MARKET INSIGHTS

The Point of No Returns:

Trading Halts and Other Market Structure Considerations in China A-Shares Research

September 2021

D E Shaw & Co

Introduction

Dongxu Optoelectronic Technology Co. Ltd. is an electronic component manufacturer headquartered in Beijing, China. Its stock has been listed since 1996, trades on the Shenzhen stock exchange, and, from June 2014 to June 2020, was a constituent of the CSI 300® Index.

Dongxu typifies issuers in the China A-shares market in at least one important way: over time, its stock has been subject to numerous trading halts. As reflected by the frequency and duration of the periods captured by the orange bars in *Figure 1*, trading in the stock has been halted, sometimes for extended periods, even as the broader market remained active.[1]

Investors in equity markets are accustomed to receiving daily security prices, but, as Dongxu illustrates, such data do not always exist for issuers in the China A-shares market. Because return data are fundamental when deriving quantities such as volatility, beta, and factor loadings, active decisions regarding how to treat trading halts for purposes of computing stock returns can affect key downstream elements of the investment process as well as portfolio outcomes.

As we'll show, given the unusual prevalence of trading halts in Chinese equity markets, those active research decisions are particularly important when conducting quantitative research on China A-shares.

Figure 1: Dongxu Cumulative Return
(2010–2020)

Sources: Bloomberg (Dongxu pricing data); the D. E. Shaw group. Applicable data are used with permission of Bloomberg.

[1] We chose Dongxu (or "Tungshu"—ticker: 000413:CH) as an example because we believe its lengthy trading experience is broadly representative of the frequency and magnitude of the key issues surrounding trading halts and pricing in the China A-shares market.

Understanding the Structure of the China A-Shares Market

Historically, Chinese stock exchange authorities—as well as company CEOs—have had the right to suspend the trading of individual stocks for a variety of reasons. These rules continue to evolve and are generally intended to limit potential threats to orderly market function, such as insider trading and panic-driven trading activity.

As *Figure 2* shows, individual issuers of China A-shares have been subject to trading halts at least an order of magnitude more frequently than constituents of other global equity markets. For more than half of these security-level halts in China, the restriction remained in place for longer than one day.

Although their prevalence has abated over the past few years as regulators introduced changes intended to reduce the frequency and duration of interruptions, trading halts remain a salient feature in the historical track records of China A-shares. Additionally, halts are often associated with events or information material to the price of a stock. When computing security- and market-level returns, it is therefore important to identify halts cleanly in the historical data.

From an allocator's perspective, interruptions in trading and pricing data might not appear to impact the overall return and risk profile of a long-term, strategic allocation. But over shorter horizons, they are material to active portfolio managers.

Trading halts are just one of a number of structural features of Chinese equity markets that active managers need to address. Others include limit-up/limit-down rules, which are frequently binding and often interact directly with halts,[2] as well as the effects of reverse mergers, state-owned enterprises, and limited free float.

When discussing these aspects of Chinese equity market structure, market participants typically do so in terms of *implementing and executing* a given investment strategy. What many overlook is the extent to which those structural elements influence the very foundational aspects of developing and designing an investment approach.

Figure 2: Stock Halts as a Percentage of Total Market
(2013–2020)

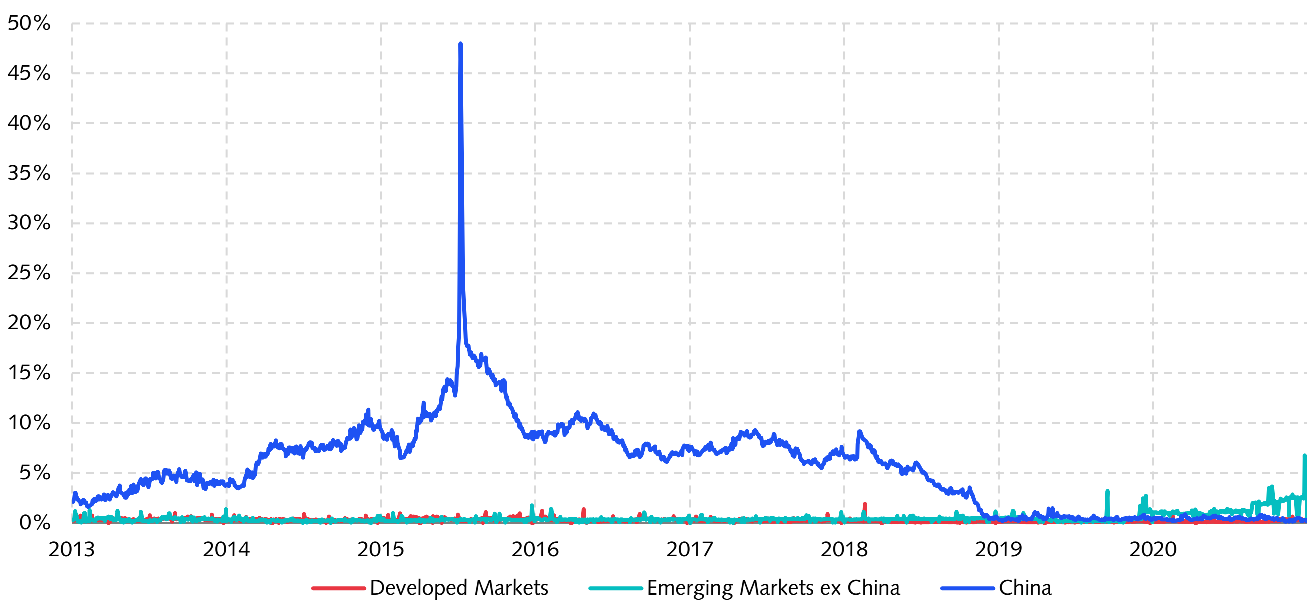

The above graph reflects the dollar average daily volume ("DADV") of stock halts as a percentage of the applicable market. The figures reflected are calculated as the expected DADV of halted stocks divided by the DADV of the relevant market. The expected DADV of halted stocks is calculated by propagating the last realized DADV adjusted by a cumulative multiplier that is proportional to the market-wide DADV.

The data reflected for "Emerging Markets ex China" represent countries that currently constitute approximately 85% of the weight of applicable markets.

Source: the D. E. Shaw group.

[2] Because of relatively tight daily up and down limits on price changes—typically +/-10%—it often takes a number of trading days for a previously halted security to reach its new equilibrium price.

Computing Returns Is Elemental, but Not Always Elementary

In a liquid and efficient market, it is common to assume that all securities are priced each trading day, and that a security's price reflects the vast majority of known information. Those assumptions do not hold for securities whose trading activity is suspended or whose price movement is limited—factors that complicate daily return computation.

To illustrate, we consider a stylized example in which the trading universe consists of only two equally weighted stocks, Stock A and Stock B, each of which moves in concert with the market.

Assume that the two stocks are at price parity before entering a six-day trading period that begins with an exogenous event—say, a natural disaster or geopolitical event—that causes a 25% decrease in the fundamental price of the market.[3] Additionally, assume that for reasons unrelated to this event, trading in Stock A is suspended for the first three days of the period, and a 10% daily price change limit applies to both stocks in the universe. *Table 1* summarizes the hypothetical daily returns of Stocks A and B over the six days following the exogenous event. (In this stylized example, we ignore compounding, idiosyncratic returns, and other potentially confounding factors.)

Because of the trading halt, Stock A's price doesn't mark at all for the first three days after the event. After the halt is lifted, its price decreases by 10%, 10%, and 5% on consecutive days—catching up to the market's 25% decline while still subject to daily limits of +/-10%.

By way of contrast, Stock B, which is not subject to a trading halt, decreases 10%, 10%, and 5% on consecutive days immediately following the event and then doesn't move on subsequent days, as no new information arrives that would affect the market.

In aggregate over the six-day period, both stocks fall by 25%, matching the overall decline in the market. Despite this straightforward appearance, however, an active manager would need to make some key methodological decisions. First among them, how should market returns be computed for each day? To account for the trading halt of Stock A, we consider below four potential approaches that market participants might adopt, and which result in the market returns shown in *Table 2*:

1) **Zero Return:** Often adopted by index providers (*e.g.*, MSCI Inc.), this approach assumes a zero return for Stock A on halted days. The full-period market return reflects the true loss due to the exogenous event, but the effect of assuming a zero return for Stock A on days 1–3 is that the market's decline on each of those days, as well the volatility computed for the market, are artificially low.

Table 1

	Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Total Return	Volatility (all days)	Volatility (only traded days)
Stock A	halt	halt	halt	-10%	-10%	-5%	-25%	4.9%	2.9%
Stock B	-10%	-10%	-5%	0%	0%	0%	-25%	4.9%	4.9%

Table 2

Market Return Approach	Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Total Return	Volatility (all days)	Volatility (only traded days)
1) Zero Return	-5.0%	-5.0%	-2.5%	-5.0%	-5.0%	-2.5%	-25.0%	1.3%	1.3%
2) Perfectly Efficient Market	-25.0%	0.0%	0.0%	0.0%	0.0%	0.0%	-25.0%	10.2%	10.2%
3) Adjust Universe	-10.0%	-10.0%	-5.0%	-5.0%	-5.0%	-2.5%	-37.5%	3.1%	3.1%
4) Exclude Halted Days	n/a	n/a	n/a	-5.0%	-5.0%	-2.5%	-12.5%	2.5%	1.4%

[3] Although our example is purely hypothetical, during the summer of 2015, trading was halted in approximately 50% of the China A-shares universe following a market drawdown of approximately 43%. Such high-magnitude events may be rare, but they do occur.

2) **Perfectly Efficient Market:** This approach attributes the entire 25% decline resulting from the exogenous event to the day news of the event became known. This assumes near omniscience—noise-free estimates of each stock's market beta and precise estimates of changes in market fundamentals—as well as the assumption that trading halts or limit-up/limit-down rules have no effect.
3) **Adjust Universe:** This approach dynamically excludes suspended stocks from the universe and estimates market returns accordingly. Although conceptually simple and perhaps accurate for each isolated day, the derived full-period market returns do not line up with the cumulative weighted returns of the underlying securities or, in most cases, those of the index provider.
4) **Exclude Halted Days:** In this approach, for days when a stock is suspended, one either excludes those days altogether or relies on an index provider to define relevant market returns. In our stylized example, in which an index provider is not involved, days 1–3 would be treated as missing observations, potentially ignoring useful information.

We explore below how the choice of methodology—often assumed as a given in the research process—affects risk estimation and return forecasting.

Risk Estimation is Subject to Estimation Risk

Our stylized example hints at the challenges that arise when calculating security-specific risk measures, especially as we apply more realistic assumptions to the factors that influence stock prices. Managers must account for estimation risk; that is, the reality that inputs to calculating a given measure—like volatility, beta, or factor loading—are themselves subject to uncertainty. As shown in *Table 1* and *Table 2*, estimates of stock- and market-level volatility can vary with the treatment of trading halts in the data.

Observed betas are also dependent on how we address trading halts in computing stock- and market-level return series. Taking our example further, we see that the four possible approaches to calculating market returns outlined above, plus the variable of including or excluding halted days when calculating individual stock returns, result in eight different estimates of Stock A's market beta (ranging from -0.93 to 2.0).

Returning to Dongxu Optoelectronic, *Figure 3* illustrates how estimates of risk measures can be sensitive to choices in return calculation methodology at the stock level alone, using the CSI 300® Index as a proxy for market return.

The estimates of volatility and beta that follow rely on three distinct approaches to assembling stock-level return series. *Table 2* above refers to the first two of these approaches as "Zero Return" and "Exclude Halted Days,"[4] respectively. The third is a proprietary "Adjusted/Efficient" approach, which seeks to move in the direction of, if not to approximate, the approach referred to above as "Perfectly Efficient Market."[5]

The ultimate impact of applying divergent risk estimates to a broader portfolio depends on the magnitude of resulting risk estimate variation across securities and through time. To get a sense of this relationship, we can compare the "Zero Return" and "Adjusted/Efficient" approaches to calculating security-level betas for constituents of the CSI 300® Index to the index itself. We find that in roughly 8.8% of daily, stock-level observations between 2005 and 2020, these approaches produce more than a 0.02 difference in resulting beta estimates.

Even small deviations in individual securities' beta estimates can compound quickly when seeking to manage portfolio-level risk precisely, especially during periods of market stress. Furthermore, this form of estimation risk applies not only to beta, but also to the estimation of a security's loadings to sector, industry group, and risk and style factors.

[4] Our application of the "Exclude Halted Days" approach entailed excluding (1) halted days, (2) any subsequent days with a 10% or larger price change (either positive or negative), and (3) one additional day.

[5] For reference, the third approach in *Table 2*, "Adjust Universe," does not apply to *Figure 3* because it is used to define the *market return*. By contrast, the example in *Figure 3* is focused on alternatives for calculating *stock-level measures relative to the market*, where we have chosen, for simplification purposes, to use the CSI 300® Index as a proxy for the market return.

Figure 3: Calculated Risk Measures for Dongxu

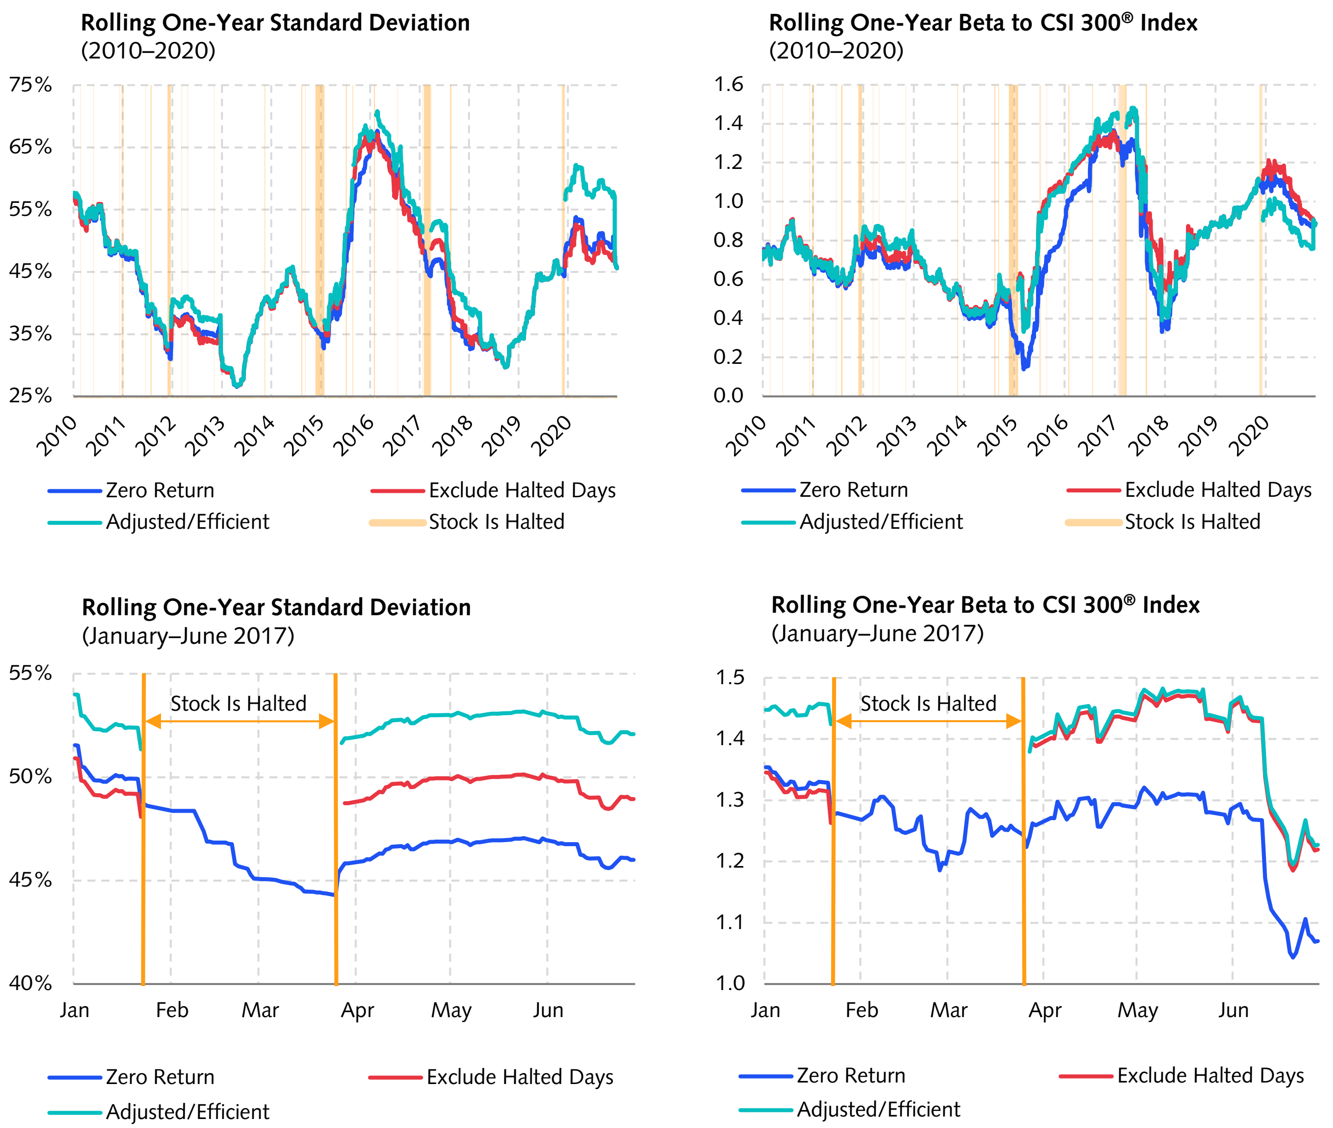

Source: the D. E. Shaw group. The charts above reflect application of the "Zero Return," "Exclude Halted Days," and "Adjusted/Efficient" approaches described above to the computation of the standard deviation or beta figures indicated.

Past Returns Matter When Forecasting Future Returns

The effects of trading restrictions on return computation also complicate the process for forecasting returns.

Consider an investment approach that deploys four equity style factors: low volatility, value, momentum, and size. Construction of these factors requires time series data of share prices and returns, and as discussed above, such data depend on active choices regarding how to handle trading halts.

For a single style factor and a single stock, using these different approaches to assemble time series factor loadings does, in fact, yield significantly different results. *Figure 4* shows estimates of Dongxu's loading to the low volatility factor using each of the "Zero Return," "Exclude Halted Days," and "Adjusted/Efficient" approaches. Especially during periods of higher market volatility, such as 2015–2016, we see meaningful differences in loadings across the approaches.

Figure 4: Dongxu Loading to Low Volatility Factor by Approach
(2010–2020)

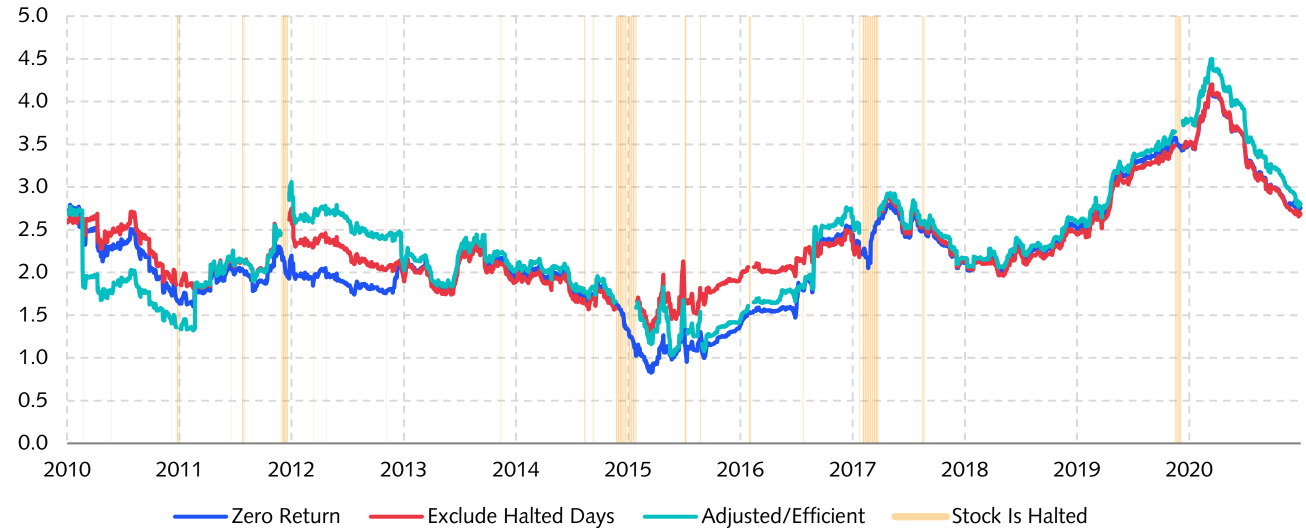

Source: the D. E. Shaw group. The chart above reflects application of the "Zero Return," "Exclude Halted Days," and "Adjusted/Efficient" approaches described above to the computation of factor loadings.

Figure 5: Long-Short Portfolio Cumulative Returns by Approach
(2010–2020)

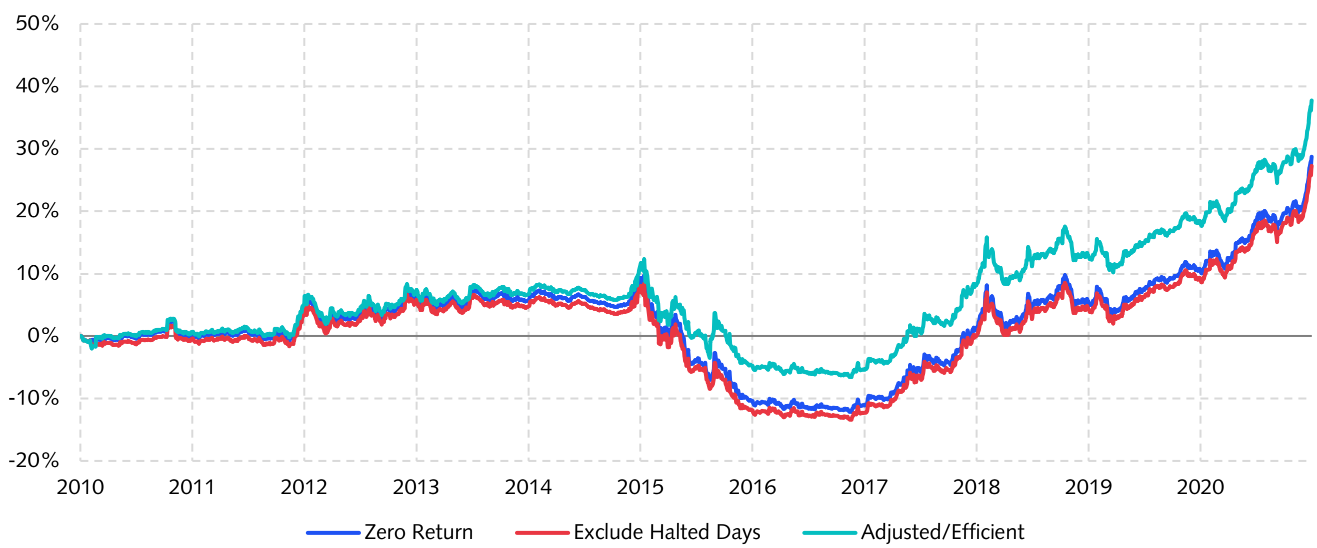

The absolute results of a long-short portfolio are qualitatively similar to excess results of a long-only implementation. We have chosen long-short to illustrate these effects in order to avoid complications introduced by the benchmark in a long-only approach. The hypothetical portfolio reflected here is intended for illustrative purposes and is simplified by, among other things, ignoring factors such as availability and borrow costs.

Source: the D. E. Shaw group.

When combining all four factors on an equal-weighted basis to construct a simplified long-short portfolio, *Figure 5* shows that significant performance differences accumulate by the end of the period. *Figure 5* reflects simulated annualized returns from 2010 through 2020 that range from +2.2% to +2.9% and Sharpe ratios that range from 0.55 to 0.70.

To summarize, active research decisions that feed into security and market return calculations are not trivial. In fact, different decisions can yield materially different outcomes.

Conclusion

For many investment applications, it is easy to overlook the potential challenges of assembling time series returns. But for active managers that deploy a quantitative approach and seek to deliver consistent risk-adjusted performance, research depends on high-quality data measured over short horizons, as this enables better calibration of risk estimates and return forecasts.

In the China A-shares market, data idiosyncrasies such as those for Dongxu are widespread across securities and over time, requiring caution, a deep understanding of market structure, and—as with any application of quantitative techniques—relentless discipline in assembling clean data to support both research and implementation. As we have seen in a number of respects, active decisions made when calculating and applying return series data for China A-shares can have a meaningful impact on portfolio risk and results.

D E Shaw & Co