Multiperiod Static-Dynamic Factor Attribution

Morningstar Quantitative Research
Jan. 26, 2022

Contents

Gordon Sutton, Ph.D.
Senior Quantitative Analyst
gordon.sutton@morningstar.com

Patrick Caldon, Ph.D.

Lee Davidson, CFA
Head of Manager and Quantitative Research
+1 312 244 7541
lee.davidson@morningstar.com

Introduction

With the growing popularity and accessibility of factor investments, it is important to analyze fund managers' performance through the lens of factors. Do fund managers profitably time their exposures to risk factors? How much return comes from long-term positioning, as against short-term tactical changes? In this paper, we attempt to address these questions with a new multiperiod static-dynamic factor attribution methodology.

The methodology is broadly applicable to assessing most funds. In particular, the application of this methodology goes substantially beyond the obvious application—checking whether managers who claim to time factors can indeed generate excess returns through their short-term tactical exposure adjustments.

We can more generally seek to understand where the returns for any fund have come from. If the returns come from long-term exposures to particular factors, then any ability to maintain these returns will depend upon whether the factor is consistently priced. If the returns instead come from short-term allocations, and independently we have formed the view that these short-term allocation decisions were lucky rather than skillful and were unrelated to the overall investment process, we might not expect these returns to be replicated in the future. Conversely, if our view is that these excess returns were driven by the fund manager's process, we might expect the excess returns to continue.

The static-dynamic attribution methodology is also useful for understanding passive investing strategies, as passive strategies can give rise to unintended factor exposures. For instance, many exchange-traded funds give exposure to particular sectors. But these ETFs also have exposures to other changing factors. Are the returns from these ETFs coming from the constant exposures to particular sectors or catching a wave as particular sectors change their character? The observation generalizes to any ETF with nonconstant factor exposures—the vast majority of them. We can use this information to assist in assessing whether the historical return patterns will continue in the future.

At the heart of the methodology, a portfolio's daily factor returns are attributed to arising either from its average factor-exposure position or from deviations from this position. A technique that deals with the instance in which there is only monthly data is discussed later in this paper. Returns from the long-term exposure position are attributed to "static," and returns from the deviations are attributed to "dynamic." When the exposures and factor returns tend to deviate up and down together, this is considered good timing, and there is a positive dynamic attribution.

A multiperiod static-dynamic attribution is created from the daily static-dynamic attributions. The multiperiod static-dynamic attribution sums to the compounded return, accounts for cross-period interactions between factors, and is consistent with the standard multiperiod attribution in that each factor's static and dynamic attributions sum to the factor's standard attribution.

For both the single-period and multiperiod attributions, the structure of the static-dynamic attribution is

$$return = stock\ specific + \sum_{k=1}^{K} static_k + \sum_{k=1}^{K} dynamic_k$$

By splitting the standard multiperiod attribution to each risk factor into static and dynamic components, the methodology provides additional multiperiod measures of how the returns were achieved. By visualizing the evolution of the dynamic-attribution contributions, exposures, and premia, the methodology also provides further insight into when exposure deviations led to additional returns. Such time-series graphs can thus direct further analysis into determining which transactions were associated with impactful exposure changes, aiming to ascertain whether the exposure movements were deliberate or byproducts.

In the absence of trading, a portfolio's risk-factor exposures still change because stock price movements alter the portfolio's stock allocation weights and each stock's risk-factor exposures evolve. Therefore, it is the role of fund managers to be aware of this and to reposition the funds to where they want them to be. That is, doing nothing can be interpreted as an active decision.

We proceed by providing some background to the problem, then provide an overview of the mathematics behind the methodology, covering standard multiperiod attribution and the extension to multiperiod static-dynamic attribution. We then demonstrate that the method behaves as wanted on a set of artificial examples and on a more realistic example in which an investor switches from a value fund to a growth fund. We lastly apply the method to two funds separately and interpret the resulting static-dynamic attributions.

Background

The method is inspired by a paper by Andrew Lo, in which the output of a factor-based risk model is used to split a portfolio's expected daily returns into security selection, factor-timing, and risk-premia components. The first two components are considered active sources of return, and risk-premia is considered a passive source.

The idea is based on the observation that for any two time-series, the expected value of their product can be decomposed as

$$E[A_t \times B_t] = Cov(A_t, B_t) + E[A_t] \times E[B_t]$$

Applying this to a factor-based risk-model, with form

$$return_t = specific_t + \sum_{k=1}^{K} exposure_{tk} \times premia_{tk}$$

gives

$$
\begin{aligned}
E[return_t] &= E[specific_t] + \sum_{k=1}^{K} Cov(exposure_{tk}, premia_{tk}) + E[exposure_{tk}] \times E[premia_{tk}] \\
&= security\ selection + \sum_{k=1}^{K} factor\ timing_k + risk\ premia_k
\end{aligned}
$$

The methodology is elegant and accessible; however, it does not compound nicely into a multiperiod attribution. Over longer time periods, in which investors may be interested, it has problems. Since the method decomposes expected returns, it does not address cross-period interactions, and it fundamentally is decomposing simple returns rather than compound returns.

Our aim has been to develop a methodology that provides a multiperiod static-dynamic attribution. To that end, we want the following natural properties to hold:

- When a portfolio's exposure to a risk factor is correlated with the factor's premia, that is, they deviate up and down together, the risk factor will have a positive dynamic attribution.
- If a portfolio's exposure to a risk factor is kept constant, the dynamic attribution for the factor will be zero.
- The multiperiod static-dynamic attribution is consistent with the compounded return and with the standard multiperiod attribution, such that
  - The sum of all static-dynamic attribution components equals the total compounded return; and
  - Each factor's static and dynamic attributions sum to the factor's standard attribution.

Further, we want the standard multiperiod attribution methodology to be:

- consistent with the compounded return;
- commutative; that is, independent of the time ordering; and
- symmetric in its treatment of factor and stock-specific returns, so that the stock-specific does not become a dump for unattributed returns.

Methodology

We assume that there is a linear model of returns of the underlying securities in a fund, such as is produced by a cross-sectional multifactor-model regression. The Morningstar Risk Model and other vendors' multifactor models are suitable examples. However, this methodology is applicable to any similar linear model of returns. We progress to a standard attribution, from which the static-dynamic attribution is a small jump.

Linear-Return Model

The Morningstar Risk Model daily cross-sectional regressions decompose daily returns, r_t, into a sum of K risk-factor products, each with structure exposure × premia, and a stock-specific return, e_t, as follows

$$r_t = \sum_{k=1}^{K} x_{t-1,k} \times b_{t,k} + e_t$$

where $x_{t-1,k}$ is the risk-factor-k exposure at the end of day $t-1$, and $b_{t,k}$ is the risk-factor-k premia for day t.

In a multifactor model, the cross-sectional regression is evaluated for each day at the stock level, but the equation also holds at the portfolio level, after performing a weighted sum over a portfolio's holdings. We will focus on the portfolio level.

The model provides single-period attributions to K+1 buckets, corresponding to K risk factors and the stock specific, which we write as

$$r_t = \sum_{k=1}^{K+1} f_{t,k}$$

For some funds, holdings are reported daily; however, for many funds, they are only updated monthly or even quarterly. Also, some securities are not covered by the risk model. The attribution methodologies are based on the reported holdings that are covered, with the securities assumed to be held until the next set of holdings is reported. The attributions are created from the modeled portfolio returns and modeled daily factor exposures. In this paper, we refer to the modeled compounded return as the "compounded return."

Standard Attribution

The standard attribution methodology combines a sequence of such single-period attributions into a multiperiod attribution that sums to the compounded return. Our approach to standard attribution is to expand the equation for compounded return into a sum of cross-period return products:

$$\begin{aligned} R + 1 &= \prod_{t=1}^{T}(1 + r_t) \\ &= 1 + \sum_{t_1} r_{t_1} + \sum_{t_1<t_2} r_{t_1} r_{t_2} + \sum_{t_1<t_2<t_3} r_{t_1} r_{t_2} r_{t_3} + \cdots + \sum_{t_1<t_2<\cdots<t_{T-1}<t_T} r_{t_1} r_{t_2} \cdots r_{t_{T-1}} r_{t_T} \end{aligned}$$

and then attribute each product across the K+1 buckets. An example of a cross-period return product is $r_{t_1} r_{t_2} r_{t_3}$, which is a small component of the combined returns from times t_1, t_2, and t_3. Each r_t is a sum of attributions to the K+1 buckets, so each cross-period return product can be further expanded to a sum of many cross-period risk-factor-return products. For example, the three-time cross-period return product from times 1, 2, and 3, that is, $r_1 r_2 r_3$, expands to the sum of $(K+1)^3$ cross-period risk-factor-return products:

$$r_1 r_2 r_3 = \sum_{i=1}^{K+1} \sum_{j=1}^{K+1} \sum_{k=1}^{K+1} f_{1,i} f_{2,j} f_{3,k}$$

When all the elements of a product are from the same factor, that is, $i = j = k$ in the example, there would be little disagreement that the product should be entirely attributed to the one factor involved. However, when there is more than one risk factor contributing to the product, it is not so obvious how to distribute the product between them, and there is no definitively correct answer.

So, the conceptual complexity is deciding how to share, or attribute, each cross-period risk-factor-return product between the risk factors contributing to the product. Our approach is to split each product equally among its contributors. For example, a four-time cross-period risk-factor-return product involving risk factors (1, 2, 3, 1) is attributed with proportions (½, ¼, ¼) to risk factors (1, 2, 3).

The equal-splitting approach allows our three-time cross-period return-product example, $r_1 r_2 r_3$, to be factored into $3(K+1)$ terms as

$$\begin{aligned} r_1 r_2 r_3 &= \sum_{i=1}^{K+1} \sum_{j=1}^{K+1} \sum_{k=1}^{K+1} f_{1,i} f_{2,j} f_{3,k} \\ &= \frac{1}{3} \sum_{i=1}^{K+1} f_{1,i} r_2 r_3 + \frac{1}{3} \sum_{j=1}^{K+1} r_1 f_{2,j} r_3 + \frac{1}{3} \sum_{k=1}^{K+1} r_1 r_2 f_{3,k} \\ &= \frac{1}{3} \sum_{j=1}^{K+1} \left[(r_2 r_3) f_{1,j} + (r_1 r_3) f_{2,j} + (r_1 r_2) f_{3,j} \right] \\ &= \frac{r_1 r_2 r_3}{3} \sum_{j=1}^{K+1} \left(\frac{f_{1,j}}{r_1} + \frac{f_{2,j}}{r_2} + \frac{f_{3,j}}{r_3} \right) \end{aligned}$$

After settling on an approach to splitting each cross-period risk-factor-return product, there just remains the mountainous accounting task of tallying all the attributions, which becomes surmountable via a dynamic-programming-like approach. The result is that the multiperiod attribution, A_k, to each of the $K+1$ buckets conveniently has the form

$$A_k = \sum_{t=1}^{T} L_t \times f_{t,k} \qquad 1 \le k \le K+1$$

where L_t is the link function for time t. L_t only depends on the r_t sequence, rather than the attribution of each r_t to buckets. The relationship of L_t with r_t is slightly nonlinear and concave up, which produces an L_t sequence that is close to constant and positive. The practical implication of L_t being approximately constant is that each time period receives approximately the same weight in the multiperiod attribution. These link-function properties are demonstrated in Exhibits 1a and 1b, respectively showing the link-function as a time series and the link-function vs daily returns.

Exhibit 1a Link Function — Parnassus Mid Cap Growth Investor

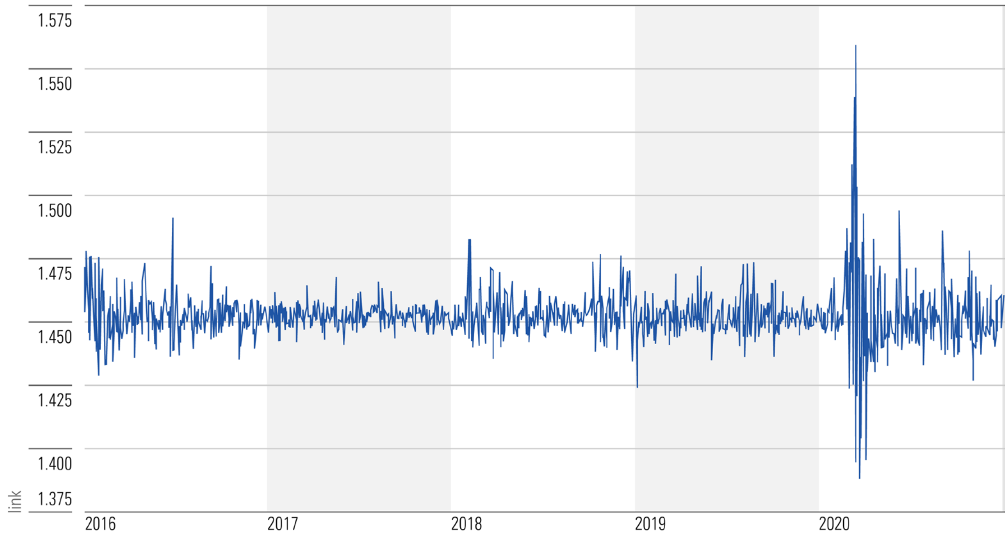

Source: Morningstar. Data as of Aug. 13, 2021.

Exhibit 1b Link Function vs. Daily Return — Parnassus Mid Cap Growth Investor

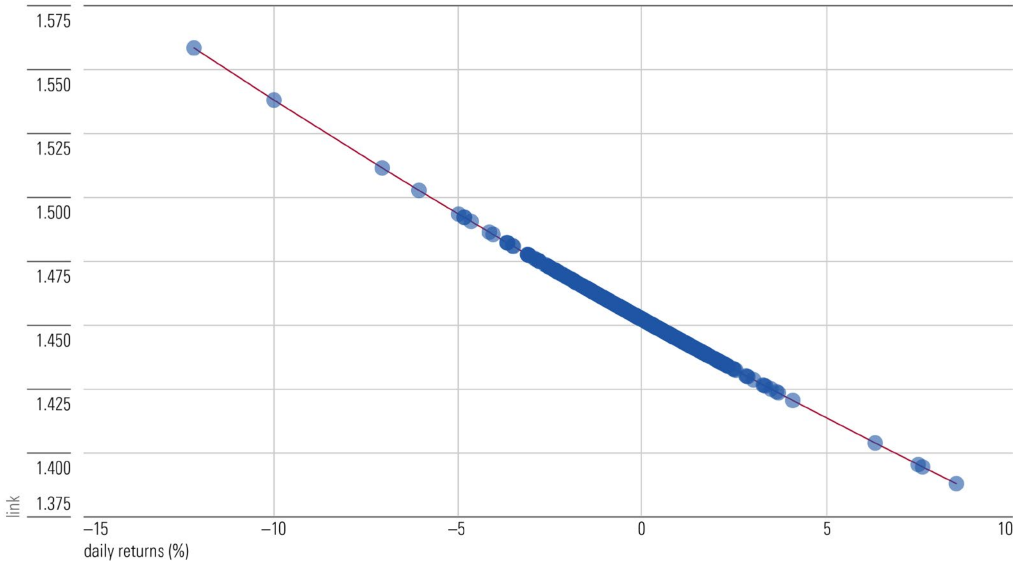

Source: Morningstar. Data as of Aug. 13, 2021.

The methodology inputs a $T \times (K + 1)$ array of daily bucket returns $f_{t,k}$ and outputs a multiperiod attribution to the $K + 1$ buckets. By construction, the multiperiod attribution sums to the compounded return, is independent of the order of daily returns, and treats the K+1 buckets, including the stock-specific bucket, symmetrically.

While we have developed this methodology independently, the method appears to have been previously presented, with an alternative notation, in the Reztsov working paper (2011).

Multiperiod Static-Dynamic Attribution

We now consider the extension to multiperiod static-dynamic attribution. The basic idea is to define a static-exposure position and then attribute the daily returns that the static-exposure position generates to static and the balance of the portfolio's daily returns, because they deviate from this position, to dynamic.

This follows the same basic idea in Lo's 2008 paper of splitting factor returns into passive and factor-timing components. Lo decomposes the expected daily returns and thereby avoids contemplating cross-period returns. However, one of our aims is to achieve a multiperiod attribution, so the mechanics differ. Mathematically, compounding the expected daily return will be higher than the compounded return. That is, deviations from the expected return reduce the compounded return. For example, $1.1 \times 0.9 = 0.99$, which is less than 1, the average return squared.

Lo treats the exposures and premia symmetrically in the decomposition equations, based on the assumption that they are both stochastic processes. When decomposing expected daily returns, this does not matter much, but when attributing multiperiod returns, it does have an impact.

Our premise is that a fund/portfolio manager has the option to control the exposures, but the premia are not controllable. So, while the factor returns are the product of exposures and premia, the exposures and premia need not be, or rather should not be, treated interchangeably in the methodology. Instead, we consider being passive as holding the constant static-exposure position. This means the daily passive returns arising from the static-exposure position will be time-varying and that, on a given day, if a factor's exposure equals the static exposure, there will be zero attribution to dynamic on that day.

Another possibility would be to define the daily static attributions as the average exposures times the average premia, as appears in Lo's decomposition of expected daily returns. However, this would be inconsistent with our standard attribution methodology. Since the link function is negatively correlated with returns, having constant daily static attributions would lead to the multiperiod static attributions being too high. Further, it could result in a portfolio with constant exposure to a risk factor having a nonzero multiperiod dynamic attribution to that risk factor. That is, although the daily dynamic attributions may sum to zero, the link-function-weighted sum may be far from zero.

So, starting with a risk-model's daily cross-sectional regression, we define a static-exposure position, $\bar{x}_k$, as the portfolio's expected exposure over the considered time span, giving

$$\bar{x}_k = \frac{1}{T}\sum_{t=1}^{T} x_{t-1,k}$$

The expectation of the exposures—that is, the static-exposure position—could be defined differently, such as an industry average or an average over a different time span, but the within-time-span average of the fund's own exposure is our current method. We can then calculate the daily static-factor returns generated from this static-exposure position for factor k as

$$f_{t,k}^{static} = \bar{x}_k \times b_{t,k}$$

The daily dynamic-factor returns, $f_{t,k}^{dynamic}$, are obtained by subtracting the static-factor returns from the total factor returns for each of the K factor buckets. That is, $f_{t,k}^{dynamic} = f_{t,k} - f_{t,k}^{static}$, or in terms of exposure deviations from $\bar{x}_k$,

$$f_{t,k}^{dynamic} = (x_{t-1,k} - \bar{x}_k) \times b_{t,k}$$

Recalling that the link function only depends on the sequence of total daily returns, the final step is to apply the link function to each of the now $2K$+1 buckets (K static, K dynamic, and the stock specific), giving the multiperiod attributions

$$A_k^{static} = \sum_{t=1}^{T} L_t \times f_{t,k}^{static} \qquad 1 \le k \le K$$

$$A_k^{dynamic} = \sum_{t=1}^{T} L_t \times f_{t,k}^{dynamic} \qquad 1 \le k \le K$$

$$A_{k+1} = \sum_{t=1}^{T} L_t \times f_{t,k+1}$$

Note that the stock-specific attribution remains as given by the standard attribution.

Since the same link function is used throughout, it can be seen that the multiperiod static-dynamic attribution is consistent with the standard attribution, that is, $A_k^{dynamic} + A_k^{static} = A_k$, and is thus also consistent with the compounded return.

The daily returns are split into $f_{t,k}^{static}$ = E[exposure] x premia, and $f_{t,k}^{dynamic}$ = (exposure – E[exposure]) x premia. Taking expected values gives $E\left[f_{t,k}^{static}\right]$ = E[exposure] x E[premia], and $E\left[f_{t,k}^{dynamic}\right]$ = Cov(exposure, premia), which are properties that we want and are also terms found in the decomposition given in Lo. The multiperiod attribution weights each day's returns by the link function to account for the interactions that arise when the daily returns are compounded. As such, the multiperiod static-dynamic attribution is not just a simple scaling of $E\left[f_{t,k}^{static}\right]$ and $E\left[f_{t,k}^{dynamic}\right]$. However, since the link function is close to constant and only slightly nonlinear, the multiperiod static-dynamic attributions are approximately a simple scaling of $E\left[f_{t,k}^{static}\right]$ and $E\left[f_{t,k}^{dynamic}\right]$, and the dynamic attribution for each factor does closely reflect the correlation between the exposures and premia, as wanted.

Exploratory Examples

To show that the multiperiod static-dynamic attribution behaves as we want, we start with three artificial single-factor examples based on the risk factors: value-growth, consumer cyclical, and momentum, over the five-year time span from 2016-20 inclusive. The premia are taken from the Morningstar Global Industry Standard Risk Model.

Example 1a

In our first exploratory example, we consider a step-change in exposure to the value-growth risk factor. Higher exposures imply the fund is more growth- and less value-oriented. We assume we have a portfolio that is exposed to just the value-growth risk factor and has no stock-specific returns. For the exposures, we create a unit step change at September 2018 and set its mean to the average value-growth exposure of the fund Federated Hermes MDT Small Cap Core IS over the time span. The components contributing to the static-dynamic attribution are shown in Exhibit 2.

The style of Exhibit 2 is repeated throughout the examples. The yellow solid line is the exposure, the yellow dotted line is the mean exposure, and the area between is shaded yellow to highlight when the exposure is above/below the mean. The blue line is the cumulative-sum premia, and the red line is the cumulative-sum contributions to the multiperiod dynamic attribution. We note that while compound premia might seem more normal, it is the cumulative-sum premia that more directly helps interpret the cumulative-sum contributions to dynamic.

Exhibit 2 Value-Growth — Exploratory Example: Step-Change Exposure

Source: Morningstar. Data as of Aug. 13, 2021.

Starting with a quick interpretation, the exposure is below average prior to the step change, when the value-growth returns were mixed, and above average after the step change, when the value-growth returns were generally positive. We would like this well-timed exposure-increase to be reflected in the attribution, and it is. As shown in Exhibit 3, the total returns are 3.7%, which is attributed as negative 0.4% to static and 4.1% to dynamic.

Exhibit 3 Step-Change Exposure — Exploratory Example: Static-Dynamic Attribution

	Total	Static	Dynamic
Value-Growth	3.696	-0.433	4.129
Total	3.696	-0.433	4.129

Source: Morningstar. Data as of Aug. 13, 2021.

Looking into the details, the daily contributions to dynamic are the product of the exposure deviation from the mean, the daily premium, and the link function on the day. Before the exposure step change, the exposure is below its mean and constant, near negative 0.5, and the link function is near 1 and always approximately constant. The cumulative contributions to dynamic thus move in the opposite direction to the cumulative premia, at approximately half the rate. After the exposure step change, the exposure is above its mean, near 0.5, so the opposite holds. That is, instead of the contributions to dynamic and premia approximately mirroring each other, they approximately move in tandem.

The daily contributions to static are the product of the average exposure, the daily premium, and the link function on the day. The average exposure is negative, while the cumulative premia is positive, so the net attribution to static will be negative.

Example 1b

For the next exploratory example, we assume all returns are generated by the consumer-cyclical sector risk factor. We consider a square-wave exposure to consumer cyclical that changes between one and zero, with an approximately 22-month period, resulting in an average exposure over the time span of 0.45. The components contributing to the static-dynamic attribution are shown in Exhibit 4.

Exhibit 4 Consumer Cyclical — Exploratory Example: Square-Wave Exposure

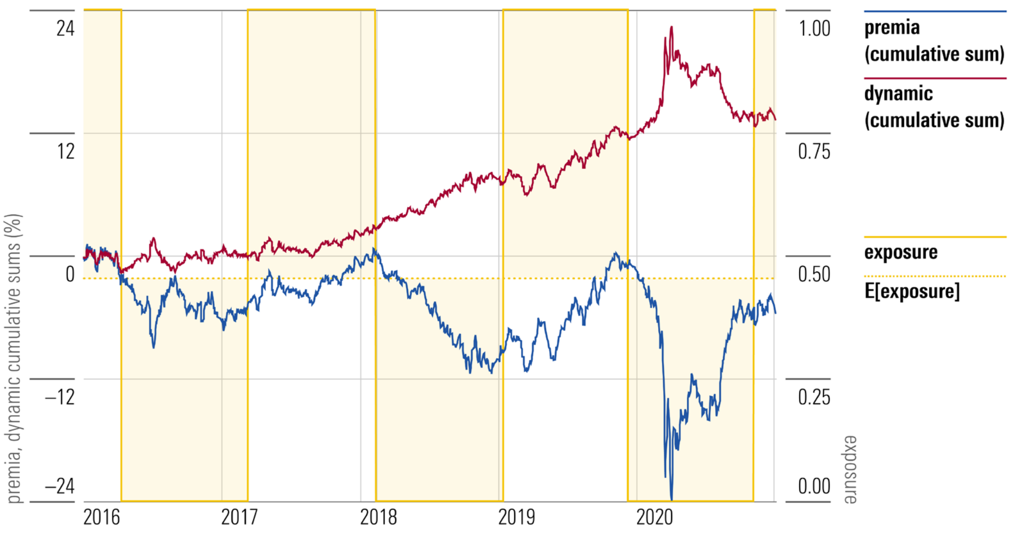

Source: Morningstar. Data as of Aug. 13, 2021.

Each of the positive-exposure periods coincides with a period of generally positive consumer-cyclical premia, and the troughs coincide with periods of generally negative consumer-cyclical premia. From this correlation, we would expect the attribution to dynamic to be positive, and indeed it is, which is quantified in Exhibit 5.

It is worth noting that while the attribution to static is negative 2.9%, the return from maintaining the average exposure over the time span is negative 2.85%. The reason for this small difference is that the higher net return delivered by the square-wave exposures magnifies the daily static returns through the cross-period interactions.

Exhibit 5 Square-Wave Exposure — Exploratory Example: Static-Dynamic Attribution

	Total	Static	Dynamic
Consumer Cyclical	10.270	-2.909	13.179
Total	10.270	-2.909	13.179

Source: Morningstar. Data as of Aug. 13, 2021.

Example 1c

For the third exploratory example, we use the momentum risk factor to explore whether defining being passive in terms of exposure, without directly addressing the premia, is problematic. This time we use a constant premium, setting it to the mean momentum premium from the Global Industry Standard Risk model over the time span. For exposures, we take the momentum exposures from the fund Federated Hermes MDT Small Cap Core IS and add 0.5. The components contributing to the static-dynamic attribution are shown in Exhibit 6, and the attribution is given in Exhibit 7.

Exhibit 6 Momentum — Exploratory Example: Constant Premium

Source: Morningstar. Data as of Aug. 13, 2021.

The constant premia create zero-correlation between the momentum exposures and premia. We want this to produce zero attribution to dynamic. While the daily contributions to dynamic are nonzero, over the whole time span they net to effectively zero, as wanted.

It is observable that the red line is wobbling up and down. It is doing so because it is effectively integrating the exposure deviations from the mean exposure, which integrate to zero by construction. The local maximums and minimums of the contributions to dynamic can be seen to correspond to exposure zero-crossings. Each day in the integration is weighted by the link function, and the link function is slightly nonlinear, so this zero-integration is approximate. However, the link function is close enough to constant and its nonlinearity is slight enough that the integration is effectively zero, as wanted.

Exhibit 7 Constant Premium—Exploratory Example: Static-Dynamic Attribution

	Total	Static	Dynamic
Momentum	3.942	3.942	0.000
Total	3.942	3.942	0.000

Source: Morningstar. Data as of Aug. 13, 2021.

From the three exploratory examples, we can see that the static-dynamic attributions are reflecting the presence, or absence, of correlations between factor exposures and premia, while producing a consistent multiperiod attribution. So, we conclude that the static-dynamic attribution is generally behaving as wanted.

Example 2

We next consider a more realistic example of a fictitious investor deciding to switch from a value fund to a growth fund. We assess the static-dynamic attribution over the same five-year time span, 2016-20 inclusive, with the investor switching from LSV Value Equity to MainStay Winslow Large Cap Growth A at the end of August 2018.

This time, the attribution calculations include all the stock exposures and premia, as produced by the risk model, and we evolve the portfolio holdings between rebalancing/reporting dates by assuming the stocks are held constant.

Exhibit 8 gives a summary of the multiperiod static-dynamic attributions for the five risk factors with highest dynamic attribution and for the entire portfolio. The Portfolio Total row tallies the attributions for all 33 risk factors and the stock specific. The stock-specific returns are not attributed to any factor, so only contribute to the total column. The top five risk factors contribute 34.6% to the 35.5% Portfolio Total static-dynamic attribution.

Exhibit 8 Static-Dynamic Attribution—Investor Switches Fund

	Total	Static	Dynamic
Value-Growth	10.853	-0.492	11.345
Australian Dollar	29.108	21.903	7.205
Technology	27.7	21.047	6.654
Energy	2.04	-3.6	5.64
Yield	4.24	0.507	3.734
Portfolio Total	127.048	107.063	35.523

Source: Morningstar. Data as of Aug. 13, 2021.

We explore the value-growth, technology, and energy risk factors, which together contribute 23.6% to the dynamic attributions. Graphs of the components contributing to their static-dynamic attributions are given in Exhibits 9, 10, and 11, respectively. Note that the dynamic cumulative sums have been plotted with different magnifications for each graph, as indicated. We leave the Australian dollar to the next example.

Exhibit 9 Value-Growth — Investor Switches Fund

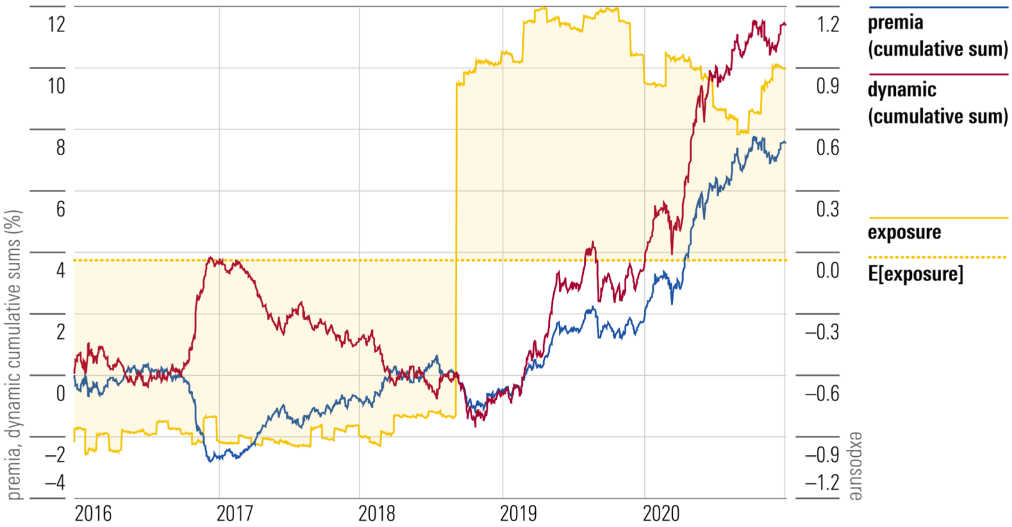

Source: Morningstar. Data as of Aug. 13, 2021.

Exhibit 10 Technology — Investor Switches Fund

Source: Morningstar. Data as of Aug. 13, 2021.

The fictitious investor's aim was to switch from value to growth, that is, increase the value-growth exposure. The fund switch also caused an increased exposure to technology and a decreased exposure to energy. Looking at Exhibit 9-11, the switch caused an approximate step-change in the exposure for all three factors, up for value-growth (-1, 1), up for technology (0.15, 0.4), and down for energy (0.12, 0.01). Also, the premia accumulated during the periods before and after the switch were approximately (0, 8) for value-growth, (14, 36) for technology, and (10, -40) for energy. So, without looking at the detail, in each case, the lower/higher factor exposure occurred with the lower/higher accumulated premia. That is, for each factor, the exposures were positively correlated with their corresponding premia, so we would like, and would expect, each factor to have a positive dynamic attribution. This can be observed in

Exhibit 8, for which the attributions are calculated accurately. The static-dynamic cumulative sums for value-growth, technology, and energy are repeated in Exhibit 12 so they can be seen on the same scale.

The methodology continues to behave as wanted.

Exhibit 11 Energy—Investor Switches Fund

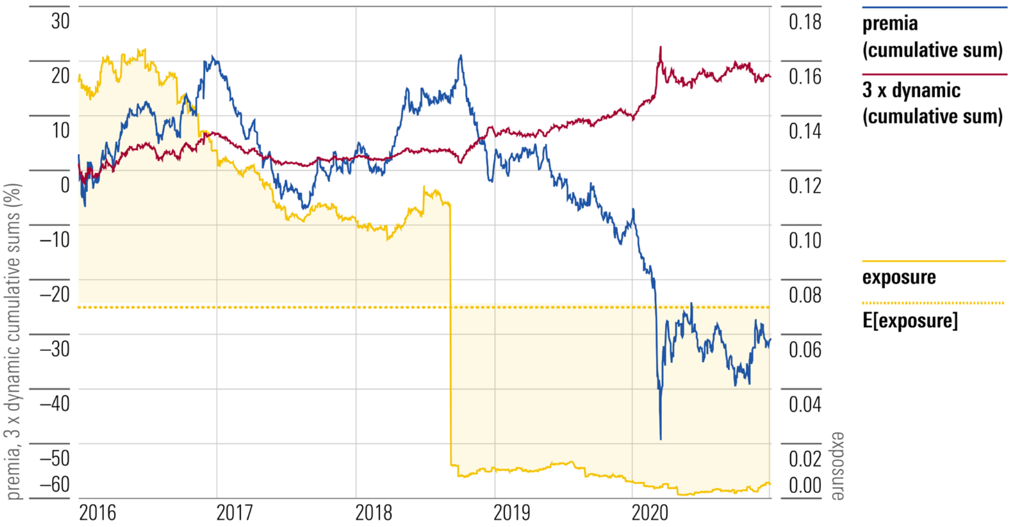

Source: Morningstar. Data as of Aug. 13, 2021.

Exhibit 12 Dynamic Attribution Contributions—Investor Switches Fund

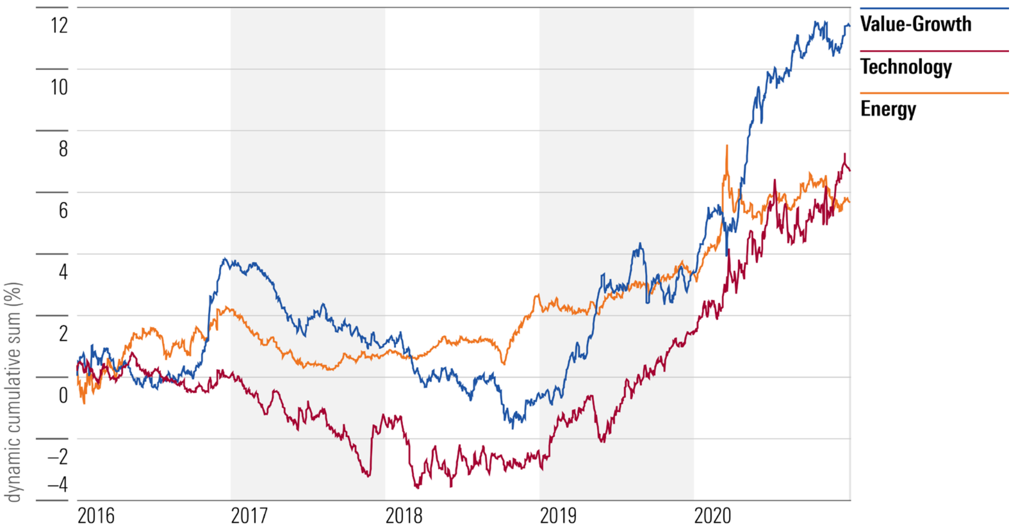

Source: Morningstar. Data as of Aug. 13, 2021.

Example 3

Now that we have demonstrated that the methodology is able to identify well-timed exposure positions, we consider a single fund, T. Rowe Price New Horizons Fund. We again calculate the static-dynamic attribution over the five-year time span, 2016-20 inclusive. The portfolio returns and attribution analysis are based on first evolving the reported portfolio holdings between rebalancing/reporting dates by assuming the stocks are held constant.

The static-dynamic attributions of the risk factors with the five highest and five lowest dynamic attributions are presented in Exhibit 13, along with the stock-specific attribution and the fund total.

Exhibit 13 T. Rowe Price New Horizons Fund: Static-Dynamic Attribution

	Total	Static	Dynamic
Australian Dollar	29.262	20.333	8.929
Healthcare	10.816	6.309	4.508
Technology	22.952	19.567	3.385
British Pound	8.321	6.143	2.178
Value-Growth	13.778	11.923	1.855
Euro	-7.027	-6.015	-1.013
Quality	-1.711	-0.585	-1.127
Developed North America	23.147	24.330	-1.183
Volatility	-9.504	-7.299	-2.205
Japanese Yen	-9.916	-6.628	-3.287
Specific	58.448	0.000	0.000
Fund Total	190.568	114.264	17.856

Source: Morningstar. Data as of Aug. 13, 2021.

The fund almost tripled its value over the five years, generating 191% total return, or 23.8% per year. Of the 191% total return, 17.9% was attributed to dynamic. As a high-level measure of manager performance, we can define the timing ratio as "dynamic attribution"/"total returns," which can be compared between funds. For T. Rowe Price New Horizons Fund, the timing ratio is 17.9/190.6 = 9.4%. The stock-specific and dynamic attributions can both be considered the result of active management and sources of active returns. As done in Lo, we define the active ratio as "active returns"/"total returns." For T. Rowe Price New Horizons Fund, the active ratio is (17.9 + 58.4)/190.6 = 40%.

We will delve into what generated the static-dynamic attributions for the three factors with highest dynamic attributions: the Australian dollar, healthcare, and technology. The components contributing to their static-dynamic attributions are respectively shown in Exhibits 14, 15, and 16.

Note that the Australian dollar factor is essentially a proxy for the Chinese economy in the context of this model. To interpret the Morningstar Risk Model currency factors, understand that the exposures are betas to currency indexes. So, an exposure to a currency factor means that the stock's returns are correlated with the exchange-rate movements. These betas can arise directly or indirectly. They can arise directly from two possible sources:

- Investing in foreign companies, which to some degree is investing in a foreign economy; and
- Investing in a domestic company where, perhaps through subsidiaries or selling into foreign markets, it has economic exposure to a foreign currency.

In some circumstances, there are confounding effects that generate the betas indirectly. This makes simple causational interpretations problematic. In particular, the Australian dollar exposure is widely used as a proxy for the performance of the Chinese economy—for this reason, many companies have their returns correlated with the Australian dollar but without any relationship to the Australian economy.

The static-dynamic attributions from exposures to the Australian dollar factor—probably a proxy for the Chinese economy—are explored in Exhibit 14. The average exposure is negative 0.36, which generated a static attribution of 20.3% because of the poor performance of the Australian dollar factor over the five-year time span. The changes in exposure through the time span lead to an 8.9% dynamic attribution.

Exhibit 14 Currency Australian Dollar—T. Rowe Price New Horizons Fund

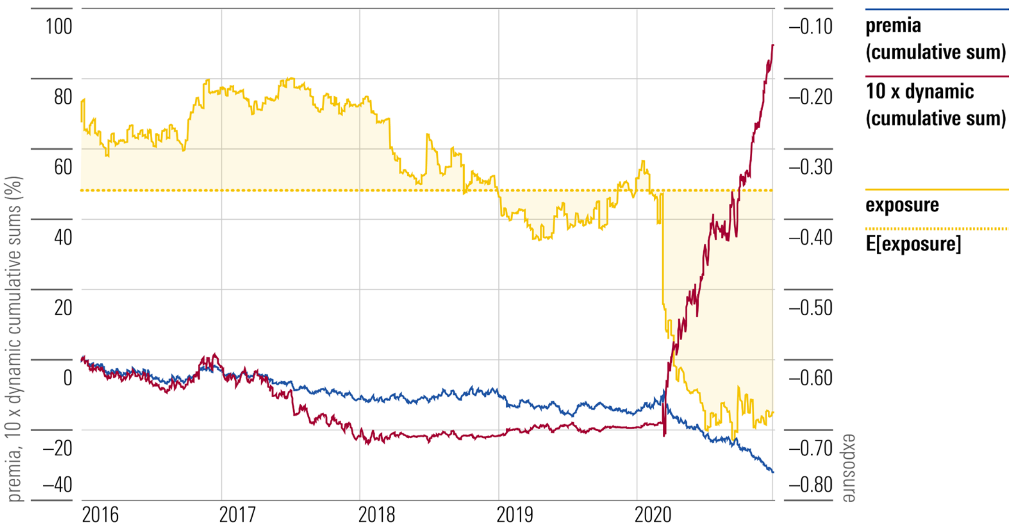

Source: Morningstar. Data as of Aug. 13, 2021.

The T. Rowe Price New Horizons Fund is a small-to-mid cap U.S. fund, with only a tiny proportion of its assets based in China. As such, it may seem surprising that there is a 29% attribution to the Australian dollar exposure. However, the currency factors cover more than the direct revenue exposures. For example, the exposures may be due to the cost of raw materials being correlated with the Australian dollar movements or the general interlinking of the world economy. The currency exposures are the betas of regressing each stock's historical returns against seven major exchange-rate historical returns and an intercept, but what causes the betas, or correlations, is not revealed and will differ between stocks. The attribution suggests that there are indirect mechanisms at work.

We can explain the dynamic attribution by splitting the five-year span into four periods: pre-2017, 2017, 2018 to early 2020, and post-early-2020. The Australian dollar performance during these periods can be simplified to flat, down 10%, flat, and down 20%, while the corresponding exposures were: slightly above average, above average, mixed, and well-below average, which resulted in approximate dynamic-attribution contributions of: 0.0%, negative 2.0%, 0.0%, and 10.9%. The exposure reduction in the final year predominantly occurred at the onset of the Australian dollar premia commencing a long negative run, so the contributions to dynamic persisted through the year.

The dynamics of the T. Rowe Price New Horizons Fund's exposure to healthcare is explored in Exhibit 15. Note that the dynamic cumulative sum has been multiplied by 5. During 2016, the exposure was ramped from below average back to near average, while healthcare returned negative 20%. Increasing the exposure positioned the fund well for the rebound and contributed positively to the dynamic attribution. From 2017 to 2019, the exposure was kept close to the fund's average, so although healthcare returned 25%, little was attributed to dynamic. During 2019, the exposure was ramped up further, while the returns fluctuated, resulting in a modest contribution to dynamic for 2019. At the start of 2020, the exposure was at its peak. Soon after, the novel coronavirus hit and healthcare rocketed, so the exposure increases through 2019 had positioned the fund well. The fund cashed in its healthcare position and reduced its exposure back to its average position just before healthcare began a sustained fall that continued through the remainder of the year. Increasing the exposure before the peak and getting out close to the peak contributed positively to dynamic.

There were several instances of apparently good pre-positioning before a shift occurred. Whether they were intentional or not, the result was a 4.5% attribution to dynamic on top of a 6.3% attribution to static.

The analysis suggests to ask questions about what strategic decisions were made in early 2019. It turns out that there was a change in manager around this time. Arguably, the healthcare-exposure ramping commenced prior to the manager change, but it appears to have continued strongly after the change, so perhaps both managers can claim some credit.

Exhibit 15 Healthcare—T. Rowe Price New Horizons Fund

Source: Morningstar. Data as of Aug. 13, 2021.

The timing of the fund's exposure to technology is explored in Exhibit 16, in which the dynamic cumulative sum is multiplied by 10. Technology performed strongly during the five-year time span, returning 50%. The fund's average exposure was 0.22, producing a 19.6% attribution to static. A further 3.4% return was attributed to dynamic technology exposures, giving a total technology attribution of 23.0%. A broad dynamics interpretation is that dropping the exposure mid-2016 to below average resulted in both missing the 2017 positive run and avoiding the 2018 fall, which contributed negative

1.0% to dynamic, and increasing the exposure several times in late 2018 to above average resulted in capturing the stronger positive run from 2019 to 2021, which contributed 4.4% to dynamic.

Exhibit 16 Technology—T. Rowe Price New Horizons Fund

Source: Morningstar. Data as of Aug. 13, 2021.

Exhibit 17 presents the dynamic cumulative sums for the Australian dollar, healthcare, and technology on the same scale. For each of the three factors explored, more than half of the contributions to their dynamic attributions were added during the last year, which includes the emergence of COVID-19. For these risk factors, T. Rowe Price New Horizons appears to have favorably navigated the emergence of COVID-19. The changes in the fund's exposure positioning for healthcare and technology appear to have commenced prior to COVID-19. So, the exploration raises the question of what management decisions were made in late 2019 and early 2020. They appear to have been successful, but were they due to good luck or good management?

Exhibit 17 Dynamic Attribution Contributions—T. Rowe Price New Horizons Fund

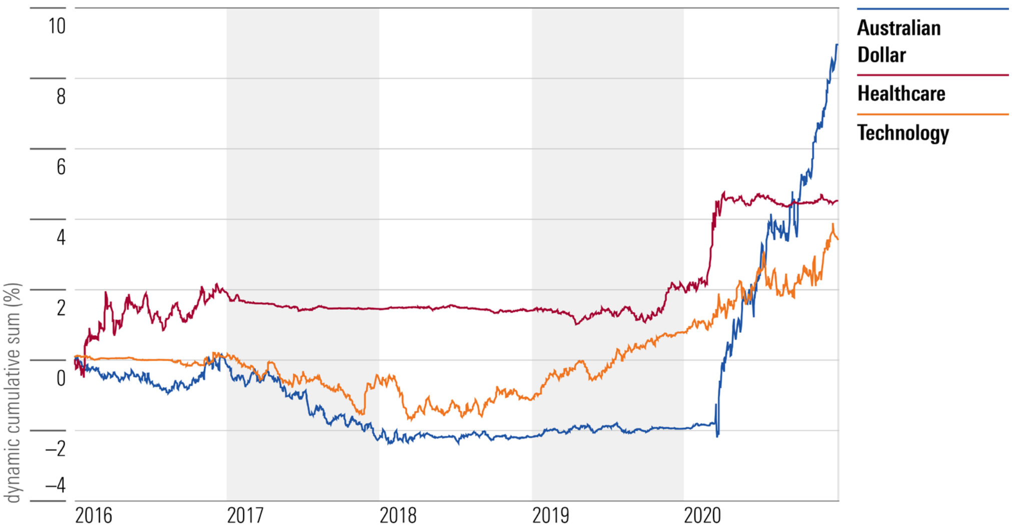

Source: Morningstar. Data as of Aug. 13, 2021.

Example 4

As a final example, we consider the fund Federated Hermes MDT Small Cap Core IS. To avoid any effects associated with COVID-19, we consider the earlier five-year time span, 2015-19 inclusive.

The static-dynamic attributions of the risk factors with the five highest and five lowest dynamic attributions are presented in Exhibit 18, along with the stock-specific attribution and the fund total. The five-year return was 62%, equating to a 10.1% annual return. The overall attribution to dynamic was negative 1.5%, giving a timing ratio of negative 2.4%. The net dynamic attribution to the style factors was 4.3%, but this was offset by a negative 5.8% net dynamic attribution to the currency factors. The stock-specific attribution was 11.6%, giving an active ratio of 18.7%. In this case, good stock selection is more apparent than good factor timing.

We delve into how the interplay between exposures and premia generated the static-dynamic attributions for the three factors with highest dynamic attributions: momentum, size, and the British pound sterling exchange rate. The components contributing to their dynamic attributions are respectively shown in Exhibits 19, 20, and 21.

Exhibit 18 Federated Hermes MDT Small Cap Core IS: Static-Dynamic Attribution

	Total	Static	Dynamic
Momentum	5.821	1.624	4.197
Size	10.734	8.504	2.230
British Pound	1.782	0.173	1.609
Emerging Latin America	1.742	0.654	1.088
Energy	-2.528	-3.600	1.072
Australian Dollar	1.652	2.773	-1.121
Liquidity	1.728	3.155	-1.426
Japanese Yen	-8.128	-6.584	-1.544
Euro	-3.439	-1.218	-2.221
Canadian Dollar	-0.176	2.292	-2.468
Specific	11.562	0.000	0.000
Fund Total	61.948	51.891	-1.504

Source: Morningstar. Data as of Aug. 13, 2021.

In Exhibit 19, the fund's exposure to the momentum risk factor can be seen to swing approximately yearly, reaching positive 0.6 and dropping to negative 0.8. Until mid-2018, the exposure swings were well-timed, being positive during periods of generally positive momentum returns and negative during periods of generally negative momentum returns. From mid-2018 to mid-2019, the momentum exposures were mistimed, with positive exposure deviations during a momentum drop and negative exposure deviations during the recovery. Over the time span, the exposure dynamics generated an extra 4.2% return compared with having constantly held the static-exposure position.

Exhibit 19 Momentum — Federated Hermes MDT Small Cap Core IS

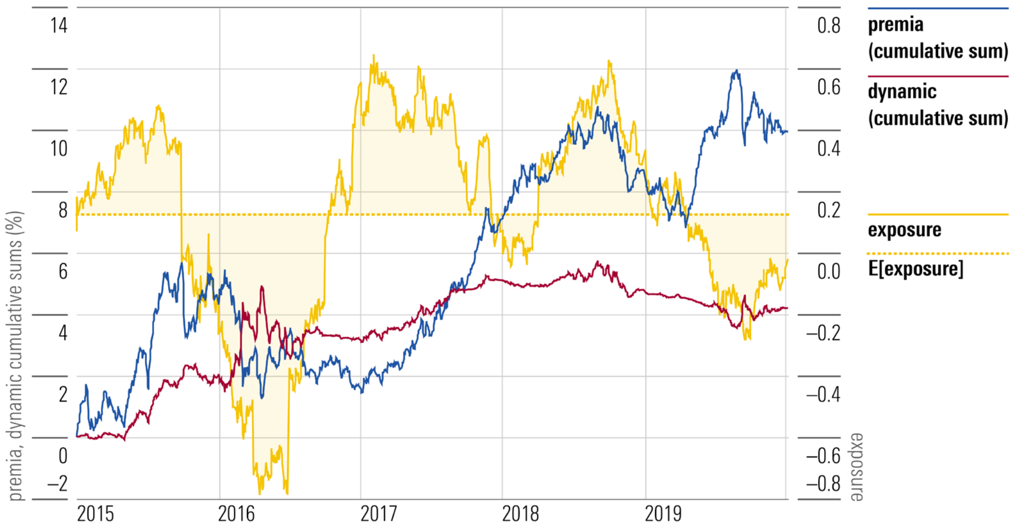

Source: Morningstar. Data as of Aug. 13, 2021.

The fund is concentrated in small-cap stocks but still varied the weighted-average size of the companies in its holdings, leading to variation in its size exposure, as can be seen in Exhibit 20. The fund generally timed the size exposure adjustments well, gaining an extra 2.2% dynamic attribution on top of the 8.5% attributed to the static-exposure position. The periods of good timing were due to the higher exposures during 2016 and the lower exposures during the second half of 2018. Other periods had less impact on the dynamic attribution because either the exposures were closer to the static exposure or the cumulative premia were relatively flat.

Exhibit 20 Size — Federated Hermes MDT Small Cap Core IS

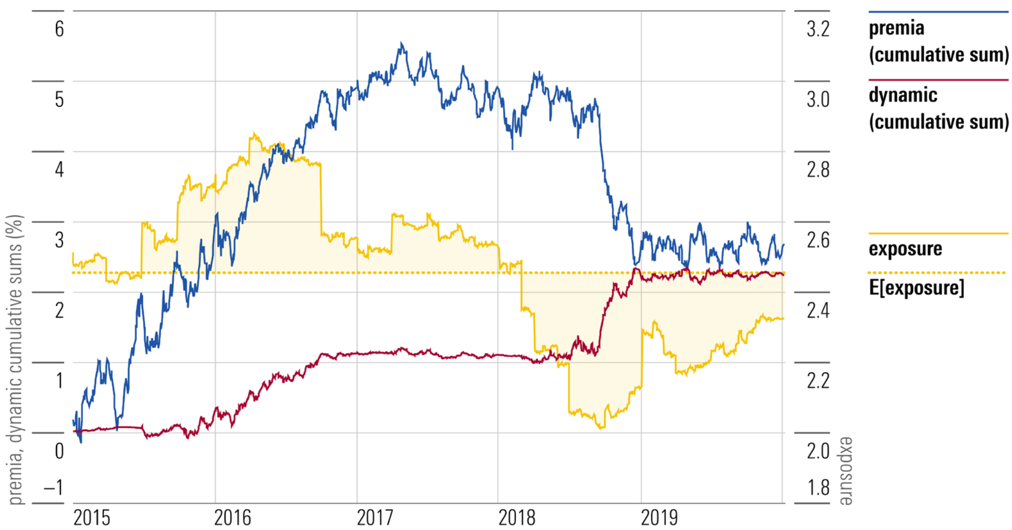

Source: Morningstar. Data as of Aug. 13, 2021.

The British pound is explored in Exhibit 21. The fund was initially unexposed to the British pound and, during the first year, moved to a relatively constant exposure position near negative 0.2. This made the initial position above average, which was also when the British pound premia were doing well, so this period contributed positively to the dynamic attribution. After mid-2015, the British pound exposures and cumulative premia were both close to their means and mostly uncorrelated, so further contributions to the dynamic attribution were small.

Exhibit 21 British Pound—Federated Hermes MDT Small Cap Core IS

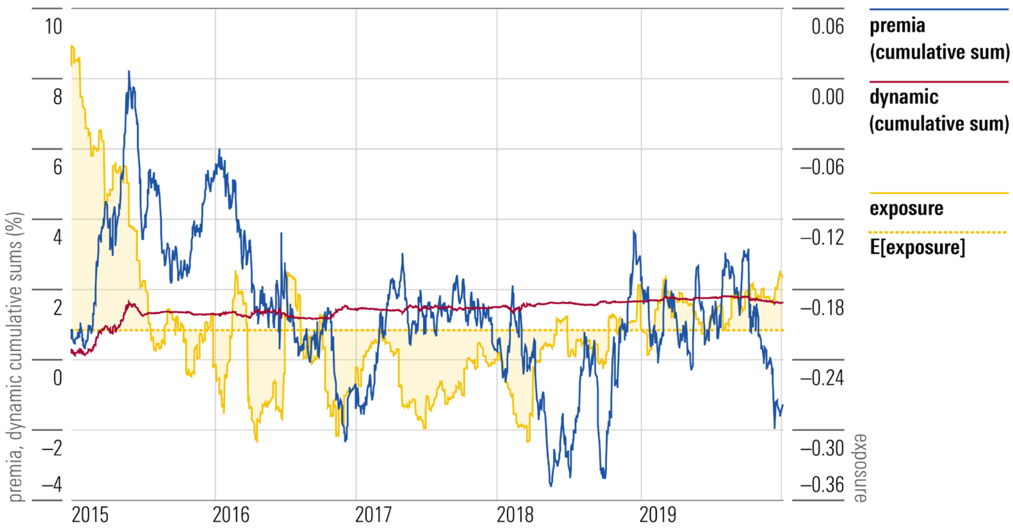

Source: Morningstar. Data as of Aug. 13, 2021.

To show more detail, the cumulative static-dynamic contributions for momentum, size, and the British pound are repeated in Exhibit 22.

Exhibit 22 Dynamic Attribution Contributions—Federated Hermes MDT Small Cap Core IS

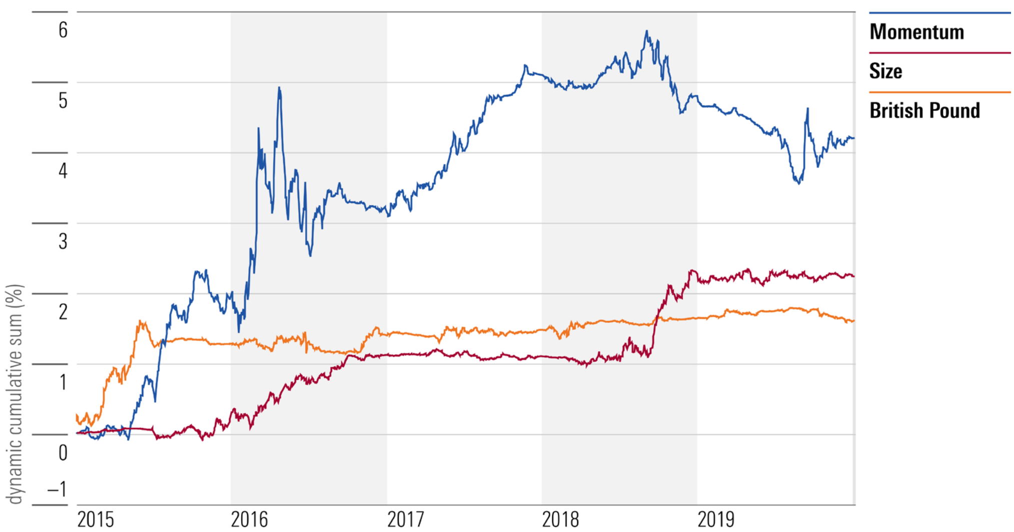

Source: Morningstar. Data as of Aug. 13, 2021.

Conclusion

We have presented a new static-dynamic attribution methodology that enhances our standard multiperiod attribution of returns to factors by further decomposing each factor's multiperiod attribution into static and dynamic components. The method extends the attribution of expected daily returns presented in Lo to a multiperiod attribution. The methodology splits the attribution to each factor into either arising from holding a static average-exposure position or arising from the dynamic exposure deviations from the average exposure. The idea is that if a fund increases its exposure to a factor just before the factor does well, this is attributable to good exposure timing, and vice versa.

We justified our approach and presented the properties that we would like from a multiperiod static-dynamic attribution. The attribution to dynamic corresponds closely to the correlations between each factor's exposures and premia; if they tend to deviate up and down together, a positive dynamic attribution is generated.

There are some potential pitfalls, like with any attribution technique. Nevertheless, although it is not always interpretable, it can still have value. For example, it can guide what questions to further explore. There are many alternative approaches to attribution. Hence, any contributions and insights are worthwhile. Our multiperiod static-dynamic method provides an extra measure of performance that can help locate and analyze changes in exposure positions.

We demonstrated that our methodology detects factor-exposure timing when it appears to be present and generally behaves as we want, first on a set of single-factor exploratory examples and then on a realistic, though constructed, example in which an investor abruptly changes a portfolio from a value fund to a growth fund.

We then explored the performance of two funds over five-year time spans. While the resulting dynamic attributions within a single fund were smaller than in the constructed examples, we were able to identify the changes to each explored factor's exposure positions that generated the positive dynamic attributions.

The static-dynamic attribution methodology can also be used to compare the performance of fund managers. When assessing multiple funds of the same nature across the same analysis time span, the timing ratio and active ratio can be used as relative measures of management skill. The timing ratio gives the percentage of returns attributable to factor-exposure dynamics, and the active ratio gives the percentage of returns attributable to active management, which encompasses both factor-exposure dynamics and specific stock selection.

The static-dynamic attribution is an additional tool that can help with understanding what is going on in a fund. Do the manager's stated objectives line up with how the returns are observed to be generated? It creates a launching point for asking managers questions. Why did they change their exposure position when they did? Was the timing by accident or design?

While we can detect what we call factor timing, the question remains as to whether positive factor timing is due to luck or skill. Further research will be needed to explore this within a statistical framework. For the time being, we can use the numbers to initiate questions.

References

Lo, A.W. 2008. "Where Do Alphas Come From?: A Measure of the Value of Active Investment Management," *Journal of Investment Management,* Vol 6., No. 2, P. 1-29.

Reztsov, A. 2011. "Excess Returns–Arithmetic and Geometric?," available at SSRN 1858937.

About Morningstar® Quantitative Research

Morningstar Quantitative Research is dedicated to developing innovative statistical models and data points, including the Morningstar Quantitative Rating, the Quantitative Equity Ratings and the Global Risk Model.

For More Information

+1 312 244-7541
lee.davidson@morningstar.com

MORNINGSTAR®

22 West Washington Street
Chicago, IL 60602 USA