is 81% and 86% correlated to the adjacent B026 and B058, respectively. Thus, increasing or decreasing its correlation can be easily offset by decreasing or increasing the correlations to its adjacent systems.

b) Sectors

As discussed in Section 6d, sector allocations were changed three times: while adjusting the MFS portfolio prior to Stage I, during Stage I and during Stage III. The latter two were done concurrently with market removal.

There are two kinds of parameters here that one should be concerned with. First, sector weights. These are the weights that are used to calculate trade sizes for each individual market in the sector. It is useful for individual market to market comparisons, for making sure that individual markets are not over-represented in the portfolio. Second, the relative weight of each sector as a whole. This is simply a sector weight multiplied by the number of markets in the sector. For example, in the final portfolio, stocks have a sector weight that is roughly 60% of the fixed-income weight. This means that each individual metal system/market pair will be represented with 60% as much weight as each currency system/market pair. However, whereas there are 6 equity markets, there are 9 fixed-income markets. Thus, as a total sector, fixed-income is represented (1/.6) * (9/6) = 2.5 times more than equities.

The following table presents the evolution of both weight categories through all three stages. The sum of the weights of all sectors is the same in each line within each of the two blocks. In the first block it equals to that of MFS, in the second it equals 100%.

Sector Weight Evolution						
	FX	**Stocks**	**Bonds**	**Energy**	**Metals***	**Commodities**
Sector Weights Per Market:						
Actual MFS	3.75	3.00	5.00	4.75	3.25	5.00
Before Stage I	3.82	3.50	7.17	1.80	6.32	2.15
Before Stages II & III**	4.04	4.83	6.93	2.24	4.53	2.18
New CTA Replication	3.90	3.97	6.35	1.55	5.75	3.23
Relative Total Sector Weights:						
Actual MFS	27%	15%	31%	12%	7%	8%
Before Stage I	23%	17%	38%	4%	14%	3%
Before Stages II & III**	29%	22%	32%	5%	9%	3%
New CTA Replication	28%	15%	37%	4%	11%	5%

**Lead, Nickel, and Zinc are each half-weighted in the final construction*

***Stage II changes the system weights holding the sector weights constant; therefore, Stages II & III are combined here.*

The progression of both parameters is not monotonic, since discarding individual markets changes both of them abruptly. Also, changes in relative sector weights are, naturally, affected by all other sectors. Yet, we can see that most parameters didn't change dramatically even between MFS and the final portfolio. They got to the levels very close to the final ones already during the approximate optimization prior to Stage I, before any market discarding. The sector that changed most was Energy. For FX, on the contrary, both parameters remained almost the same.

We would like to note that this table was created after the replication process, and the data presented in it were not taken into account during the process. The stability of the parameters in the table is thus a semi-independent check that the changes made were quite conservative.

The following charts visually present the difference between the initial and final sector allocations from the table:

MFS	Replication Portfolio

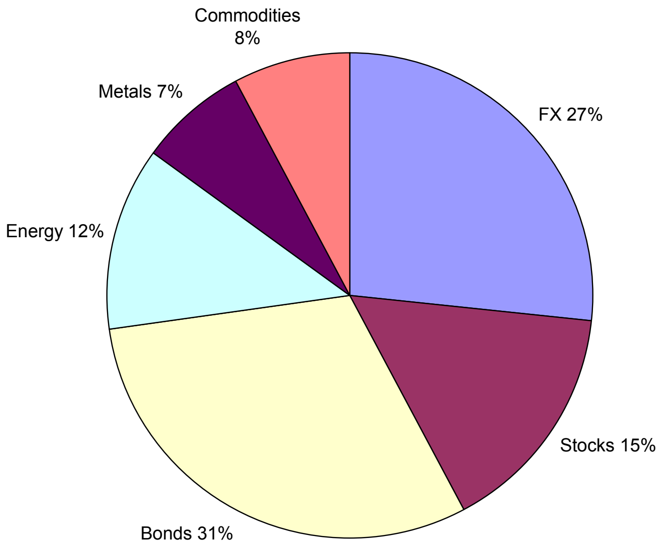

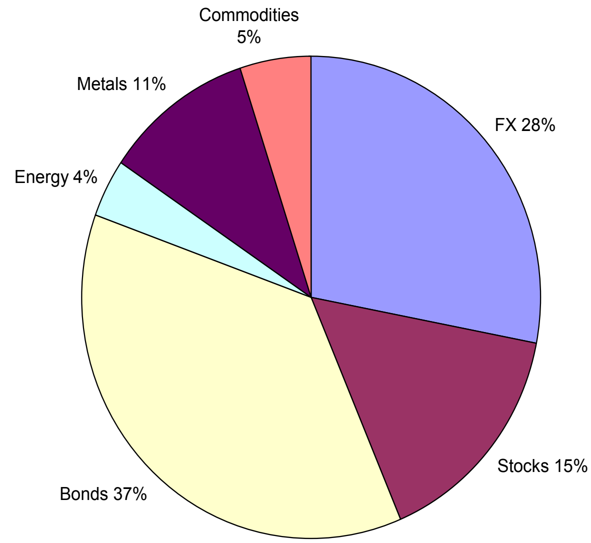

c) Markets

As discussed above, all markets in each sector have the same weights. Therefore, the market composition consists of just a list of markets by sector. The table on the next page shows the market composition evolution of the portfolio.

The full list contains the 61 markets that are the 51 MFS markets plus 6 additional markets. The additional markets are listed last in each column and are italicized. Those markets that are not highlighted are the markets remaining in the final Manager X replication portfolio. The markets highlighted in yellow are those that were discarded at the first stage of the portfolio construction. The markets highlighted in green were discarded at the final stage. Fourteen markets were discarded at the Stage I and seven more were discarded at Stage III, as described in Section 6d.

As mentioned in the main text of the report, all efforts were taken to keep as many markets as possible. The decision was made only when the market decreased correlation and performance, or significantly hurt one of them without increasing the other one.

Market Portfolio: Evolution from MFS to Manager X Replication Portfolio*					
FX	**Stocks**	**Fixed Income**	**Energy**	**Metals**	**Commodities**
Dollar Index	Russell	Eurodollar	Nat. Gas	Gold	Wheat
USDGBP	Nasdaq	Euribor	Light Crude	Copper	Soybeans
EURUSD	DAX	Short Sterling	Brent Crude	Aluminum	Corn
USDCHF	STOXX	Aus. Bills	Gas Oil	Nickel	Coffee
USDJPY	CAC	US 5Y	Heating Oil	Lead	Cattle
USDCAD	IBEX	US 10Y	Unleaded Gas	Zinc	Cotton
AUDUSD	Nikkei	US 30Y		*Silver*	*Sugar*
CADJPY	Topix	Bunds			
GBPJPY	Hang Seng	Bobl			
EURJPY	Taiwan	Gilts			
CHFJPY	*FTSE*	Japan. Bond			
EURGBP	*S&P 500*	Aus. 10Y			
NZDUSD		*Aus. 3Y*			
USDZAR					
USDMXN					
GBPCHF					
Original Number of Markets:		**61**			
16	12	13	6	7	7
After Stage I (First Removal)		**47**			
14	9	9	5	5	5
Final Number of Markets:		**40**			
12	6	9	4	5	4

**Markets highlighted in yellow were discarded at Stage I, in green - at Stage III. Italicized markets were not in the MFS portfolio.*