



A version of this article appeared in the Summer 2011 issue of *Journal of Investing*.

Executive Summary

The transition management industry standard for measuring total costs is unquestionably implementation shortfall (IS) as codified by the T Standard and as further formalised in the T charter. The adoption of the IS methodology as the starting point for calculating transition performance and the subsequent standardisation of this approach through the “T” initiatives have marked key milestones in the evolution of transition management as a formalised enterprise.¹ Naturally, no methodology, regardless of its elegance, can capture every dimension of performance measurement for portfolios in dynamic markets. IS, of course, is no exception. As has been well documented, while IS may represent the most comprehensive and sound representation of the total cost of a given transition, it cannot in isolation shed light onto whether a reported IS figure is evidence of a good result or not, and says little to the question of whether the transition manager was able to add additional value through implementation. Different transitions will naturally incur different levels of cost, so although minimising IS must always be the overriding objective, a low shortfall number does not necessarily indicate a good result, nor a large shortfall number a poor result.

Trading Benchmark Choice and Transition Management Performance Attribution

By

Graham C. Cook, CFA, FRM
Vice President and Transition Manager
BNY Mellon Beta Management

This article discusses the role and importance of trading benchmarks in transition management, particularly with respect to the interpretation of performance and the attribution of implementation shortfall. First, the need for finer detail on implementation shortfall is reviewed. Next, the role of trading benchmarks within a transition management implementation exercise is discussed along with the impact that the choice of benchmark can have on attribution. Then the factors influencing the choice of benchmark are explored followed by a critical appraisal of commonly used benchmarks. The insights arising from this critique are then developed to present a conceptual framework for determining the optimal execution trajectory for orders that contain no *a priori* informational content. In the concluding section, this is applied to formulate new benchmarks for evaluating the quality of equity trading and the measurement of market impact. Both of these measures are based on participation rates, namely the negligible impact participation rate (NIPR) and the optimal trajectory participation rate (OTPR). Finally, the issue of unobserved latent market impact is broached, along with suggestions for how this can be identified and factored into a full attribution of implementation shortfall.

¹ The “T-Standard” and subsequent “T-Charter” industry initiative formalise a code of best practice in the clear and transparent presentation of pre trade estimates and post trade performance, as well as making compulsory certain important disclosures for compliance with the codes.

Transaction Cost Attribution in Transition Management

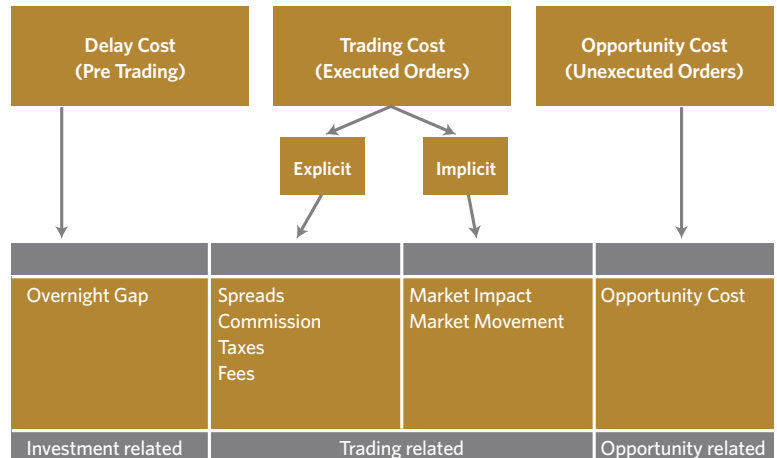
Essentially, the orders that are required for a fund restructure are the result of long-term investment decisions, and not subject to the short-term alpha capture considerations that are important in active portfolio changes.

It is important to note that when describing an order as having no *a priori* informational content, what we really mean is that there is no information contained within the orders that is likely to show itself in any meaningful way within the time taken to complete the order. This is clearly demonstrated within the context of a portfolio restructure or transition exercise. In many cases, these involve the removal of an asset manager or asset class expected to underperform in the medium to long term, and the replacement by one that is expected to outperform over the same period. Whilst there may well be useful information here, its realisation is expected to occur over many months or years, and rarely over the hours or days in which the resulting orders are implemented. Essentially, the orders that are required for a fund restructure are the result of long-term investment decisions, and not subject to the short-term alpha capture considerations that are important in active portfolio changes.

Typically, a transition manager is tasked with implementing the client investment decision in as cost-efficient and risk-controlled manner as possible. During the proposal phase, the transition manager will likely estimate for the client the expected transaction cost for the restructuring with a variance around that central case estimate that is derived from the predicted tracking error of the legacy portfolio to the target over the horizon of the transition. Following the completion of trading, the transition manager will report back to the client an IS number with varying degrees of performance attribution.

Exhibit 1 below shows how total implementation shortfall can be broken down in a way consistent with Andre Perold's original definitions, and along the lines suggested by Robert Kissell.²

Exhibit 1 - The Anatomy of Implementation Shortfall



Source: BNY Mellon Beta Management, 2011.

Note: For illustrative purposes only.

² See Perold (1988) and Kissell (2006).

Indeed, as is typical for any transition where the target and legacy assets are not closely related, more of the final IS result could be put down to the largely random factor of realised tracking risk than any other component cost.

Since the fixed (explicit) costs are well known but account for only a small proportion of total costs (typically less than a fifth), they provide little insight as to how well costs were managed. Implicit costs on the other hand account for a much larger proportion of total costs. As these costs are highly variable, they offer a greater potential for cost management.

One approach for such an evaluation compares the actual cost relative to the estimated cost. This is done in a variety of ways,³ including:

- **Z Scores** — can be used to compare the difference between the estimated and realised costs, where the denominator is the period tracking error between the legacy and target portfolios over the expected implementation period. This Z score can then be evaluated against standard, published statistical tables, and its significance reported.
- **“Within Range” Analysis** — here the actual shortfall result is compared to the range of outcomes detailed in the pre trade report.
- Alternatively, the **T-Ratio** could be computed. This effectively measures how closely the actual result in terms of implementation shortfall was to the central estimate of the cost, and is then scaled by the tracking error around this central point estimate.⁴

The problem here is in comparing the final result against an estimate that often has very low explanatory power due to the frequent dominance of the realised tracking risk in determining the absolute level of implementation shortfall. Indeed, as is typical for any transition where the target and legacy assets are not closely related, more of the final IS result could be put down to the largely random factor of realised tracking risk than any other component cost. That’s not to say that the choice of transition manager or their performance doesn’t matter, just that it takes a large number of observations to distinguish skill from luck.

Alternatively, the raw IS could be normalised for difficulty by applying an adjustment factor that takes into account the inherent potential for market impact and tracking risk in the transition portfolios along the lines suggested by Kritzman et al (2006). Under this methodology, the realised implementation shortfall is adjusted by a ratio calculated by adding the ratio of realised and expected explicit costs to the ratio of realised and expected tracking error, where realised tracking error is estimated by subtracting the expected market impact from the realised implicit costs. Other approaches in this vein would include stripping out market wide movements in prices to arrive at a market-adjusted shortfall. This is achieved by subtracting a beta-adjusted market return from the total implicit costs.³ Neither of these approaches give any granular detail as to where the costs came from, and whether these were in

³ See Kissell (2008), pp32-34, for more detail and a good review of the range of approaches.

⁴ See Clay (2007).

line for each component (explicit costs, delay, market impact, market movement or opportunity costs) that could provide the asset owner with insights into the skills of the transition manager or even attribute to the investment, trading and opportunity categories of the process that Kissell (2006) suggests.

A more common presentation of the breakdown of implementation shortfall for an equity restructure would focus on the attribution between the four components of implicit costs: delay, market movement, market impact and opportunity costs.

A more common presentation of the breakdown of implementation shortfall for an equity restructure would focus on the attribution between the four components of implicit costs: delay, market movement, market impact and opportunity costs. The focus is placed on the implicit costs since the explicit costs are generally fixed, whereas the implicit costs are variable and therefore may offer opportunities for reduction through careful management of implementation. When implementation shortfall is calculated in its strict form and in accordance with the T standard the delay cost is an artefact of the metric, and cannot be managed;⁵ it is, to all intents and purposes, a random element. Opportunity cost generally has only a negligible contribution to the total cost in a typical transition, since wherever possible, all trades tend to be completed in full, and there is little discretion whether or not to execute an order except in the face of extreme adverse price movements. In order to gain additional insight into the costs, therefore, the key attribution is between market impact and market movement.

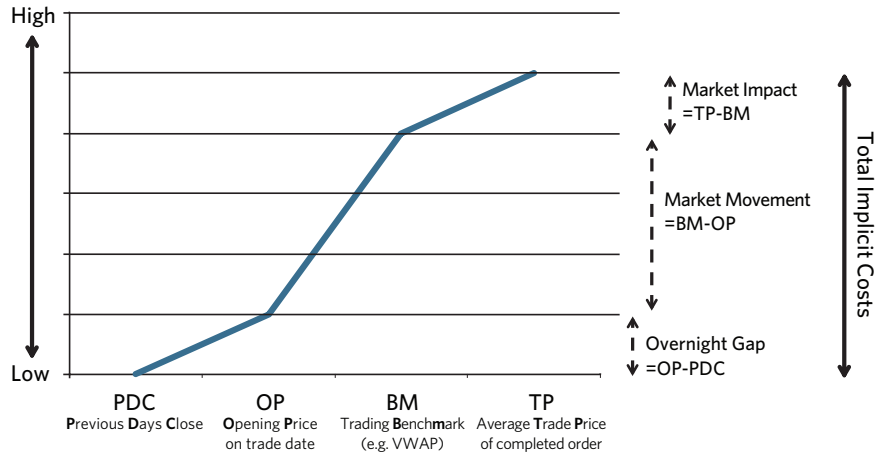
A serious hurdle to overcome in performing this attribution is that during a period of trading, neither market impact nor market movement (realised tracking risk) are directly observable. Remember that market impact is defined as the difference between the price achieved through trading and the price that would have prevailed had trading not taken place. Similarly, market movement is defined as the change in price of the security that we are trading that is due to (exogenous) factors unconnected to the trading of the order during the period of execution. For a transition, this is the realised tracking error between the transition and target portfolios. To be able to discern market impact from market movement, it is necessary to utilise proxies for one or other of the variables. The proxy for market impact can be derived from a model of expected costs, which is common in the academic literature, or, more common in practice, by using a measure of trading performance against a passive trading benchmark.

⁵ Some might argue that overnight gap could be mitigated through hedging (such as with a futures overlay) at the previous nights close, but if we follow the T Standard and IS in their strict forms, the base benchmark prices are set for the day prior to any trading starting which the hedge would violate.

Exhibit 2 illustrates the implicit cost of executed orders in this framework.

Given how these two components of implicit costs are commonly viewed (i.e., market movement is unavoidable and outside of the control of the trader whilst market impact is seen as the direct responsibility of the trader), the choice of the trading benchmark, and therefore how the attribution is presented, can have a significant effect on the perceived performance of the trade execution.

Exhibit 2 - Implementation Shortfall Attribution



Source: BNY Mellon Beta Management, 2011.

Note: This illustration is based on a proposed conceptual framework around the attribution of implicit trading costs, and does not reflect actual trading.

The base point for measuring implementation shortfall is invariably the previous nights close. From this, the overnight gap, in addition to the explicit costs, is essentially formulaic and incontrovertible.

The remaining attribution is purely a function of the trading benchmark chosen. Given how these two components of implicit costs are commonly viewed (i.e., market movement is unavoidable and outside of the control of the trader whilst market impact is seen as the direct responsibility of the trader), the choice of the trading benchmark, and therefore how the attribution is presented, can have a significant effect on the perceived performance of the trade execution. And yet this choice receives little attention, especially within transition management where the default trading benchmark chosen by the transition manager to de-compose his shortfall result tends to be either volume weighted average price (VWAP) or a version of arrival price, regardless of the actual trading strategy employed for the transition.

Trading Benchmarks in Transition Management

First and foremost, it needs to be made clear that what is being described here refers to metrics to be used to help decompose implementation shortfall, *not* metrics intended to replace IS. For that reason, it bears repeating that in the absence of an otherwise specified objective, the goal of transition management is to minimise (if not eliminate) the implementation shortfall incurred.

Benchmarks for measuring trading performance should also be weighted averages rather than single points for any order that is larger than the average bid or ask volume depth, i.e., any order that would require more than one trade to complete.

Within transition management, trading benchmarks have two main functions and a corollary function:

1. To measure the success or quality of trading; i.e. how effectively did my transition manager's trading desk execute?
2. To inform the trading strategy or style to be employed.

Additionally, in the above framework, such benchmarks can play a role in the attribution of trading costs but are useful only in situations where certain criteria are met.

In the next section, the factors that *should* drive the selection of the trading benchmark, if they are to effectively achieve all three functions listed above, are identified. Likewise, a summary of some of the pros and cons of using each of the common trading benchmarks is presented as a prelude to the following section in which the concept of the NIPR benchmark is introduced.

Factors in Determining an Appropriate Trading Benchmark

The ideal trading benchmark should be consistent with the trading strategy employed. Note that trading strategies and benchmarks are not one and the same, although it is essential they are closely aligned if the benchmark is to be of use. This would also include being relevant to the trading period, i.e. the trading benchmark is calculated for the same period during which trading takes place.

Benchmarks for measuring trading performance should also be weighted averages rather than single points for any order that is larger than the average bid or ask volume depth, i.e., any order that would require more than one trade to complete.

Benchmarks should also be transparent and replicable from public price information. Their method of calculation should be explicitly set in advance, and ideally, they should be achievable.

Finally, and most importantly if any meaningful insights are to be attained, the benchmark should be independent of the trading to which it is being compared. That is to say that the benchmark should neither be susceptible to gaming nor be influenced by the trading being implemented. To ensure strict compliance with this, it would be necessary to either have a predetermined benchmark, the level of which would be known before trading starts, or else to strip out the trades being analysed from the benchmark calculation.

The main advantage of using pre-trade benchmarks is that they are the only class of benchmark that is truly free of the potential for gaming since they are determined in advance and are not influenced by the act of implementation itself.

Common Trading Benchmarks

As has been noted extensively elsewhere, trading benchmarks can be categorised into three classes:

- **Pre-Trade**, such as the previous day's close, the open and arrival price.
- **Post-Trade**, such as the close.
- **Trading Period Averages**, such as VWAP, time-weighted average price (TWAP), participation-weighted average price (PWAP) or the high-low-open-close average (HLOC).

The important advantages and challenges are broadly applicable to each specific class of benchmarks, with a few quirks, so let's take each of these in turn.

Pre-Trade

The most widely used pre-trade benchmarks are Open, Previous Close and Arrival Price with the former two being most commonly used in transition management. The pre-trade benchmark fits most closely with the framework of implementation shortfall and usually corresponds to some proxy of decision price.

The main advantage of using pre-trade benchmarks is that they are the only class of benchmark that is truly free of the potential for gaming since they are determined in advance and are not influenced by the act of implementation itself. The flipside to this however is that they do not represent the market conditions during trading and therefore are not particularly useful for measuring the trade execution performance of the transition manager. In addition, they use a single point as reference, and are unrelated to the execution strategy for a transition, since these can rarely be executed instantaneously except through a guaranteed principal bid. The exception to all of this is the Open benchmark, which is not strictly a pre-trade benchmark when trading can take place on the open. As such it is highly susceptible to influence, with the transparency of market impact on the open completely lost in markets where the opening price is set through an auction. One further point to note with using the Open as a trading benchmark is that when attribution is calculated in the way outlined above, all implicit costs are attributed to market impact.

Post-trade benchmarks are rarely applied in transition management but are often used by institutional investors (such as index funds) with inflows and outflows to the portfolio that are themselves assumed to take place at the close.

Post-Trade

Invariably, the benchmark used will be the closing price on the day of trading, sometimes referred to as the terminal or exit price.

Post-trade benchmarks are rarely applied in transition management but are often used by institutional investors (such as index funds) with inflows and outflows to the portfolio that are themselves assumed to take place at the close. As Madhavan⁶ notes, a post-trade benchmark promotes trading at the close, when prices are most sensitive to impact. Not only does such trading incur higher impact costs, the visibility of these costs is largely lost since the impact is incorporated into the benchmark. In addition, implementation shortfall is likely to be much higher and more variable as the orders incur a full day's market movement risk.

In contrast to the Open benchmark, using the above framework for attribution with the close as the benchmark attributes all of the implicit costs to market movement, which is likely to be extremely erroneous given the point just made on market impact.

Trading Period Averages

This class of benchmarks implicitly acknowledges both that an institutional order is not typically executed in full instantaneously, and that single point benchmarks are prone to correspond to extreme and unrepresentative price levels contrary to the objective of a robust benchmark. Trading period averages attempt to overcome these issues by providing a level that is representative of the market conditions facing a trader over the period in which the order is executed. Most are weighted averages to further improve the integrity of the benchmark. The main exception to this is the open-close-high-low average. Much used in years gone by when tick level data was rare and computing power a fraction of what is available currently, this benchmark is rarely used by sophisticated institutional investors of late. This is the most arbitrary of this class of benchmarks, and as such, the most difficult to consistently achieve. As intimated above though, it is most useful in markets where high quality and high frequency data are not available to facilitate the accurate calculation of weighted benchmarks.

By far the most common of the trading period average benchmarks in use today is the ubiquitous VWAP, largely due to the seminal paper of Berkowitz, Logue and Noser back in 1988, coincidentally the same year that Andre Perold published his ground breaking paper introducing the concept of implementation shortfall. In it, the authors showed that the VWAP was the best unbiased approximation for the price available to a passive trader with the simple objective of completing an order during the trading day. This, then, provides the baseline from which to compare active and passive traders alike, with skilled active traders expected to beat the benchmark on a consistent basis.

⁶ See Madhavan (2002), pp34-35.

The two main criticisms of the VWAP benchmark are that it is easy to game, and that it disincentivises traders to achieve best execution through a careful trade off between market impact and the risk of adverse market movement, instead leading them to spread even the most liquid orders out over the day in order to minimise the deviation from the benchmark price.

In VWAP's favour is its simplicity, as well as its widespread availability, which aides transparency (and therefore integrity), and the fact that it is by far the most prevalent benchmark which itself yields the comfort of the herd. However, there are serious pitfalls in an uncritical use of VWAP as either a benchmark or as a trading strategy, which often follows as a direct consequence of its use as a benchmark. These issues can compromise the achievement of best execution and the interpretation of trading performance. In addition, care needs to be taken in the calculation of VWAP if its use is to yield valid and useful insights.

The two main criticisms of the VWAP benchmark are that it is easy to game, and that it disincentivises traders to achieve best execution through a careful trade off between market impact and the risk of adverse market movement, instead leading them to spread even the most liquid orders out over the day in order to minimise the deviation from the benchmark price. Unconstrained illiquid orders are likely to have very significant market impact, not just on the executions, but on the benchmark itself. Where these orders are a significant percentage of the total trading volume for the day, they effectively set the benchmark, negating a very important criterion for an effective benchmark, namely independence of influence from what it is measuring.

The debate about how much VWAP can be gamed by traders is well documented and does not really need repeating here. Except to the extent of influencing the benchmark noted above, it is not as much of an issue in transition management, given the need to complete the required trades in full; most of the gaming revolves around not completing orders when the current price is worse than the VWAP price.

Summary of Potential Deficiencies of VWAP as a Trading Metric:

- Unrelated to optimal trading strategy
- Uses whole day as trading benchmark for small liquid orders
- Easy to game
- Meaningless where the order represents a high percentage of the traded volume

In our opinion, the clearest view comes when working in volume-time space, since it is the combination of volume and time which determines the potentiality for costs.

Much misunderstood or overlooked is the fact that there are many ways to calculate VWAP. The selection of the most appropriate methodology can help alleviate some of the issues noted above that detract from its effectiveness. The basic definition of VWAP is the average price of all trades executed that day weighted by the number of shares traded at that price. Variations on basic VWAP are generally achieved by either omitting certain trades from the calculation or altering the time period from the standard full trading day.

The most frequently omitted trades from the VWAP calculation are arranged crosses, block, off-exchange or out-of-hours trades, on the basis that these were not available to the trader, and their inclusion would invalidate the objective of a fair representation of the traders opportunity set.

Where the trade to be executed is a substantial fraction of total trading, and therefore likely to significantly influence the VWAP price, 'Exclusive VWAP' (where the traders own executions are omitted) has the dual benefit of providing a more independent benchmark and encouraging price discipline that might otherwise be lost if the trade were large enough to effectively set VWAP if included.

Equally interesting is the recent rise in the use of interval VWAP, which in its purest form addresses the issue of unnecessarily spreading liquid trades out over a whole day (encouraged by the use of basic VWAP) by calculating VWAP only over the time period in which the orders are actually executed. This may still not represent the optimal period for execution, but it does remove the incentive to unnecessarily delay trading liquid orders solely to avoid the potential for significantly undershooting the benchmark.

Participation Rate Based Benchmarks

The author finds that none of the above benchmarks are particularly robust and useful measures for either gauging execution quality or discerning between market impact and market movement.

In order to see the issues clearly, we need to be analysing the problem from the correct perspective and framework. Most previous academic studies have focused on working from either a time or volume perspective. In our opinion, the clearest view comes when working in volume-time space, since it is the combination of volume and time which determines the potentiality for costs.⁷ Neither volume nor time individually captures the dimensionality of the market environment unless all orders are so small that they can be instantaneously executed — an extremely rare occurrence with transitions. The level of volume for any given security varies over time, both intra-day and over successive time periods.

⁷ "Volume-Time" is not an abstract cosmological concept. Rather, it is simply a reflection of the dynamic nature of markets. Almgren Thum, Hauptmann, and Li (2005) describe it in this manner: "The level of market activity is known to vary substantially and consistently between different periods of the trading day: this intraday variation affects both the volume profile and the variance of prices...Volume time, τ , ... represents the fraction of an average day's volume that has been executed up to clock time t ."

...the best way to yield insights into how a trade interacts with the market is in terms of volume-time, i.e., the time interval in which a certain percentage of a day's volume is executed, for which the key metric is the participation rate of the order.

However, the daily volume profile, i.e., the distribution of volume over a trading day, is typically remarkably stable. In such an environment, the best way to yield insights into how a trade interacts with the market is in terms of volume-time, i.e., the time interval in which a certain percentage of a day's volume is executed, for which the key metric is the participation rate of the order. It is easy to see why this would be the case when we think about how a trade interacts with the market as a whole in the price formation process. Each buy order creates upward pressure on the price of a security, and likewise, each sell order creates downward pressure. A market clearing equilibrium is found when the buy and sell orders are balanced. When a new order arrives, it disturbs this equilibrium, creating a net upward or downward force. The strength of this force depends on the relative size of the order to the volume being transacted in that time interval, which is the definition of the participation rate. And so the greater the participation rate, the greater the market impact. This perspective is gaining in popularity in academic studies, with a number of papers in the last couple of years looking at the optimal trading problem through the participation rate.⁸ It is also notable that in the practical field too, participation rate based strategies are being formalised into trading algorithms and benchmarks. Until fairly recently, the use of the participation rate was more by way of a constraint than a strategy. For example, an order might be designed to trade with a VWAP strategy but be no more than 20% of volume. In line with volume (ILV) orders were more common, but rarely stated a target percentage of volume.

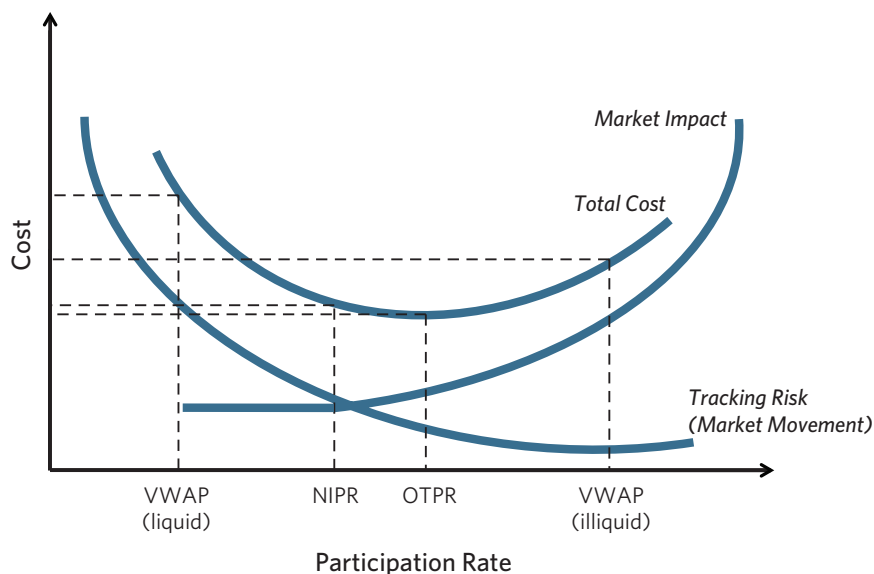
Exhibit 3 below will be familiar to many as the diagram often wheeled out to illustrate the trade-off between market impact and tracking risk and the identification of an optimal trading period. The x- axis is normally time, but here it is the participation rate, and so the market impact and market movement (tracking risk) functions are reversed so that tracking risk (adverse market movement) decreases as the participation rate increases and market impact increases as the participation rate increases. Note that market impact at low levels of participation rate is fairly constant: it does not increase materially as we move along the lower part of the x- axis, and then grows exponentially at higher rates. There are several reasons for this. Firstly, at very low levels of participation, orders are easily absorbed by the market, and the shock waves created are not easily discernable from the regular ebb and flow of the market. Only at much greater levels do we genuinely have impact, and as we increase the participation rate from here, orders then have incrementally larger impact and reveal more and clearer information. The second reason why impact is low and constant at low levels of participation is that prices are not continuous; they are actually discrete and change in steps equal to the tick size. At low rates, the impact is either nothing, or one tick size.

⁸ See, for example, Algrem et al (2005), Gomes & Waelbroeck (2010) and Fraenkle et al (2010).

...NIPR is the closest we have to a zero impact trade schedule and forms the basis of a truly passive trading benchmark that can also be reliably used to attribute between endogenous market impact and exogenous market movement.

A number of recent empirical studies have confirmed this functional form of the impact schedule,⁹ and also that at low participation rates there is little or no incremental change in the impact. Furthermore, some of these studies have shown that the functional form is uniform across the major markets with the main discrepancy being Japan which is distorted by unusually large tick sizes and the 'up tick'¹⁰ rule.

Exhibit 3 - Cost as a Function of the Rate of Trading



Source: BNY Mellon Beta Management, 2011.

Note: This illustration is based on the concept of varying participation rates on the market impact cost and tracking risk of a transition trade, and does not reflect actual trading.

Two critical levels of participation are identified in Exhibit 3 and are labelled as the negligible impact participation rate (NIPR) and the optimal trajectory participation rate (OTPR).

The NIPR marks the inflection point when the market impact function ceases to be flat, and so incremental changes in the rate will increase the impact. The OTPR marks the participation rate which minimises total cost. Both of these points have significant applications for the analysis of transition costs. The author believes that the NIPR is the closest we have to a zero impact trade schedule and forms the basis of a truly passive trading benchmark that can also be reliably used to attribute between endogenous market impact and exogenous market movement. It defines the participation rate (and by extension the timeframe) at which orders can be executed as quickly as

⁹ See, for example, Algreem et al (2005), Farmer et al (2010), Gatheral (2010) etc.

¹⁰ The up-tick rule is a mechanism applied in a number of markets that restricts the short selling of securities to only occur on an "up-tick", i.e. at a price higher than the last traded price, or at the current price if that was higher than the last executed price prior to that.

At its most basic level, the search for an optimal execution trajectory involves the trade-off of two types of implicit trading cost: market impact and market movement.

possible without incurring impact costs and thus with minimal adverse market movement risk. Trading at a faster rate is in effect a choice to incur market impact in preference to market movement risk. Trading at a slower rate is in effect choosing (unnecessarily) to take on market movement risk since orders will be left unfilled for a longer time, thereby increasing the chances that prices move against you.

One illuminating insight from this perspective is that a full day VWAP strategy is actually a de facto participation rate strategy, spreading orders over the day, in line with the volume profile. The difference to explicit participation rate strategies is that the rate is not chosen directly as a strategic choice in balancing market impact and tracking risk but indirectly and in a counter-intuitive fashion, from the liquidity of the order, as measured by the percentage of an average day's volume. So a liquid stock that could be completed quickly with no impact, for example one that represented 3% of the average daily traded volume (ADV), would have a participation rate of 3% whilst a less liquid security with higher potential impact costs, say 30% ADV, would have a much higher participation rate. In our view, per the illustration in Exhibit 3, only if (by accident or coincidence) the liquidity measure matches the OTPR would VWAP be an appropriate optimal strategy, and likewise, only if it matched the NIPR would VWAP be an appropriate benchmark from which to attribute implicit costs between market impact and market movement.

Determining the Optimal Execution Trajectory

To develop a methodology for determining the optimal execution trajectory, we adopt the standard convention of applying mean variance analysis. This is consistent with modern portfolio theory, provides a simple analytical framework within which to operate, and is likely to be sufficiently approximate to the actual distribution in the timescales applied so as to allow meaningful conclusions to be drawn.

At its most basic level, the search for an optimal execution trajectory involves the trade-off of two types of implicit trading cost: market impact and market movement. We can safely ignore the explicit costs of execution (commission, taxes and fees) which are normally *ad valorem* in nature rather than transaction based and so do not impact upon the analysis.

...only with trading period average benchmarks can execution performance and market impact be reasonably estimated and a fair and unbiased attribution of implicit costs be achieved when analysing the results of a transition assignment.

As depicted in Exhibit 3, both of the implicit costs involved are a function of the pace of trading. Market impact increases as the pace of trading increases whilst opportunity cost decreases as the pace of execution increases. Consequently, we have an inherent trade-off between the two and, importantly, a concave total cost curve that has a global minimum where the total cost is at its lowest somewhere between the two extremes. It is this sweet spot that we are seeking in deriving an optimal trajectory. In the case of a transition, Exhibit 3 shows the costs at a macro level and *does not* provide details at the micro (i.e. single order) level, including the interactions for each individual order, both between each other and with respect to the transition portfolio as a whole. At the micro level, and in the absence of information contained in the trades, once the optimal aggregate participation rate (OTPR) is determined, the most important consideration relates to each individual order's contribution to total tracking (market movement) risk. An optimal implementation strategy would recommend faster execution of those securities with the highest marginal contribution to tracking error (MCTE) between the transition and target portfolio and possibly slower execution of those with very low or negative MCTE so that realised tracking risk is minimised.

The NIPR Trading Benchmark

The aim of the NIPR benchmark is remarkably simple — to facilitate the accurate and fair attribution of total implicit costs between market movement and market impact by providing a truly passive benchmark that is relevant to the trading undertaken and avoids some of the shortcomings of the two most common trading benchmarks, namely arrival price and VWAP.

As detailed above, post-trade benchmarks are not useful for either measuring market impact or in attributing between market movement and market impact. Pre-trade benchmarks are also severely restricted in their ability in this regard. When the previous day's close is used as the trading benchmark, the differential between the average executed price and the benchmark is not market impact but rather a combination of delay, market impact and market movement. When the Open is used as the trading benchmark, any market impact incurred at the open is attributed to delay and the market movement component is incorporated into market impact thus overstating the true impact.

It is this author's contention that only with trading period average benchmarks can execution performance and market impact be reasonably estimated and a fair and unbiased attribution of implicit costs be achieved when analysing the results of a transition assignment. Even here, care needs to be taken to ensure that the benchmark being used is appropriate for each calculation. We believe that for execution performance analysis, a truly passive benchmark that also represents both the execution strategy and trading horizon is required. To calculate market impact and thus be able to attribute between the implicit costs, we require a benchmark that is not materially influenced by the

Most discussion around the measurement of market impact has focused on the impact which is realised during the period of trading. This in reality only captures one part of the actual impact that occurs and leaves another potentially very important component hidden.

executions and so is not contaminated with the market impact of the executions, i.e. a truly passive and independent benchmark. One way to do this is exclusive VWAP, but, in our view, a more robust benchmark would be the NIPR.

To calculate the NIPR benchmark, we first identify the negligible impact participation rate, π_N , and then calculate the participation weighted average price (PWAP) for the first X/π_N shares traded in the market that day (or after the order is received if later), where X is the total order size. For example if π_N is 10% and our order is to buy 100,000 shares of a stock, the NIPR benchmark would be the average price of the first million shares traded that day. Many transaction cost analysis (TCA) packages available (for example those provided by Instinet and ITG among others) already contain PWAP benchmarks where one can stipulate the participation rate. The main obstacle is in identifying the NIPR. Anecdotally, many brokers or traders will assert that participation rates of 20-25% would not have material impact; however, our experience would indicate that the NIPR actually lies in the 10% range for liquid (continuously traded) securities in most developed markets. One research project underway is to empirically validate this for the major global equity markets and identify any systematic variations due to, for example, capitalisation or index membership.

As mentioned earlier, if we are implementing orders that have no informational content, and thus no alpha, then, in the absence of strong momentum, we believe there is no justification for incurring market impact on the basis of capturing alpha. Rather, our main concern should be avoiding impact and controlling tracking error risk by first executing orders that have the largest marginal contribution to portfolio tracking error.

When using NIPR or a proxy for it, this paper seeks to demonstrate that we can validly attribute between observable market impact and market movement, but we are still missing one potentially important, but usually overlooked component of market impact: unobserved market impact.

Post-Trade Price Reversion as a Component of Market Impact

Most discussion around the measurement of market impact has focused on the impact which is realised during the period of trading. This in reality only captures one part of the actual impact that occurs and leaves another potentially very important component hidden. To see this, we need to look from a slightly different perspective and measure how the securities that we are trading react *once the order is completed* to reveal the hidden impact being exerted — in essence an economic opportunity cost. For example, consider the simple scenario of a trader with a buy order for 10 million XYZ Inc. shares, which open at \$5. The trader sits on the offer trading at \$5 until completed

Implementation shortfall measures the total cost incurred in a transition, but further useful information can be gleaned by decomposing its component parts using a rigorous approach detailed in this paper.

with an average price of \$5. Immediately after completing the order, the share price drops 6% to \$4.70 and stays there for some time afterwards. Whilst most transaction cost measures would show both no market impact and no market movement cost, it is quite clear that the trader was impacting the price by supporting it thus stopping its fall to the level that it eventually settled on once the trader's upward pressure had been released. This is latent market impact and ought to be taken into account when measuring transaction costs and trader performance. One way to do so, as alluded to above, is to measure the performance of the security immediately following completion of the order, which is sometimes referred to as "snap back" analysis. Much of the academic literature on the subject suggests that the temporary component of market impact (of which the author considers this to be) has dissipated within 30 minutes or so for most liquid securities, making this an appropriate period over which to measure the latent impact through price reversion. However, we still need to be careful to ensure that we do not capture exogenous market movement, and so it would be prudent to attribute to latent market impact only the price reversion on those securities whose actual realised participation rate was greater than the NIPR.

Conclusions

Implementation shortfall measures the total cost incurred in a transition, but further useful information can be gleaned by decomposing its component parts using a rigorous approach detailed in this paper. We believe:

- Market impact and realised market movement risk are unobservable directly but can be deduced;
- Trading performance (execution quality) and market impact are not the same thing;
- Using trading performance as a proxy for market impact can seriously underestimate market impact's size, depending on the trading benchmark used;
- Basic VWAP has serious flaws that call into question its usefulness as a trading benchmark. However, there are modifications that can alleviate some of the issues;
- In determining the optimal trading trajectory, the most useful dimension is neither time nor volume but volume-time; and
- Participation rate based benchmarks are superior metrics for analysing both trading performance (OTPR) and the attribution of implicit trading costs (NIPR).

Graham C. Cook, CFA, FRM

Graham is a Vice President and Transition Manager with BNY Mellon Beta Management. Prior to this he was a Director and Head of the Portfolio Implementation and Restructuring Group at Instinet, and Joint Head of Portfolio Solutions with Morgan Stanley in Australia. He has also held quantitative investment manager roles with Royal London Asset Management and First Quadrant in the United Kingdom. Graham graduated with a BA in Financial Economics from Coventry University and an MSc in Corporate and International Finance from the University of Durham, is a CFA charter holder, Chartered Financial Risk Manager (FRM) as well as holding the Investment Management Certificate (IIMR) and Retirement Provision Certificate (PMI). With 18 years of industry experience, Graham has published several articles on the subject of transition management and has presented at numerous industry events.

Appendix 1

Example of Attribution and Comparison
Using Different Benchmarks

Buy 10m Stock XYZ Shares 13th January 2011

Reference Prices

Previous close	PDC	1.7100	= Decision Price
Open	OPN	1.7130	

Trading Benchmarks

VWAP	VWAP	1.7122
NIPR	NIPR	1.7099
Interval VWAP	IVWAP	1.7106
Execution Price	Px	1.7127

Total Implicit Costs $(PDC - Px) / PDC$ = Delay + Market Impact + Market Movement

Attribution Components

Overnight Gap	$(PDC - OPN) / PDC$
Market Movement	$(OPN - BM) / PDC$
Market Impact:	$(BM - Px) / PDC$

NIPR Benchmark		Full Day VWAP Benchmark		Interval VWAP Benchmark	
Total Implicit Costs	-0.16%	Total Implicit Costs	-0.16%	Total Implicit Costs	-0.16%
Overnight Gap (Delay)	-0.18%	Overnight Gap (Delay)	-0.18%	Overnight Gap (Delay)	-0.18%
Trading Related Costs	0.02%	Trading Related Costs	0.02%	Trading Related Costs	0.02%
Market Impact	-0.16%	Market Impact	-0.03%	Market Impact	-0.13%
Market Movement	0.18%	Market Movement	0.04%	Market Movement	0.14%

Source: BNY Mellon Beta Management, 2011.

Note: This illustration uses the real value of a security and the real impact of attribution based on real benchmark data but the security was not actually traded. The NIPR benchmark used in this illustration assumes that the NIPR is 10%.

Disclosures

Mellon Transition Management operates as a branded entity contracting under The Bank of New York Mellon, a subsidiary of The Bank of New York Mellon Corporation. Certain brokerage services are provided by MBSC Securities Corporation, a BNY Mellon Company and FINRA member.

The statements and opinions expressed in this article are those of the author as of the date of this article, and do not necessarily represent the view of The Bank of New York Mellon Corporation, BNY Mellon, Mellon Transition Management or any of their respective affiliates. This article does not constitute investment advice, and should not be construed as an offer to sell or a solicitation to buy any security or make an offer where otherwise unlawful. The Bank of New York Mellon Corporation, BNY Mellon, Mellon Transition Management and its affiliates are not responsible for any subsequent investment advice given based on the information supplied.

Past performance is not a guide to future performance. The value of investments and the income from them is not guaranteed and can fall as well as rise due to stock market and currency movements. When you sell your investment you may get back less than you originally invested. No investment strategy or risk management technique can guarantee returns or eliminate risk in any market environment. No investment process is free of risk and there is no guarantee that the investment process described herein will be profitable. Investors may lose all of their investments.

This is not intended as investment advice but may be deemed a financial promotion under non-U.S. jurisdictions. The information provided is for use by professional investors only and not for onward distribution to, or to be relied upon by, retail investors. The information herein is subject to change, without notice, at the discretion of Mellon Transition Management and Mellon Transition Management does not undertake to revise or update this information in any way.

The information is not intended and should not be construed as legal or tax advice. Parties should independently investigate any investment strategy or manager, and should consult with qualified investment, legal, and tax professionals before making any investment.

Neither Mellon Transition Management nor any other party involved in or related to compiling, computing or creating the included data makes any express or implied warranties or representations with respect to such data (or the results to be obtained by the use thereof), and all such parties hereby expressly disclaim all warranties of originality, accuracy, completeness, merchantability or fitness for a particular purpose with respect to any of such data. The investment products and services mentioned here are not insured by the FDIC (or any other state or federal agency), are not guaranteed by any bank and may lose value.

Certain information included in this article reflects a hypothetical portfolio and accordingly, has certain inherent limitations. Unlike the results shown for an actual performance record, any simulated or hypothetical results do not represent actual trading. Also, because trades have not actually been executed, simulated results may have under- or over-compensated for the impact, if any, of certain factors, such as lack of liquidity. Simulated or hypothetical trading programs in general are also subject to the fact that they are designed with the benefit of hindsight. No representation is being made that any account will or is likely to achieve the profits or losses similar to those being shown.

References

- Almgren, R., C. Thum, E. Haptnann, and H. Li. "Equity Market Impact" Risk, (July 2005) pp57-62
- Bertsimas, D., and A.W. Lo. "Optimal Control of Execution Costs" *Journal of Financial Markets*, 1 (1988) pp1-50
- Berkowitz, S., D. Logue and E. Noser. "The Total Cost of Transactions on the NYSE," *Journal of Finance*, Vol.43, No.1 (1988) pp 97-112
- Clay, M., "The T Ratio – An Information Ratio for Transition Events" *Journal of Performance Measurement*, Vol.11, No.4 (Summer 2007) pp26-34
- Collie, R.C., "Transition Trading and the T Standard" *Journal of Trading*, Vol. 2009 No.2 (Transition Management Guide, Summer 2009) pp41-44
- Dickson, J. and R. Benson. "In Search of the Optimal Trading Frontier" *Transaction Performance*, Vol. 2002, No.1 (Spring 2002), pp47-53
- Domowitz, I., J. Glen and A. Madhavan. "Liquidity, Volatility and Equity Trading Costs Across Countries and Over Time, *ITG Working Paper*, 2001
- Farmer, J.D., A.Gerig, F. Lillo and H. Waelbroeck. "A New Approach to Understanding the Market Impact of Large Trading Orders" Unpublished working paper, preliminary draft.
- Fraenkle, J., S.T. Rachev and C. Scherrer. "Market Impact Measurement of a VWAP Trading Algorithm" Unpublished working paper, Karlsruhe Institute of Technology, 2010
- Freyre-Saunders, A., R. Guobuzaita and K. Byrne. "A Review of Trading Cost Models: Reducing Transaction Costs" *Journal of Investing*, Vol.13, No. 3 (Fall 2004) pp93-115
- Gatheral, J. "No dynamic arbitrage and market impact" *Quantitative Finance* Vol. 10, No7 (2010) pp749-759
- Gomes, C., and H Waelbroeck. "Effect of Trading Velocity and Limit Prices on Implementation shortfall". *Pipeline Financial Report*, September 2008
- Gomes, C., and H Waelbroeck. "Transaction Cost Analysis to Optimise Trading Strategies", *Journal of Trading*, Vol.5, No.4 (Fall 2010) pp29-38
- Keim, D. and A. Madhavan. "The Cost of Institutional Equity Trades: An Overview" *Financial Analysts Journal*, Vol .54, No. 4 (July 1998) pp50-69
- Kissell, R. "The Expanded Implementation shortfall: Understanding Transaction Cost Components" *Journal of Trading*, Vol .1, No.3 (Summer 2006) pp6-16
- Kissell, R. "Transaction Cost Analysis: A Practical Framework to Measure Costs and Evaluate Performance" *Journal of Trading*, Vol.3, No.2 (Spring 2008) pp29-37
- Kritzman, M., S. Myrgren and S. Page. "Implementation shortfall, From Concept to Practice", *Journal of Portfolio Management*, Vol .33, No.1 (Fall 2006) pp25-30
- Madhavan, A. "VWAP Strategies" *Transaction Performance*, Vol .2002, No.1 (Spring 2002) pp32-39
- Perold, A. "Implementation shortfall: Paper Versus Reality" *Journal of Portfolio Management*, Vol.14, No.3 (Spring 1988) pp4-9
- Rakhalin, D., and G. Sofianos, "Choosing Benchmarks vs. Choosing strategies: Part 2- Execution Strategies: VWAP or Shortfall" *Journal of Trading*, Vol.1, No.1 (Spring 2006) pp26-32
- Schwartz, R.A., and B. Steil. "Controlling Institutional Trading Costs" *Journal of Portfolio Management*, Vol.28, No.3 (Spring 2002) pp39-49
- Sofianos, G. "Choosing Benchmarks vs. Choosing Strategies: Part 1- Execution Benchmarks: VWAP or Pre Trade Prices" *Journal of Trading*, Vol.1, No.1 (Spring 2006) pp71-74
- Wagner, W.H., and M. Edwards. "Best Execution" *Financial Analysts Journal*, Vol.49, No.1 (1993) pp65-71
- "T Standard Version 2.0" Russell Investments, December 2008

BNY Mellon Asset Management is the umbrella organization for BNY Mellon's affiliated investment management firms and global distribution companies. BNY Mellon is the corporate brand of The Bank of New York Mellon Corporation. • The statements and opinions expressed in this article are those of the authors as of the date of the article, are subject to change as economic and market conditions dictate, and do not necessarily represent the views of BNY Mellon, BNY Mellon Asset Management International or any of their respective affiliates. This article is of general nature, does not constitute investment, legal or tax advice, is not predictive of future performance, and should not be construed as an offer to sell or a solicitation to buy any security or make an offer where otherwise unlawful. The information has been provided without taking into account the investment objective, financial situation or needs of any particular person. BNY Mellon Asset Management International Limited and its affiliates are not responsible for any subsequent investment advice given based on the information supplied.

Past performance is not a guide to future performance. The value of investments and the income from them is not guaranteed and can fall as well as rise due to stock market and currency movements. When you sell your investment you may get back less than you originally invested. • While the information in this document is not intended to be investment advice, it may be deemed a financial promotion in non-U.S. jurisdictions. Accordingly, where this document is used or distributed in any non-U.S. jurisdiction, the information provided is for use by professional investors only and not for onward distribution to, or to be relied upon by, retail investors. • Products or services described in this document are provided by BNY Mellon, its subsidiaries, affiliates or related companies and may be provided in various countries by one or more of these companies where authorized and regulated as required within each jurisdiction. However, this material is not intended, and should not be construed, as an offer or solicitation of services or products or an endorsement thereof in any jurisdiction or in any circumstance that is otherwise unlawful or unauthorized. **The investment products and services mentioned here are not insured by the FDIC (or any other state or federal agency), are not deposits of or guaranteed by any bank, and may lose value.**

• This document should not be published in hard copy, electronic form, via the web or in any other medium accessible to the public, unless authorized by BNY Mellon Asset Management International Limited.

In Australia, this document is issued by BNY Mellon Asset Management Australia Limited (ABN 56 102 482 815, AFS License No. 227865) located at Level 6, 7 Macquarie Place, Sydney, NSW 2000. Authorized and regulated by the Australian Securities & Investments Commission. • In Brazil, this document is issued by BNY Mellon Serviços Financeiros DTVM S.A., Av. Presidente Wilson, 231, 11th floor, Rio de Janeiro, RJ, Brazil, CEP 20030-905. BNY Mellon Serviços Financeiros DTVM S.A. is a Financial Institution, duly authorized by the Brazilian Central Bank to provide securities distribution and by the Brazilian Securities and Exchange Commission (CVM) to provide securities portfolio managing services under Declaratory Act No. 4.620, issued on December 19, 1997. • Investment vehicles may be offered and sold in Canada through BNY Mellon Asset Management Canada Ltd., a Portfolio Manager, Exempt Market Dealer and Investment Fund Manager. • In Dubai, United Arab Emirates, this document is issued by the Dubai branch of The Bank of New York Mellon, which is regulated by the Dubai Financial Services Authority. • In Germany, this document is issued by WestLB Mellon Asset Management Kapitalanlagegesellschaft mbH, which is regulated by the Bundesanstalt für Finanzdienstleistungsaufsicht. WestLB Mellon Asset Management Holdings Limited is a 50:50 joint venture between BNY Mellon and WestLB AG. WestLB Mellon Asset Management Kapitalanlagegesellschaft mbH is a wholly owned subsidiary of this joint venture. • If this document is used or distributed in Hong Kong, it is issued by BNY Mellon Asset Management Hong Kong Limited, whose business address is Level 14, Three Pacific Place, 1 Queen's Road East, Hong Kong. BNY Mellon Asset Management Hong Kong Limited is regulated by the Hong Kong Securities and Futures Commission for Type 1 (dealing in securities), Type 4 (advising on securities) and Type 9 (asset management) regulated activities, and its registered office is at 6th floor, Alexandra House, 18 Chater Road, Central, Hong Kong. • In Japan, this document is issued by BNY Mellon Asset Management Japan Limited, Meiji Seimei Kan 6F, 2-1-1 Marunouchi Chiyoda-ku, Tokyo 100-0005, Japan. BNY Mellon Asset Management Japan Limited is a Financial Instruments Business Operator with license no 406 (Kinsho) at the Commissioner of Kanto Local Finance Bureau and is a Member of the Investment Trusts Association, Japan and Japan Securities Investment Advisers Association. • In Korea, this document is issued by BNY Mellon AM Korea Limited for presentation to professional investors. BNY Mellon AM Korea Limited, 21/F Seoul Finance Center, 84 Taepyungro 1-ga, Jung-gu, Seoul, Korea. Regulated by the Financial Supervisory Service. • In Singapore, this document is issued by The Bank of New York Mellon, Singapore Branch for presentation to professional investors. The Bank of New York Mellon, Singapore Branch, One Temasek Avenue, #02-01 Millenia Tower, Singapore 039192. Regulated by the Monetary Authority of Singapore. • This document is issued in the UK and in mainland Europe (excluding Germany), by BNY Mellon Asset Management International Limited, 160 Queen Victoria Street, London EC4V 4LA. Registered in England No. 1118580. Authorized and regulated by the Financial Services Authority. • This document is issued in the United States by BNY Mellon Asset Management.

BNY Mellon holds over 90% of the parent holding company of The Alcentra Group. The Group refers to these affiliated companies: Alcentra, Ltd and Alcentra NY, LLC. Only Alcentra NY, LLC offers services in the U.S. • Ankura, Insight Investment and WestLB Mellon Asset Management do not offer services in the U.S. This presentation does not constitute an offer to sell, or a solicitation of an offer to purchase, any of the firms' services or funds to any U.S. investor, or where otherwise unlawful. • BNY Mellon holds a 20% interest in Siguler Guff & Company, LP and certain related entities (including Siguler Guff Advisers LLC). • BNY Mellon Beta Management is a division of The Bank of New York Mellon, a wholly-owned banking subsidiary of BNY Mellon. • BNY Mellon Cash Investment Strategies is a division of The Dreyfus Corporation. • Hamon's services are offered in the U.S. by Hamon U.S. Investment Advisors Limited. BNY Mellon holds a 19.9% interest in Hamon Investment Group Pte Limited, which is the parent of Hamon U.S. Investment Advisors Limited. • The Newton Group refers to the following group of companies: Newton Investment Management Limited, Newton Capital Management Limited, Newton International Investment Management Limited, Newton Capital Management LLC, and Newton Fund Managers (CI) Limited. Except for Newton Capital Management LLC and Newton Capital Management Limited, none of the other Newton companies offers services in the U.S. • BNY Mellon Asset Management International Limited and any other BNY Mellon entity mentioned above are all ultimately owned by BNY Mellon.



BNY MELLON ASSET MANAGEMENT

The Alcentra Group
Ankura Capital Pty Limited
Blackfriars Asset Management Limited
BNY Mellon ARX
BNY Mellon Beta Management
BNY Mellon Cash Investment Strategies
The Boston Company Asset Management, LLC
The Dreyfus Corporation
EACM Advisors LLC
Hamon Investment Group
Insight Investment
Mellon Capital Management Corporation
The Newton Group
Pareto Investment Management Limited
Siguler Guff & Company LP
Standish Mellon Asset Management Company LLC
Urdang Capital Management, Inc.
Urdang Securities Management, Inc.
Walter Scott & Partners Limited
WestLB Mellon Asset Management