

Fixing the Peer Group Problem

HOW TO MEASURE PERFORMANCE WHEN NO SUITABLE PEER GROUP IS AVAILABLE

By Antti Raappana, CFA, Kimmo Kurki, Fernand Schoppig, and Barry Gillman, CFA

Peer group comparisons are a key element in asset manager evaluation and selection. The peer group floating bar chart is a familiar and easily understood tool that enables asset owners, consultants, and managers to assess whether a particular fund has done well or poorly against competitors with similar mandates over the same time period. Furthermore, peer groups are used as a reflection of a true range of potential portfolio returns in order to evaluate how significant the performance difference is between the portfolio and the “market average” (i.e., how skillful the manager is).

But there are inherent problems for users of this approach. First, the use of published peer groups may provide an assessment against a broad universe of comparison, but its drawbacks include survivor bias, composition bias, timeliness, and (especially for the broadest peer universes) mismatches against the mandate being measured. Published peer groups can give you an estimate of how well your portfolio has done against competitors with similar benchmarks but whose mandates may differ from yours. What they *can't* do is give a realistic assessment of the range of outcomes your manager could have achieved by making different decisions within the constraints of the mandate. In other words, using the wrong peer group gives an inaccurate picture of the true distribution of potential outcomes and might lead to false conclusions on the significance of relative performance.

Second, for some mandates, a suitable peer group may not even exist. Such is the situation for mandates that combine an “unusual” mix of countries, have constraints in a single country, or have systemic biases in terms of sectors or styles that reduce the peer group universe to few or no peers. In this circumstance, the problem is often

compounded by the lack of a suitable benchmark index. Published indexes tend to be classified either by country/region or by size/style; when the mandate does not fit neatly into one of these classifications, the investor is limited to trying to combine two or more indexes into a “custom blend.” Even when this works to create a suitable return benchmark, the investor likely loses access to the index analytics provided by index publishers and, importantly, loses the ability to evaluate the significance of return deviations.

There's no commonly accepted set of tools currently available for investors to overcome these problems. As a result, investors and their managers often end up disagreeing over whether the manager's performance is good or bad given the circumstances and mandate constraints.

The pioneering work on the first problem was developed by Ron Surz of PPCA, who created portfolio opportunity distributions (PODs)—simulated peer groups comprising all the realistic portfolios that can be constructed from the constituents of a specific index. PPCA publishes PODs that provide peer group comparisons for the US and other developed markets.

We have used the ideas behind Ron Surz's work to solve the second problem: the lack of suitable peer groups for unconventional mandates. Our challenge was to develop suitable comparisons where no usable peer groups existed, particularly for specialized small- and mid-cap mandates in regional combinations within the emerging and frontier markets. In the InvestWorks database as of 31 December 2016, there were over 2,000 products in the US large-cap equity mutual fund peer group, but there were none at all if you needed a Latin-America-minus-Brazil peer group. A similar problem existed if the mandate was Southeast

Asia with only limited exposure to Singapore. In these cases and others, we faced the problem of evaluating how much of the performance measurement was due to mandate specifications versus successful stock picking.

Our solution to the evaluation problem was to build a simulation model that addressed the lack of a representative market index and having too few (or no) competitors. We did this by, in essence, simulating the range of possible portfolio-return outcomes that could have been achieved by following the specific constraints of the selected manager. These constraints included country, market-cap size, number of stocks held, and maximum weight in a stock or sector. Using a Monte Carlo simulation, every randomly generated portfolio with those mandate constraints represented a potential peer portfolio that could have been achieved without stock-selection skills.

By running thousands of simulations to include all the available portfolios that could have been constructed using a particular mandate's constraints, a comparable peer group portfolio-return distribution was generated. We could then see how the manager's actual results over that period compared to the range of possible results that could have been achieved under that manager's specific portfolio constraints. This approach also helps in a situation where no appropriate benchmark index is available.

The advantage of these simulations is that they are tailor-made to the specifics of each manager's own stated approach and not affected at all by the number of competing managers whose mandates may not be truly comparable. This approach can provide more timely comparisons as the data-collection time is reduced and also allows for simulated peer groups to be generated for non-standard time periods.

The process for a one-period simulation consists of the five steps below:

1. Screen for stocks that fulfill the mandate constraints (exchange of listing, country of operation, size, trading liquidity) at the beginning of the period and that have been listed for the full simulation period so that a full-period total return is available. This constitutes the simulation universe.
2. Calculate the total return for the period in the evaluation currency of each stock in the universe.
3. Impose portfolio construction constraints (number of stocks, maximum/minimum exposure, country/sector exposure, maximum cash, and any other potential constraints).
4. Simulate random portfolios that fulfill those constraints from the eligible stock universe.
5. Compare actual portfolio performance to the distribution of the random portfolios.

Note that although we have used one year as our full period for the purposes of illustration in this article, typically a one-year evaluation period is not long enough to draw strong conclusions on stock-selection skill. Consistency over time is also important. We believe that an evaluation period of at least three years is required to be effective in practice. The simulation approach is also useful in situations where mandate constraints have changed from one year to another.

One important advantage of the simulation approach is that it can be used for different time horizons. In practice, there are two possible approaches to lengthening the observation period:

- A. Running the shorter, one-period simulations sequentially for the desired number of years and then linking the individual period returns to construct a multiyear return for each sample portfolio, and
- B. Running each simulation over the desired observation period, generating a total return for the full period.

We have found Method A to be more practical because it allows individual stocks to enter and exit the simulation

at year ends, enabling us to take into account stocks listing or delisting (or otherwise no longer meeting the mandate criteria) during the multiyear period.

In addition, shorter simulation periods are less prone to biases caused by extraordinary individual stock returns. Unless any rebalancing methodology during the single-period simulation is applied, there is a possibility that the allocation within any randomly generated portfolio skews unrealistically toward an exceptionally well-performing stock. This problem is more pronounced as the length of the simulation period increases because performance dispersion is likely to increase with time. In practice, this means each portfolio has close to 100% turnover after every simulation period. Although this may be a somewhat unrealistic assumption, we believe it produces more realistic results than the alternative approach, with its implicit buy-and-hold assumption that includes no rebalancing or changes to the portfolio over an extended period.

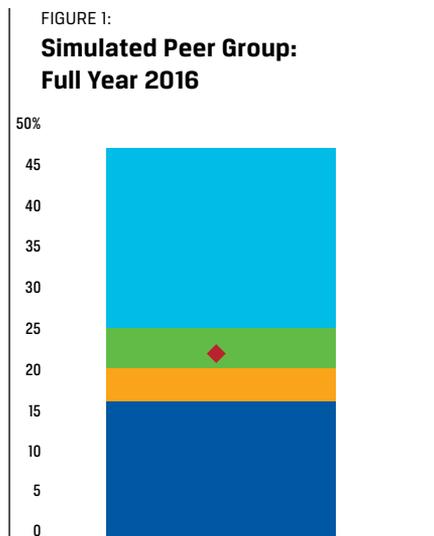
In Figures 1 and 2, we show one- and three-year examples using our selected

method. The mandate in this case is one where we could find no suitable peer group or index. The portfolio is invested in small and mid-sized stocks in selected regional emerging markets.

In this simulation for 2016, the red diamond shows that the manager's return of 23.0% was in the middle of the second quartile—above the simulated median of 21.3%—and below the breakpoint at the top of the second quartile (25.9%).

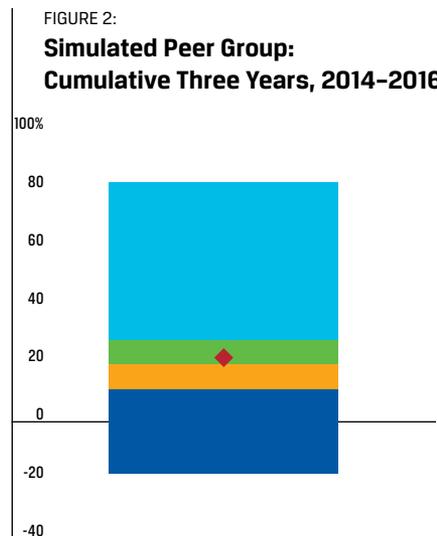
When the simulation was run for the three-year period, the manager's return (red diamond) remained in the second quartile, with a cumulative return of 20.4%—above the simulated median of 18.9%—and below the breakpoint at the top of the second quartile (27.9%). More sophisticated statistical methods can also be used to test the significance of the outperformance, given that the simulation approach creates a distribution of outcomes.

So far, we have developed a simulation that provides a peer-group measurement where none existed previously, but by definition, we have no way to test its validity because there is no



QUARTILE TABLE	
High portfolio	47.5%
First quartile through second-quartile break	25.9%
Median	21.3%
Third quartile through fourth-quartile break	16.8%
Low portfolio	0.2%

Data as of 31 December 2016.



QUARTILE TABLE	
High portfolio	81.6%
First quartile through second-quartile break	27.9%
Median	18.9%
Third quartile through fourth-quartile break	10.6%
Low portfolio	-17.9%

Data as of 31 December 2016.

suitable peer group or benchmark. To construct a real-world test to evaluate validity, we needed to test the approach against a valid and existing peer group that also has a suitable index.

For this example, we selected the US equity market (the largest and most competitive in the world), and within the market, we used the mid-cap segment as one that is reasonably homogeneous. We then simulated a US mid-cap equity peer group in order to compare it to an actual peer group. We did this for each calendar year over the period 2013–2016. The actual peer group used was InvestWorks' US Mid-Cap Equity, which uses the S&P Mid-Cap 400 Index as its benchmark.

The data from 2013–2015 show that, while there were some variances, the simulated universe was sufficiently

close to the actual median to provide validation of the approach.

However, examination of the 2016 data raises an important issue for investors, whether they are using a simulated or an actual universe. Note the large gap—over 8 percentage points—between the S&P Mid-Cap 400 and both the actual and simulated medians. One test for the validity of any peer group is whether the median differs significantly from the index; in 2016, both the actual and simulated results were far from the S&P Mid-Cap 400 return. Note that the broader Russell Mid-Cap Index was up 13.8% in 2016, more in line with the universe's medians.

These findings illustrate another important application for simulated peer groups. The key issue for clients is to decide whether their manager

should be focused primarily on beating the index or on stock selection from a broad universe. If it's the former, the manager takes a risk in straying outside the index constituents, as illustrated in 2016; if it's the latter, having a simulated peer group universe can provide understanding and validation for both the manager and the client when the index does not provide a good representation of the broad universe results. (Note that when we ran the 2016 simulation and restricted it to only the constituents of the S&P Mid-Cap 400, the median of the simulation universe was 21.0%, compared to the index's 20.7%, providing additional validation for the simulation approach when it is constrained to match an index.) This emphasizes the utility of this approach for performance assessment when there is neither a suitable peer group nor a suitable index.

In conclusion, we hope that the approach we have described provides some assistance to those in the investment community struggling with the same problem we faced: how to judge a manager's performance in the absence of a suitable peer group or index.

Antti Raappana, CFA, and Kimmo Kurki are Helsinki, Finland-based investment professionals at Danske Capital. Fernand Schoppig and Barry Gillman, CFA, are investment consultants with FS Associates in the United States.

FIGURE 3:

Test of Simulated vs. Actual Data, US Mid-Cap Equity Peer Group

Year	Simulation Median	Actual Median	S&P Mid-Cap 400
2016	12.3%	12.1%	20.7%
2015	-4.4%	-1.2%	-2.2%
2014	10.5%	9.8%	9.7%
2013	33.9%	36.4%	33.5%

Sources: Authors' simulations, Standard & Poor's, and InvestWorks, as of 31 December 2016.

Dividend Taxation and the Free-Cash-Flow Theory

By Patrick Cusatis, CFA

The taxation of dividends has changed over the years, especially since 2003. Dividends were taxed as ordinary income until the Jobs and Growth Tax Relief Reconciliation Act of 2003 lowered the tax rate to 15% for the top four tax brackets and to 5% for the bottom two tax brackets. This was a temporary tax relief; the prior tax rates were scheduled to be reinstated on 1 January 2009. However, later tax acts extended the tax relief, and dividends are currently taxed

at 15% for the middle four tax brackets and 20% for the highest tax bracket. These tax rates apply only to qualified dividends, which include most regular dividends; other ordinary dividends, such as those paid on real estate investment trusts and master limited partnerships, are still taxed as ordinary income.

Dividend taxes faced by individual investors in the United States are among the highest in the world. According to a 2015 report by the Tax Foundation,

the United States has the ninth highest marginal dividend tax rate of the 34 OECD countries. A 2015 report by Ernst and Young combines the effects of corporate taxes and personal taxes into an integrated tax rate. This report finds that the integrated dividend tax rate in the United States is the second highest among OECD and BRIC countries, behind only France. Both reports show an international tendency to increase dividend tax rates in recent years. Some