
Evaluating Capacity and Liquidity for Equity Strategies

How to tell when a fund has become too big.

Morningstar Manager Research

6 November 2018

Contents

- 1 Executive Summary
- 3 Fund-Level Capacity
- 13 Market-Level Capacity
- 17 Liquidity Analysis
- 25 Conclusion

Tom Whitelaw

Director, Equity Ratings, Global Manager Research

+1 312-244-747

tom.whitelaw@morningstar.com

Important Disclosure

The conduct of Morningstar's analysts is governed by Code of Ethics/Code of Conduct Policy, Personal Security Trading Policy (or an equivalent of), and Investment Research Policy. For information regarding conflicts of interest, please visit:

<http://global.morningstar.com/equitydisclosures>

Executive Summary

Identifying a strategy's capacity is a key aspect of evaluating its prospects. It's one thing to identify a strategy that can do well using its current process and manager, but knowing the point at which the scale of its assets under management may erode alpha generation, or force alterations to a successful process, adds complexity to the analysis. Evaluations of capacity differ across asset classes. This paper will focus solely on equity strategies. As part of this process, we will also delve into liquidity analysis to better understand some of the threats that, while not currently evident, may be baked into the strategy if a liquidity event were to occur.

Red Flags: Data Points That Can Help You Spot Capacity Issues

Morningstar collects a vast array of data for funds, usually going back to the inception dates of individual share classes. This historical data trove helps our analysts assess if a fund is taking in too much money by allowing us to ascertain even very subtle changes in an investment approach through time. The data points we consider most relevant to capacity can be broadly split into three camps, each of which we'll cover in more depth in the following sections:

- ▶ Fund-Level Capacity
- ▶ Market-Level Capacity
- ▶ Liquidity Analysis


No one data point analyzed in isolation tells us very much about capacity. Our teams instead approach the issue through multiple vectors to build a mosaic of information. The same is true of looking at just one point in time. We therefore focus on time series data. We want to know what has changed and, more importantly, when it changed and in what environment. Our analysts are looking for patterns in the data, specifically whether changes to key data points occur around historical peaks for asset growth or when inflows have been especially rapid.

Changes to the way a strategy is run may not result in a downgrade to a fund. It will however lead to a reassessment that accounts for the current and expected future characteristics of the fund. For example, if a successful small-cap manager has had to move into mid-caps as the fund's size has grown, we would assess that strategy as a mid-cap offering. While the team's skills and expertise as small-cap investors would form part of our analysis, we would want to understand if it has the requisite skill set to perform in a different area of the market.

We also regularly hear from managers that they have adapted their styles because of current market dynamics rather than capacity issues. While this may be fair, we want to be sure that the strategy in question can still profit in its former hunting ground should those dynamics revert. That may not be possible if it's significantly bigger than it once was. Here, for example, we can use liquidity analysis of historical portfolios to see if this is the case.

To help highlight how we apply this analysis in the real world, we have used a single fund, Virtus KAR Small-Cap Growth PXSGX, for each of the following exhibits. As can be seen in Exhibit 1, this fund has seen dramatic increases in AUM in very short order, making it an interesting case study. Although the fund is a U.S.-domiciled offering, the principles outlined here are applicable to assessing the capacity of equity funds irrespective of domicile.

Exhibit 1 AUM Growth at Virtus KAR Small-Cap Growth



Source: Morningstar Direct.

Fund-Level Capacity

Fund-level data give us a holistic view of the portfolio. We are not drilling down on individual positions at this stage, instead we are determining how a strategy is run from the shape of the current and historical portfolios. As the name implies, fund-level data points are distinct to the specific strategy in question. They are universal, and we look at the same data points across all the global markets, sectors, and subsectors in which we cover funds.

There's no magic number for the amount of assets that a specific style of investing can handle—it's dependent on the individuals running it, the process deployed, and the point of the economic cycle. We therefore track AUM growth alongside the data points below to see if they correlate with asset growth. This helps us to identify funds that are growing beyond their means before it becomes an obvious problem, because by then, the relative performance damage may have already occurred.

Net Inflows

Net inflow data deals with the amount of new money moving in to a strategy. While total AUM gives us the overall fund size and growth in size, by using net inflows, we exclude the effects of market growth, allowing us to see how much new money the investment team is having to put to work (or in the case of outflows, the extent of selling that has been required).

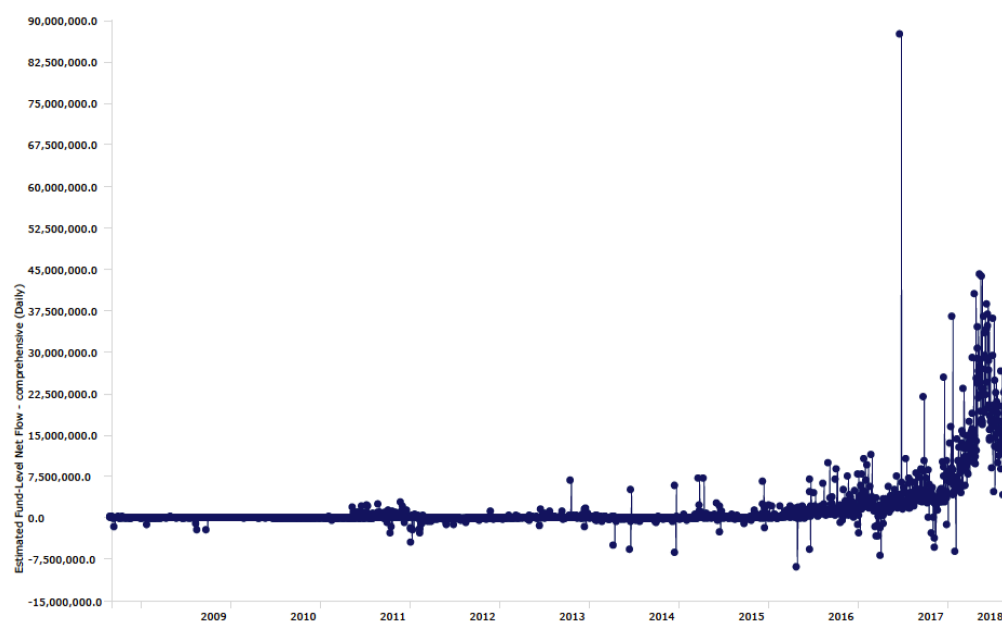
Organic Growth Rate (Speed of Growth)

Net flows also let us identify another important issue: speed of growth. It's much harder to put large sudden inflows of capital to work than it would be if those flows were spaced out over a longer period. Funds that show big monthly inflows may well have performance diluted in a rising market until that money can be put to work. Managers may opt to hold more cash, tempering performance in an up market, or purchase larger-cap, more-liquid securities than they might otherwise to keep cash levels in check. If the funds normally hold smaller-cap or less-liquid securities, their impact can get diluted by the latter tactic (the “liquidity tools” we discuss later are used to identify if this is taking place).

Virtus KAR Small-Cap Growth is a prime example of a fund that received rapid inflows. Exhibit 1 shows that from the beginning of 2017 to September 2018 the fund grew from \$500 million to over \$5 billion, while Exhibit 2 confirms that much of this growth came from inflows over a very short period. In this specific case, the growth in AUM depicted in Exhibit 1 is obviously of much greater magnitude than market appreciation in the period, so it is not hard to see the fund must have had inflows. Exhibit 2

confirms this and would be of greater help if the growth arc of a fund were less clearly in excess of market appreciation.

Exhibit 2 Monthly Inflows at Virtus KAR Small-Cap Growth



Source: Morningstar Direct.

Speed of growth is important because even if a fund's AUM remains modest, rapid inflows that are large relative to the fund's scale could water down returns. Short of closing the fund completely to further investment (a "hard close" as opposed to a "soft close" in which the fund is only closed to new accounts), there aren't many guaranteed solutions to cope with such inflows. Hard closes can have unintended deleterious consequences, as they immediately throw the fund into net redemptions (since the only flows are outbound), creating a different sort of liquidity challenge for the manager. Nevertheless, when we perceive a fund has rapid inflows, we prefer to see the fund company take action whether it be via a temporary hard- or soft-close, ceasing marketing, or turning away large individual mandates.

Rapid inflows can also create risk of a reversal. When money comes in quickly to a fund, it tends to be chasing performance (usually of the short-term variety, unfortunately) or trend-driven. That also means that when the performance reverses or the trend ends, that money can also quickly move out. This higher redemption risk leads to a higher chance that the fund will need to sell assets too quickly to get good prices for them.

Firm Assets Under Management

Firm assets under management can be more difficult to quantify as this measure requires that similar strategies run by a firm to be linked together. In the first instance, we do this by assigning them strategy IDs, which allow us to easily identify the investment vehicles that are using the same strategy at a firm, and thus how much money is run in the same style. There can also be a high degree of overlap between different strategies that a firm runs. For example, if a group has a successful emerging-markets fund and equally successful Asian and dedicated China funds, we would want to know if asset growth in the China fund was impacting the way the Asian strategy was being run, and if that in turn resulted in changes to the opportunities available to the broader emerging-markets fund. So, for some of the other data points we look at we would expand to the fund's holdings to include overlap with other strategies in their stable.

In the case of the Virtus KAR fund, the growth team also runs a core fund (with over \$6.865 billion in AUM at the end of September 2018) that tracks the Russell 2000 Index. As the portfolios share around 40% in common holdings, it's worth assessing their combined ownership of relatively illiquid names. Exhibit 3 shows all the stocks where the combined ownership by the two funds is more than 3% of the total shares outstanding in the companies in question.

Exhibit 3 Virtus KAR Small-Cap Growth and Virtus KAR Small-Cap Core Combined Weightings

Name	Ticker	Shares	Shares Outstanding*	% of Firm Owned*
Abaxis Inc	ABAX	4,310,188.49	22,870,000.00	18.85
RLI Corp	RLI	6,634,347.53	44,340,000.00	14.96
Moelis & Co Class A	MC	6,567,981.37	46,000,000.00	14.28
Autohome Inc ADR	ATHM	10,130,143.32	117,140,000.00	8.65
Pricesmart Inc	PSMT	2,384,868.67	30,460,000.00	7.83
Aspen Technology Inc	AZPN	4,824,098.81	71,010,000.00	6.79
FactSet Research Systems Inc	FDS	1,564,247.57	36,380,000.00	4.30
Old Dominion Freight Lines Inc	ODFL	3,340,485.52	82,000,000.00	4.07
Copart Inc	CPRT	9,393,529.08	233,920,000.00	4.02
MarketAxess Holdings Inc	MKTX	1,388,732.83	37,570,000.00	3.70

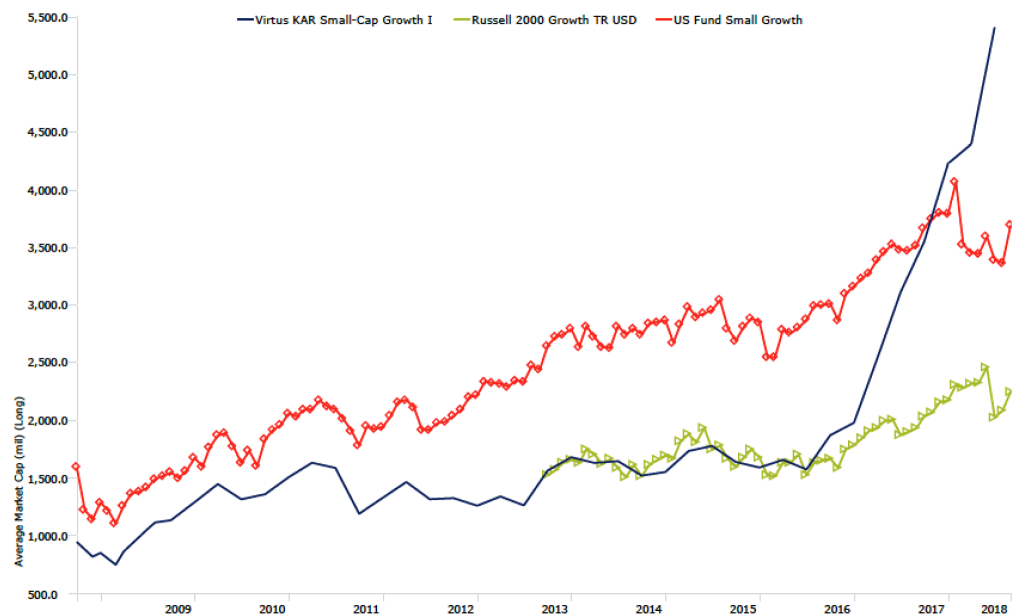
Source: Morningstar Direct. Data as of September 2018. These numbers are estimates based on the latest portfolios provided by Virtus KAR (June 30, 2018).

Average Market Cap

Changes in the average market cap of a portfolio's holdings over time can help us identify if a strategy has been stepping up the market-cap ladder. To avoid misidentifying funds whose average company size has grown simply because stock prices (and thus market cap) in that market segment have appreciated, it's important to consider this relatively. Therefore, we track any changes alongside changes in the average market caps of the stocks contained in the most appropriate index.

While we acknowledge that there are plenty of examples of managers moving up the market-cap ladder without any noticeable detriment to performance, those managers are still playing with a reduced opportunity set. For example, the past few years have been a great time to be in large caps. When smaller-cap names return to vogue, it won't be possible for many who have stepped up the market-cap ladder to head back down, owing to their burgeoning AUM. Skillful small-cap managers who start moving up the cap ladder as assets increase also tend to be moving out of their comfort zone. This could certainly be the case with Virtus KAR Small-Cap Growth as Exhibit 4 shows. Here, the fund's average market cap increased from around \$1.5 billion to almost \$5.5 billion over the course of two years, a worrying sign.

Exhibit 4 Average Market-Cap Time Series for Virtus KAR Small-Cap Growth, the Russell 2000 Growth Index, and the Small-Growth Morningstar Category



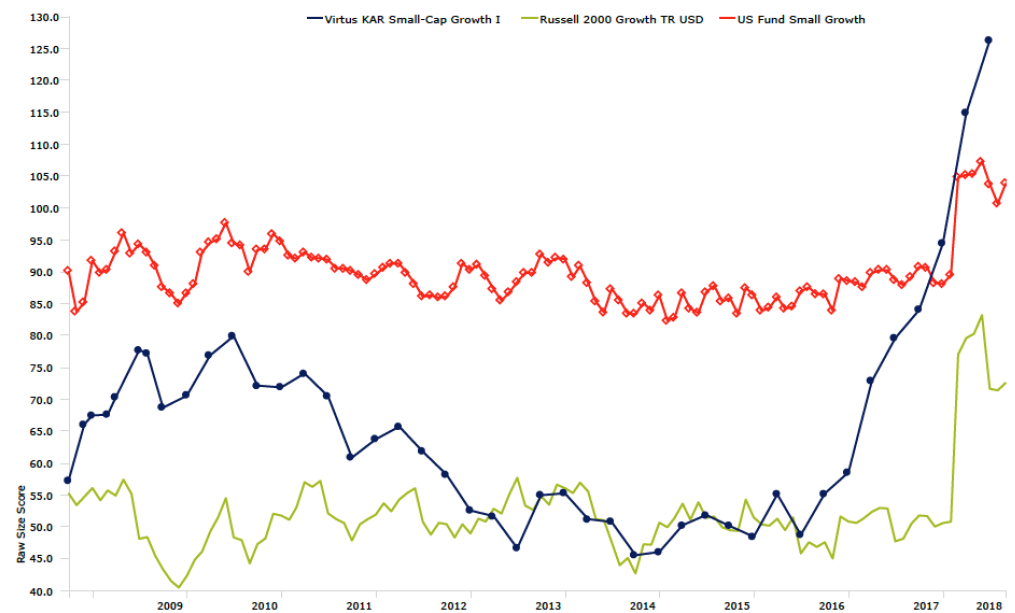
Source: Morningstar Direct.

Size Score

To get around some of the issues surrounding absolute versus relative movements in the average size of companies in a portfolio, we also use the size score. This is a more flexible system that isn't adversely affected by overall movements in the market. Giant-cap stocks are defined as the group that accounts for the top 40% of the capitalization of the Morningstar domestic-stock universe; large-cap stocks represent the next 30%; mid-cap stocks represent the next 20%; small-cap stocks represent the next 7%; and micro-cap stocks represent the remaining 3%. Each stock is given a size score that ranges from negative 100 (very micro) to 400 (very giant). The fund's size score is simply an asset-weighted average of the scores of its individual holdings. We can then track changes in that score over time to see if the portfolio has been any noticeable cap drift. Alternatively, we can review whether the percentage of assets allocated to certain size bands has changed through time. The data can similarly be viewed relative to a peer-group average or an index.

Again, Exhibit 5 shows a worrying development at Virtus KAR Small-Cap Growth, with the fund's size score rising to previously unseen levels. Its size score is also much higher than that of the average fund in its Morningstar Category, a position in which it has never previously been.

Exhibit 5 Size Score Time Series for Virtus KAR Small-Cap Growth, the Russell 2000 Growth Index, and the Small-Growth Morningstar Category



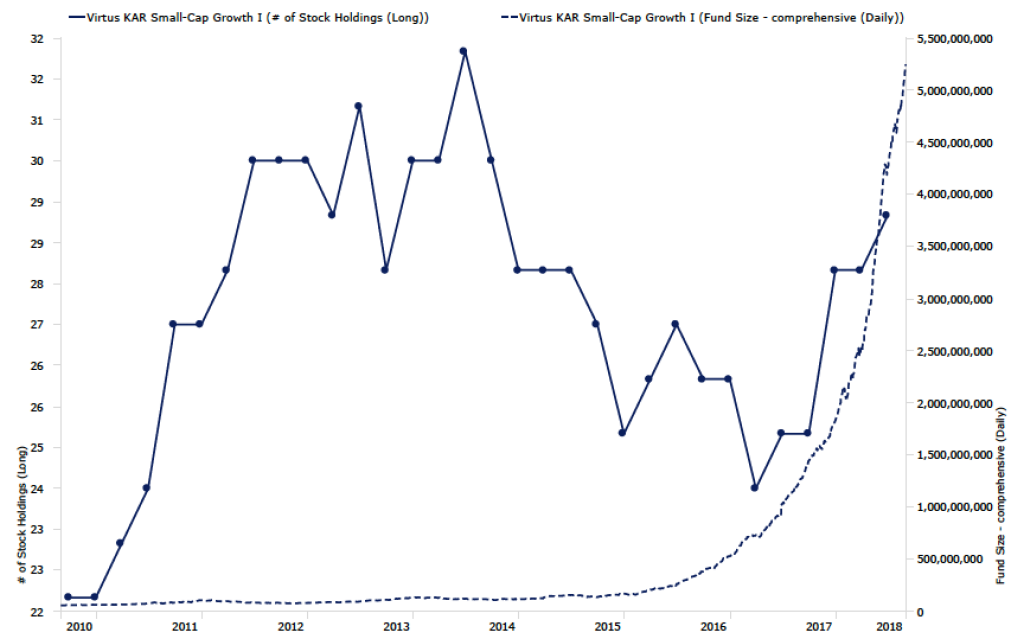
Source: Morningstar Direct.

Number of Holdings

A jump in the number of holdings in a portfolio can also be cause for concern. Good investment ideas are hard to come across, and it's therefore unlikely a manager has been keeping great ideas out the portfolio in the past. Increasing the number of holdings allows the fund to take on more assets while limiting liquidity issues, but it often signals a possible erosion of the fund's ability to generate the same levels of historic alpha.

The data here isn't too damning for Virtus KAR Small-Cap Growth. Although the number of holdings has increased as assets increased, the fund owned even more individual positions in the past when its asset base was far smaller. The conclusion here is that the managers have primarily opted to invest in more-liquid, larger-cap names to soak up inflows rather than spreading the money across more new portfolio holdings. The recent rise in the number of holdings is still significant, however, and remains something to monitor closely.

Exhibit 6 Stock Holdings and Fund Size Time Series for Virtus KAR Small-Cap Growth



Source: Morningstar Direct.

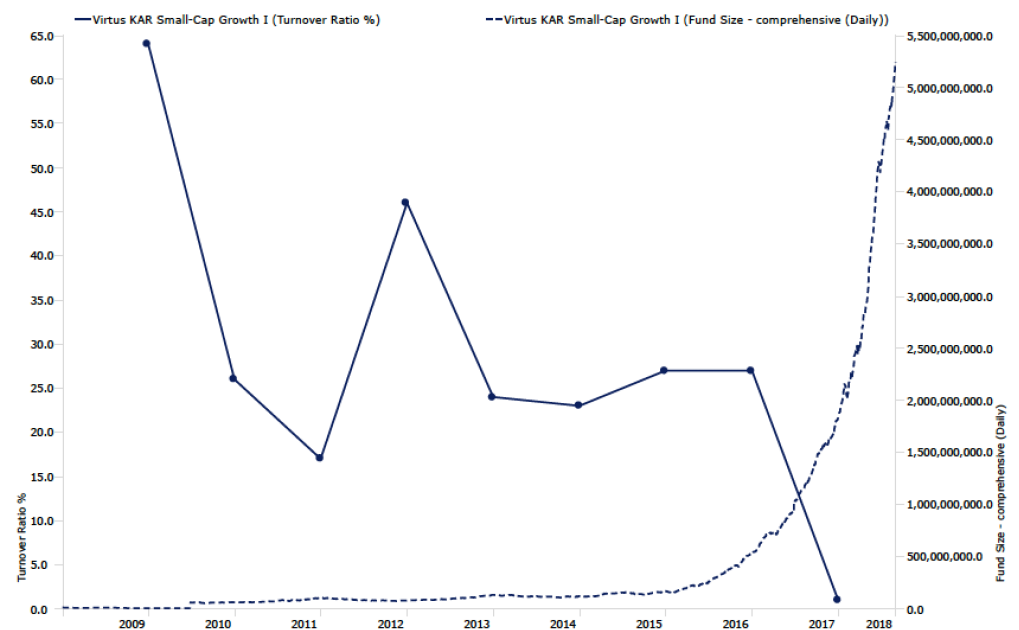
Portfolio Turnover

Portfolio turnover affects the assessment of capacity insofar as — all else being equal — higher levels of trading carry greater liquidity demands than lower levels of trading. Thus, a drop in portfolio turnover may be perceived as salutary, as it should limit trading costs. However, if a manager has had to trade less because the strategy's size has made it too costly to trade as much as they had previously, it could well be cause for concern (providing one thinks trading rapidly was a key to the manager's alpha generation).

High turnover often indicates tactical trading is taking place, something that's much harder to do at scale. Momentum strategies, for example, would be more at risk from liquidity issues. A corollary is that strategies that require less turnover may have a higher capacity than those that trade more rapidly. Further, contrarian offerings that tend to be liquidity providers could also potentially handle more money because of their proclivity to trade against the market. They would, however, be presented with at least the same problems as other strategies in the event of rapid outflows; and they could face even greater difficulties given they would generally be dumping shares for which there is limited demand in the market.

Virtus KAR Small-Cap Growth shows very interesting turnover data, with annual turnover actually falling to 0% for the past 12 months. As the turnover numerator is the lesser of purchases or sales, it tells us that while the fund is still trading (putting new money to work), its selling activity has been minimal.

Exhibit 7 Portfolio Turnover and Fund Size Time Series for Virtus KAR Small-Cap Growth



Source: Morningstar Direct.

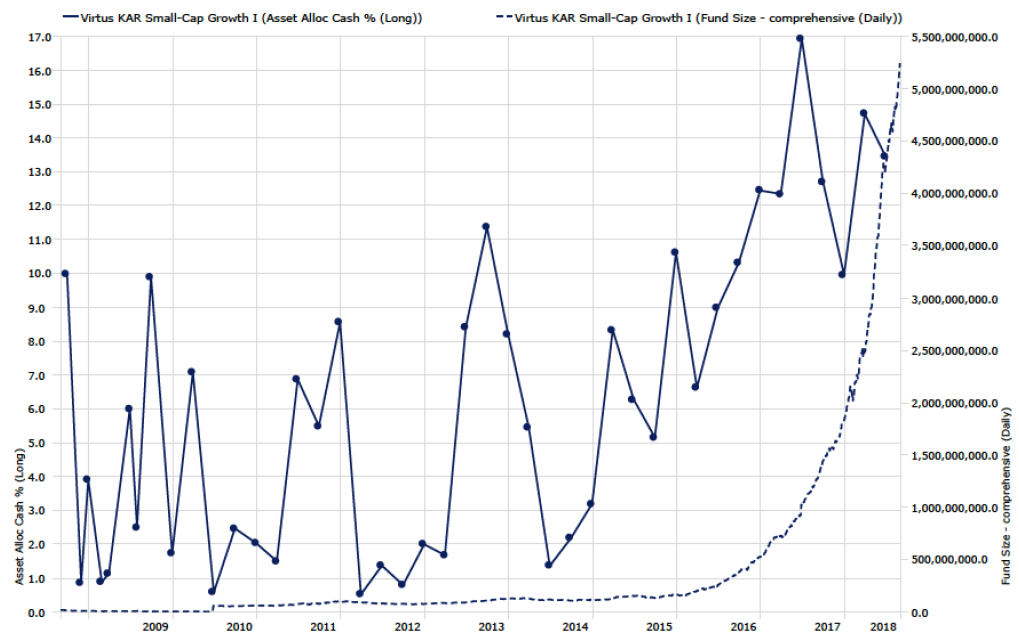
It should be noted that the denominator in the turnover calculation is average AUM for the period—so the huge inflows would have tempered the turnover figure anyway. However, a quick look at the portfolio indicates that the fund sold only one holding (MercadoLibre) over the 12 months to June 30, 2018. While this behavior is uncharacteristic (five stocks were sold in each of the three preceding 12-month periods going back to June 2014), it's not necessarily damning. But given that occurred at the same time the fund was hit with a large growth in AUM, it would be a strong area of discussion with management.

Cash Weighting

Another way to deal with taking on too much money too quickly is to allow cash in the portfolio to build. This gives a portfolio manager time to put those new assets to work, but it also waters down the performance of the rest of the portfolio in a rising market. The more illiquid names that are owned, the more cash we would expect to show in our end-of-month portfolios following large monthly inflows. There are ways around this (buying more of the portfolio's most liquid names or adding an index proxy exchange-traded fund, for example), but all likely dilute the fund's best ideas.

In Virtus KAR Small-Cap Growth's case, the cash stake has spiked at various points in its history. However, we have witnessed a noticeable rise in the strategy's cash stake following the large recent inflows. Prior to the late-2016 uptick in assets, we hadn't witnessed cash weightings at these elevated levels at any point in the fund's past, which likely tells us that the management team has struggled to put new money to work in a timely fashion.

Exhibit 8 Higher-Than-Average Cash Weightings



Source: Morningstar Direct.

Active Share

Active share measures the amount by which a fund's holdings diverge from those in the benchmark index chosen for comparison. Active share can be thought of as the summary output of the decisions taken by the portfolio manager. As such, changes in a fund's active share can give us an idea of whether the process is evolving from its past state. However, active share should be considered within the context of the other data points presented here. For example, a rise in active share may be considered good as it indicates the fund looks more different from the benchmark (it is now more active), but from a

capacity standpoint, it could mean that a small-cap fund with a small-cap benchmark has had to start buying mid-caps to deal with inflows. A drop in active share, on the other hand, may be an indication that a fund has had to move closer to its benchmark, perhaps by buying more of the larger-cap names in the index, to cope with inflows. Therefore, we look at any meaningful changes in active share to help us understand if there have been subtle changes in the process that haven't been picked up elsewhere.

Virtus KAR Small-Cap Growth's concentrated and benchmark-agnostic strategy has led to a consistently high active share, which currently stands at 98%. Therefore, this data point isn't much help here, as the portfolio is so differentiated already that even a boost up the market-cap ladder hasn't really moved the needle on active share.

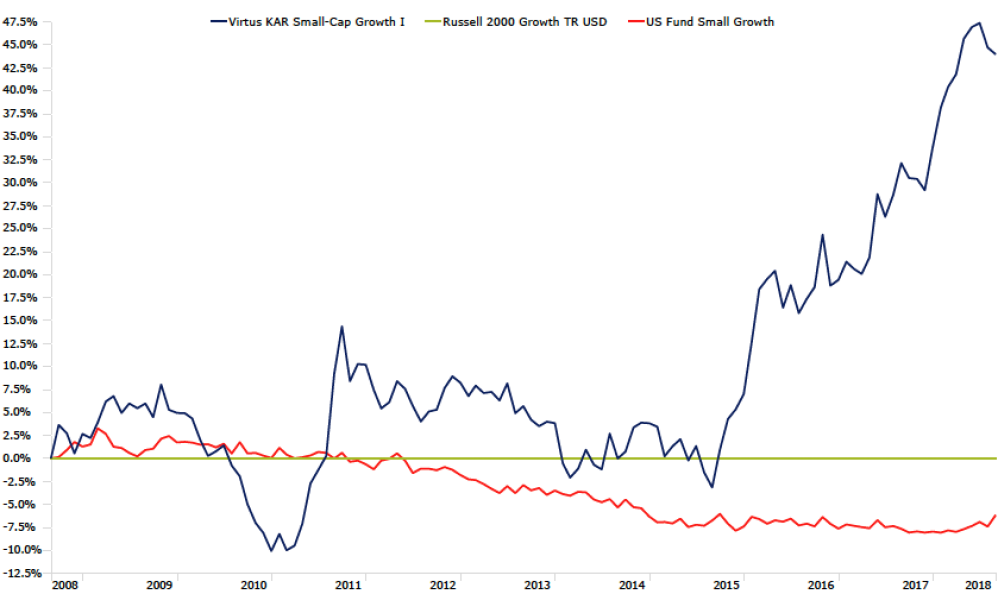
Performance

We know that strong relative performance can lead to inflows from investors chasing performance, but the point at which that may erode performance is more easily assessed by other data points. Here, we are more concerned with the relative performance of funds that have already closed owing to capacity. If a manager felt he could comfortably manage 0.5% of the available assets in his universe, the fund would likely close its doors before reaching that figure. Even so, if the fund continues to beat the market, its overall size in the market will continue to nudge up over time. It may be that the manager was conservative when he set the original close, but strong relative outperformance on a closed strategy would trigger a closer look.

Pengana Emerging Companies (an Australian smaller-companies fund that we rate highly) is a good example of this. The strategy closed its doors at roughly 0.5% of the free float of the S&P Small Ordinaries Index. It went on to significantly outperform over a number of years and was forced to provide large special distributions to return capital to shareholders in order to keep its assets under control. While handing money back shows strong stewardship, it can cause headaches for investors who receive unwanted tax bills and must find new places to park their capital.

Exhibit 9 shows the extent of Virtus KAR Small-Cap Growth's historical excess return above the index. If such market-beating performance occurs in future years, the fund's relative share of the market would continue to grow, despite it having closed to new investors. This is why we greatly prefer to see funds close earlier and provide some runway for outperformance or for existing investors topping up (the Virtus fund only closed to new investors in September 2018).

Exhibit 9 Strong Relative Outperformance for Virtus KAR Small-Cap Growth



Source: Morningstar Direct.

Market-Level Capacity

Market level refers to both the fund's investable universe and the group of comparable peer funds. Unlike fund-level data, market-level data points and their relevance are entirely dependent on the market, sector, and style of the fund in question.

Size Versus Investment Universe and Category Peers

For this data point, our analysts look at the size of a strategy relative to those that have the same investment universe. The bigger a fund is relative to its peers and the size of its opportunity set, the more market impact it will likely have. Again, though, how large a fund can get in a specific universe before it encounters problems depends on a number of factors (investment style, market liquidity, and so on). There's no hard number, but all things being equal, the biggest funds in any category would obviously be put under greater scrutiny from a capacity standpoint.

Outside of the raw data, our analysts go further and subcategorize funds depending on how they invest. For example, we don't have specific micro-cap categories, but we do have funds in our small-cap categories that focus on the micro end. Therefore, we assess them relative to similar funds that invest in the same area of the market; while those funds may be small relative to the category, we would pay closer attention to large inflows that push them above peers that also invest at the very bottom of the market-cap ladder.

In the case of Virtus KAR Small-Cap Growth, the fund benchmarks itself to the Russell 2000 Growth Index, which had a total market value of close to USD 1.17 trillion at the time of publication. If we take that as a proxy for the fund's opportunity set, it would mean approximately 0.45% of its investable universe. However, the strategy (along with the majority of those listed below) invests significantly outside of that index, so in this case the number is less relevant. Then, when simply compared with its small-growth category rivals, the Virtus KAR fund is by no means the largest—placing ninth in terms of fund size.

Exhibit 10 The Largest Strategies in the U.S. Small-Growth Morningstar Category

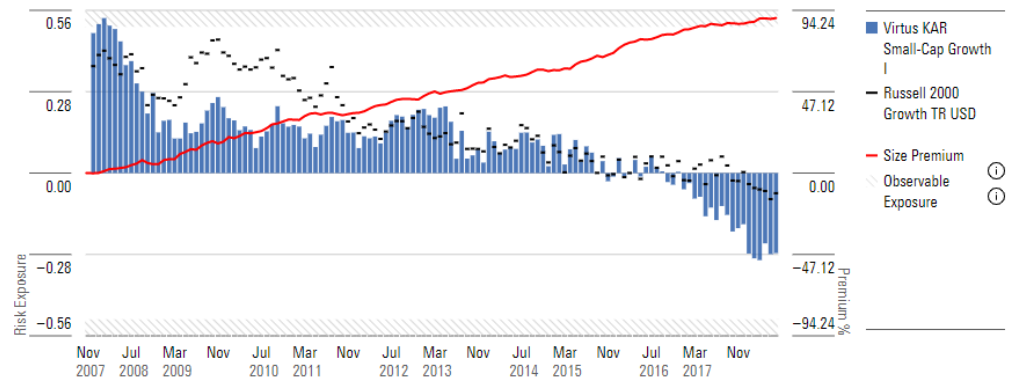
Name	Ticker	Fund Size USD	Average Market Cap		% of Russell 2000 Growth	Number of Holdings
			2016-06	2018-06		
T. Rowe Price New Horizons	PRNHX	27,249,362,902.00	4,134.65	6,519.91	2.33	192
Vanguard Small Cap Growth Index Inv	VISGX	25,729,087,094.00	3,063.29	4,374.92	2.20	653
Vanguard Explorer Inv	VEXPX	16,915,214,202.00	2,725.76	4,535.92	1.45	606
Janus Henderson Triton D	JANIX	12,100,110,271.00	3,336.83	4,783.68	1.04	124
Neuberger Berman Genesis Inv	NBGNX	11,324,293,900.00	3,311.46	4,040.00	0.97	124
T. Rowe Price Small-Cap Stock	OTCFX	10,675,877,966.00	2,111.89	2,965.11	0.91	279
T. Rowe Price QM US Small-Cap Gr Eq	PRDSX	8,065,088,362.00	3,351.72	4,548.20	0.69	305
Brown Capital Mgmt Small Co Inv	BCSIX	6,020,090,618.00	2,650.40	4,275.23	0.51	37
Virtus KAR Small-Cap Growth I	PXSGX	5,295,176,899.00	1,576.59	5,407.52	0.45	30
T. Rowe Price Instl Small-Cap Stock	TRSSX	5,116,524,545.00	2,119.25	2,969.98	0.44	276
Russell 2000 Growth TR USD	NA	1,168,954,998,240.00	1,511.85	1,849.81	NA	1,212

Source: Morningstar Direct.

What is more interesting, however, is the change in average market cap alongside fund size. We already noted the dramatic increase in average market cap in the previous section, but this brings home the point with the fund moving from owning the smallest stocks on average among the largest managers, to owning the largest stocks in the two years that it grew rapidly. Further, the strategy is the most concentrated of the largest strategies that we cover, which adds further pressure relative to its larger peers.

Size Factor Exposure

We also use the Morningstar Risk Model to help us assess capacity. The size factor measures the market cap of stocks in a fund's portfolio. Higher scores indicate greater investment in smaller-cap names. The size factor can also be used to judge style drift. Virtus KAR Small-Cap Growth has historically had a significant exposure to the small-cap factor, both absolutely and relative to its Russell 2000 Growth benchmark. As we can see below, that started to change in 2016 and 2017, just as the fund was experiencing its asset growth. The information included here is similar to what we see from a fund's size score (highlighted earlier). However, we can also examine the size premium (the red line in the chart below shows how positive or negative the exposure to the size factor has been over time).

Exhibit 11 Declining Monthly Size Factor Exposure

Source: Morningstar Direct.

Competing Capacity Estimates

Our analysts talk to a lot of managers (around 2,500 interactions each year), and we ask them all how much money they think they can run in their current style. Analysts recheck this number in future meetings to understand how their thoughts around capacity might be changing. Therefore, we have a large amount of data on how much money that fund managers believe their strategies can comfortably handle. We dig deeper when we hear capacity estimates being moved up as a fund's assets grow—it may be the case that managers upped the number because their investable universe expanded, or it may be the case that business pressures are getting the better of them.

We also use this information to help us understand if a strategy has capacity limits that are outside norms for its peers. If certain managers think they can run significantly more than their peers feel comfortable running, we want to understand exactly why, and what they might have to give up by running a larger fund than is typical.

Closure Points of Competing Funds

While capacity estimates can be flexible, data around soft and hard closures of funds isn't. It's rare for managers to turn away money, unless they think it's to their long-term detriment. Therefore, by looking at the size at which similar funds closed, we get a very clear view from managers who are dealing with the day-to-day challenges of running large books of money as to just how much is manageable. Virtus KAR Small-Cap Growth made the decision to close to new money on Sept. 28, 2018, after more than doubling its asset base since the end of 2017.

Ownership Popularity Factor Exposure

Again, we turn to the Morningstar Risk Model to understand ownership popularity. This factor looks at what other managers have been buying. A high exposure to this factor indicates that more funds have taken long positions in stocks within the portfolio, relative to those that have shorted the stock in the past three months. This is a useful tool for assessing capacity, as it can help indicate if a fund is moving into more-crowded trades. Any behavioral changes relative to a fund's long-run history can also be

telling. For example, a fund that consistently went against the grain when it was small, then increased its exposure to the ownership popularity factor after substantial AUM growth, could indicate a manager who is happy to bank his historic performance and stick with the herd going forward.

Liquidity Analysis

We have developed several innovative new tools that allow us to more easily see if a fund is running into liquidity issues. These rely on using our holdings database and Morningstar equity data to drill deeper into funds' portfolios. A fund that is growing in assets may not raise a flag as having capacity issues on any fund- or market-level screens because the manager hasn't yet changed the way it's being run. For example, if AUM increases fivefold from a small base, the manager could maintain the same number of portfolio holdings and keep topping up in the same names, meaning the average size of the companies held wouldn't increase. However, the manager would be creating an increasingly more illiquid portfolio. While this may not immediately impact the portfolio, it would make it much more difficult to unwind positions in a timely manner.

Our new screens also allow us to stress-test any portfolio in various liquidity environments to see, for example, how quickly a manager could sell a set percentage of his portfolio, or how long it would take to put a certain amount of new money to work. We can also model how a fund's liquidity profile would look at various sizes to help us understand if its capacity estimates appear sensible.

Understanding Our Liquidity Assumptions

Assessing available liquidity is incredibly difficult. It's influenced by many factors and often dries up when it's most needed. We therefore use historical lows for stress-testing liquidity. Our estimates use the lowest 30-day average daily trading volume for each individual stock going back to 2000. We then take 25% of that number — as we don't believe that one fund would be the only entity looking to sell at any one point (especially in a liquidity event). It's also a rough proxy for the maximum portion of trading volume a fund could access without alerting the market to its intentions. We then take the number of shares a fund owns in each of its holdings and divide it by our new daily trading volume estimate to arrive at the approximate number of days it would take to fully liquidate that position.

In the examples shown in Exhibit 12, the thick black line shows the estimated time to liquidate the portfolio, or its "burndown." While the burndown profile is arbitrary (it's highly unlikely that managers will need to fully liquidate their entire portfolio, and it's also unlikely that we correctly predict the available liquidity should the need arise), it does give insight into the amount of stocks that could be problematic to liquidate. For example, if a stock is still around after 10 days of trading, then the manager is likely going to have to weather whatever he was trying to avoid.


Given that our initial assumptions regarding available liquidity as depicted on the thick black line are conservative (25% of the minimum 30-day average volume for each stock since 2000), we also provide burndown lines using two more-generous assumptions. The red line estimates the burndown based off 25% of the average volume of all the stocks in the portfolio over their lifetimes, and the green line uses 25% of the latest 30 days' volume. These measures provide additional context and allow our analysts to see if the current portfolio is likely to have liquidity issues now, or if the problems are more likely to impact it during a liquidity event.

The y-axis indicates the percentage of the portfolio remaining (starting at 100%), while the x-axis indicates the trading days it would take to exit (261 trading days is one calendar year).

It is important to note that we are not modeling actual liquidations. We recognize that a manager could deploy any number of strategies to get the cash for redemptions without affecting the portfolio exposures too much. What the modeling shows is how greatly a manager might have to deviate from his preferred criteria for buys/sells and portfolio construction. For example, we wouldn't want a fund to be forced to sell only its most liquid holdings as it would change the makeup of the portfolio. We use many of the aforementioned fund- and market-level data points to help inform us as to how likely large-scale redemptions are. For example, a sharp rise in flows after a period of strong performance could quickly unwind after a period of poor performance, and we would want to make sure the portfolio could cope if flows reversed. We also look closely at the largest owners of the fund to see if there are holders who own such a significant percentage of assets that a redemption would create liquidity issues.

Number of Days to Liquidate the Portfolio

While it would be an unlikely and unfortunate set of circumstance that would lead a fund to have to liquidate its entire portfolio, understanding how long it would take for that to happen is a useful tool in gauging how liquid its holdings are. Exhibit 12 shows the liquidity profile of Virtus KAR Small-Cap Growth.

Exhibit 12 Virtus KAR Small-Cap Growth's Burndown Chart

Source: Morningstar Direct. Data as of June 30, 2018.

Virtus KAR Small-Cap Growth closed to new money in September 2018. In that month, the strategy reached \$5.295 billion in size. The data used by our liquidity screener uses the portfolio size at the last available portfolio date. For that reason, we are using the fund's \$4,257 billion size at June 30, 2018. You will see later how we are able to model the fund's actual size and scenario-test portfolios based on a manager's capacity estimates.

The black line only reaches down to roughly 13% on the y-axis after 261 trading days (a full year's worth). This indicates the manager needs more than 12 months to liquidate the entire portfolio based on our worst-case scenario (25% of each individual holding's lowest 30-day average daily volume since 2000). When we use the most recent trading volumes the picture improves, with around 95% of the portfolio liquidated after 130 days. However, selling down the last bit to achieve full liquidation would stretch the time needed to around a year. The grey lines indicate the liquidity profile of the historical portfolios, quickly illustrating that our latest portfolio is the most illiquid we have seen from this fund, a reflection of its huge growth in assets.

A full liquidation is rare, but a 20% redemption could certainly be feasible. In this example, we can see that would take just around three weeks to achieve that if liquidity dried up to historical lows. More comfortingly, we can see that the managers could liquidate over 50% in just 10 days based of recent trading volumes. It appears much of the illiquidity is in the tail of the portfolio, which is not unusual but still bears monitoring in a small-cap offering of this scale.

Highlighting Illiquid Stocks

Taking longer to sell individual holdings than one would like is one of the main downsides of taking on too much money. If something changes in an investment thesis and a fund needs to get out quickly, it can be very difficult if it owns a relatively illiquid name. For this reason, we also look at the number of days it would take to sell out of each individual position (based on the worst-case 30-day trading volume assumption, or burndown rate, below). While we can see the full liquidity profile for every stock in the portfolio, we have focused here on the five that would take the longest time to liquidate.

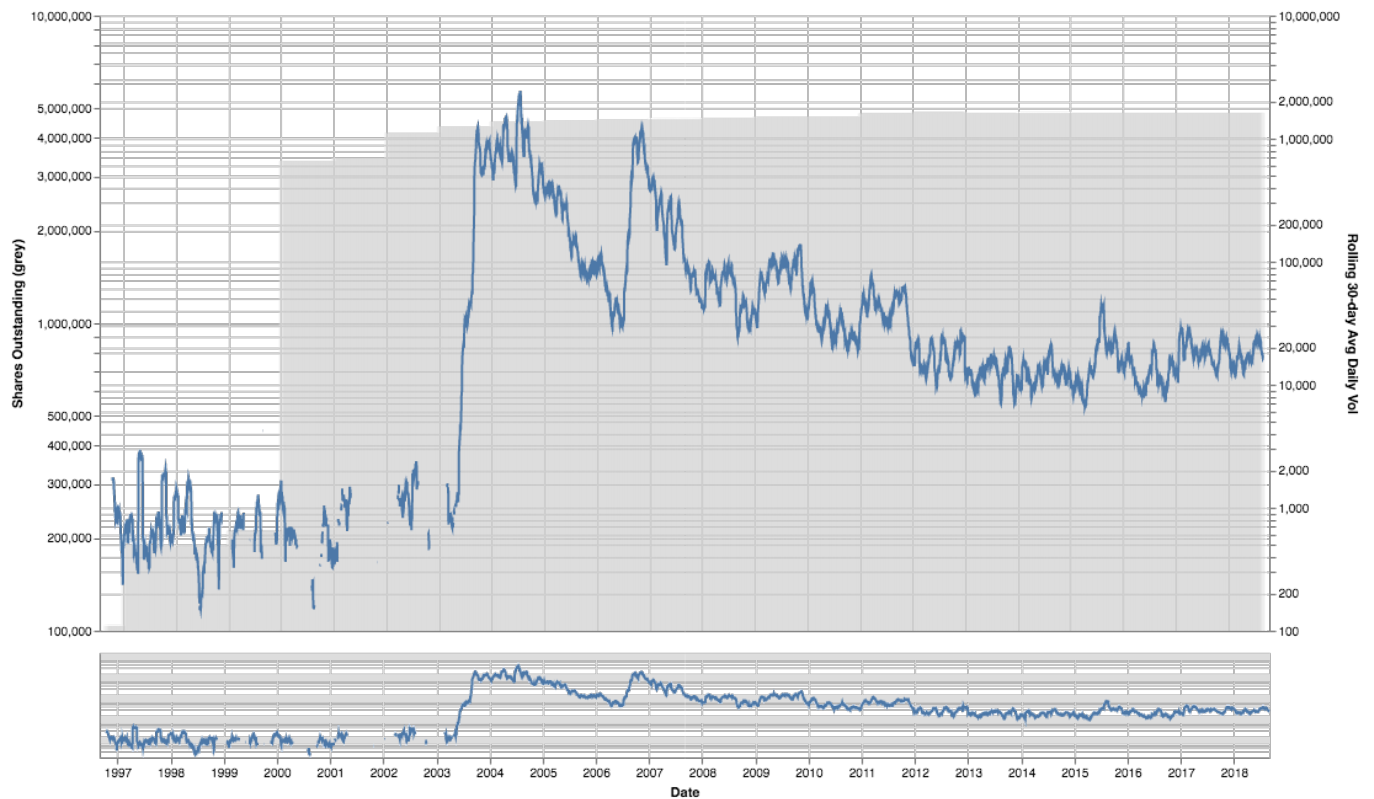
Exhibit 13 Days-to-Liquidate Analysis for Virtus KAR Small-Cap Growth

Name	Portfolio Weight	Number of Shares	Days to Liquidate	Burndown Rate	30 Day Volume	
					Average	Latest
NVE Corp	1.647	483,690	5,624	86	135,257	12,390
Omega Flex Inc	1.726	780,514	3,752	208	7,645	13,375
Pricesmart Inc	2.29	905,200	1,286	704	102,519	102,348
National Research Corp Class A	1.921	1,836,730	1,118	1,643	20,020	23,292
Old Dominion Freight Lines Inc	5.291	1,270,500	700	1,815	669,113	585,096

Source: Morningstar Direct. Data as of June 30, 2018.

It appears that NVE Corp. has had incredibly low average 30-day trading volume in its past, leading to the estimated 5,624 days' worst-case estimate for liquidation. While the lifetime average 30-day volume number is much better (meaning the stock could be sold in fewer than four days), the latest 30-day average volume of just 12,390 shows that volume in this stock is unpredictable.

We can also take a deeper dive into the historical liquidity profiles of individual equity holdings. Given the disparity between its different volume measures, NVE Corp. is a candidate for further review. Exhibit 14 shows that the latest average 30-day trading volume number seems a more reasonable assessment for the current market environment than the lifetime average number; we can adjust the burndown rate in our model to see how long it would take to liquidate at that level (39 days).

Exhibit 14 NVE Corp.'s Historical Trading Volume

Source: Morningstar Direct.

Ownership and Free-Float Stake

Looking directly at the percentage of a company that a strategy owns helps supplement the number-of-days-to-liquidate analysis by removing the need for volume assumptions. It's especially useful when assessing how a strategy might behave in a liquidity event.

Also, the more a fund owns, the more visibility the market has into what it is doing. Various filings are required to notify the public of trades. What is required depends on which market you are in and what that local regulator demands. For example, in the United States and Australia, the Securities and Exchange Commission and Australian Securities and Investments Commission, respectively, require that a fund notify the market if it acquires beneficial ownership of more than 5% of a voting class of a company's equity securities. Meanwhile, in the United Kingdom, London Stock Exchange rules dictate that an investor who acquires, exceeds, or falls below 3% (and each 1% threshold above 3% up to 100%) of a UK company must disclose his interests to the LSE and the company within two days. These requirements can put a manager at risk of front-running by other market participants who may be able to see positions being built or exited.

We look at both ownership stake, which is the absolute percentage of a company owned, as well as free-float percentage, which gives us an idea of how many of the available shares on the market a fund owns. Below, we can see this analysis for Virtus KAR Small-Cap Growth. Again, we have chosen to only display the five largest holdings, this time ordered by their overall company ownership stakes.

Exhibit 15 Underlying Ownership Analysis for Virtus KAR Small-Cap Growth


Name	Ownership Stake %	Free Float Stake %	Portfolio Weight %
NVE Corp	9.989	12.071	1.647
Fox Factory Holding Corp	9.817	14.635	4.832
The Chefs' Warehouse Inc	8.727	15.423	2.084
Omega Flex Inc	7.734	27.06	1.726
National Research Corp Class A	7.424	16.805	1.921

Source: Morningstar Direct. Data as of June 30, 2018.

NVE Corp. again comes out as a stock that we would like to know more about. The fund owns 10% of the company and over 12% of the free float. That information, coupled with its seeming illiquidity, make it an important holding for the fund. The large free-float weighting that the fund owns in Omega Flex would also lead to further analysis.

Changes in Historical Liquidity

As shown in Exhibit 12, we can use our burndown analysis to assess the liquidity profile of a fund's current and historical portfolios. While this helps us assess changes in a fund's liquidity profile through time, we would also like to understand how assuming different asset levels would have affected the fund's historical profile. For example, in Exhibit 16, we model the liquidity of Virtus KAR Small-Cap Growth's historical portfolios assuming an AUM level equal to the fund's size at the time of writing (\$5.295 billion).

Exhibit 16 Virtus KAR Small-Cap Growth's Changing Liquidity Profile

Source: Morningstar Direct. Data as of June 30, 2018.

The above analysis shows that once asset size is controlled for, today's portfolio (the thick black line) is more liquid than any of the previous portfolios (grey lines) we have seen. When we see a profile like this, we want to make sure that the managers haven't just topped up their most liquid names when new money has come through the door, as that would undoubtedly dilute the future impact of the less liquid end of its portfolio. Simple portfolio attribution can also tell us the extent to which less-liquid names added to performance when the fund was smaller. If they drove a substantial amount of performance, and if investment in less-liquid names is now restricted by the fund's scale, there would be clear reason to question the repeatability of the fund's past performance.

Analysis of a Strategy's Capacity Target


When asked how much money their strategy can handle, most managers have a number readily at hand, but the thought put into arriving at that number can vary greatly. We therefore prefer to run our own estimates where possible. One way of doing that is through modifying the above liquidity analysis. For example, if a manager is currently running \$1 billion in assets but thinks he can run \$5 billion in his current style, we can estimate the liquidity profile of today's fund as if it had \$5 billion of assets. We can also look closely at the changes in ownership and free-float weightings at that level of AUM to see how feasible it would be to run that level of money with today's portfolio. This helps give us an idea if a fund would have to change its approach to accommodate the expected future asset growth.

Number of Days to Buy

When thinking about capacity, the focus is often on how long it would take to exit positions, but it's also useful to consider how long it would take to put new money to work proportionately across a portfolio, as this can also cause style drift that could erode alpha generation. This is because a manager could invest the money into the fund's most liquid holdings and get the job done more quickly; that would change the way he invests and risk degrading the properties that have helped the strategy outperform.

To look at this analysis, we run a model that estimates how many days it would take to buy the current portfolio weighting in all of the positions in a given portfolio assuming a 10% increase in AUM. The methodology is similar to the burndown methodology discussed previously, but here we use 25% of the closest 30-day average trading volume. This gives an estimate of how long it would take to put that extra 10% of assets to work given relatively recent trading volumes.

Exhibit 17 Virtus KAR Small-Cap Growth's Number of Days to Buy



Source: Morningstar Direct. Data as of June 30, 2018.

In Exhibit 17, we can again see the impact of the fund's relatively illiquid tail of holdings, as it takes over a month to fully replicate the portfolio. This again indicates that the more illiquid ideas in the portfolio will likely get diluted by strong inflows. Interestingly in the case of this fund, it received close to 20% of new money in only a couple of months; the above exhibit helps highlight just how difficult it would have been for the managers to replicate their portfolio in that scenario.

Conclusion

There is no universal maximum amount of money a certain type of strategy can run. While grouping funds by investment style, market-cap bias, portfolio holdings, and so on does get you closer to the answer, there are many variables at play that result in very similar looking funds being able to comfortably run vastly different amounts of money.

In addition to the process and portfolio construction deployed by the manager, the amount of money that a strategy can run is affected by the point in the cycle we are at. For example, during the period following the financial crisis—that is from 2009 onward—running a large asset book hasn't proved particularly problematic. During this 10-year bull market, we haven't seen any noticeable marketwide liquidity squeezes for active equity managers, and volatility has been at historical lows. This abundant liquidity has allowed even the largest active managers to put new money to work with comparative ease. The largest (and most liquid) stocks in the MSCI World Index, for example, have also been among the best performers during this period (the 10 largest stocks in that index returned almost 20% per year more than the index as a whole over the past five years). Therefore, investing in market heavyweights like Apple, Amazon, Facebook, Google, or Microsoft if a fund has felt a capacity pinch has proved highly accretive in most cases.

While we have seen a gradual rotation out of active management and into passive strategies during this period, those outflows have been relatively measured. Indeed, as we near the end of 2018, we haven't yet seen a coordinated period of panic selling (like 2008) that adds real pressure to capacity-constrained strategies. Further, there has been no shortage of buyers willing to mop up sales that were required to meet redemptions. Nevertheless, it's called a market cycle for a reason, and it's important to look through it and not become complacent.

All the tools and techniques that we have discussed here are used to help us understand if a strategy may feel the pinch from taking on too much money. They aren't used in isolation; they simply add to our analysis of a strategy. They are also complemented by due diligence on all aspects of a fund's management—from the skills of the people involved, the process behind it, and the way the portfolio is constructed.

While total AUM is important to assessing a fund's current liquidity challenges and comparing them with those it faced in building its past record, changes in AUM owing to inflows or outflows and the speed thereof are what can trigger changes in the way a fund is managed. If we think that increased

AUM limits the manager's opportunity set or requires the manager to change his style, we will reassess the merit (or lack thereof) of the fund in its new context and change its Morningstar Analyst Rating in accordance with our findings.

As the available market and underlying liquidity of its stocks grow and shrink from day to day, our assessments of capacity are ongoing. And while we don't hold funds to our worst-case scenarios when rating them (as we understand the relatively arbitrary nature of any capacity or liquidity assumptions), we do want to make sure that management has considered the same possibilities and understands the limits of its current approaches.

The fund we chose as our example, Virtus KAR Small-Cap Growth, is a strategy that is clearly changing its stripes following dramatic increases in AUM in relatively short order. For that reason, we have held the fund's rating at Bronze, despite having gradually increased our conviction in its management over time. The inflated asset base leads us to question if the same level of outperformance is possible in the future. Yet, we are still happy to have a positive rating here, as we believe that, although it is a different proposition to what it once was, it still has the potential to at least outperform the average fund in its category.

About Morningstar Manager Research

Morningstar Manager Research provides independent, fundamental analysis on managed investment strategies. Analyst views are expressed in the form of Morningstar Analyst Ratings, which are derived through research of five key pillars—Process, Performance, Parent, People, and Price. A global research team issues detailed Analyst Reports on strategies that span vehicle, asset class, and geography. Analyst Ratings are subjective in nature and should not be used as the sole basis for investment decisions. An Analyst Rating is an opinion, not a statement of fact, and is not intended to be nor is a guarantee of future performance.

About Morningstar Manager Research Services

Morningstar Manager Research Services combines the firm's fund research reports, ratings, software, tools, and proprietary data with access to Morningstar's manager research analysts. It complements internal due-diligence functions for institutions such as banks, wealth managers, insurers, sovereign wealth funds, pensions, endowments, and foundations. Morningstar's manager research analysts are employed by various wholly owned subsidiaries of Morningstar, Inc. including but not limited to Morningstar Research Services LLC (USA), Morningstar UK Ltd, and Morningstar Australasia Pty Ltd.

For More Information

Mike Laske
Product Manager, Manager Research
+1 312 696-6394
michael.laske@morningstar.com



22 West Washington Street
Chicago, IL 60602 USA

©2018 Morningstar. All Rights Reserved. Unless otherwise provided in a separate agreement, you may use this report only in the country in which its original distributor is based. The information, data, analyses, and opinions presented herein do not constitute investment advice; are provided solely for informational purposes and therefore are not an offer to buy or sell a security; and are not warranted to be correct, complete, or accurate. The opinions expressed are as of the date written and are subject to change without notice. Except as otherwise required by law, Morningstar shall not be responsible for any trading decisions, damages, or other losses resulting from, or related to, the information, data, analyses, or opinions or their use. The information contained herein is the proprietary property of Morningstar and may not be reproduced, in whole or in part, or used in any manner, without the prior written consent of Morningstar. To license the research, call +1 312 696-6869.