

C A M B R I D G E A S S O C I A T E S L L C

DEMYSTIFYING QUANTITATIVE INVESTING

2006

Catherine Piez
Paul Aleckson

Copyright © 2006 by Cambridge Associates LLC. All rights reserved.

This report may not be displayed, reproduced, distributed, transmitted, or used to create derivative works in any form, in whole or in portion, by any means, without written permission from Cambridge Associates LLC ("CA"). Copying of this publication is a violation of federal copyright laws (17 U.S.C. 101 et seq.). Violators of this copyright may be subject to liability for substantial monetary damages. The information and material published in this report are confidential and non-transferable. This means that authorized members may not disclose any information or material derived from this report to third parties, or use information or material from this report, without prior written authorization. An authorized member may disclose information or material from this report to its staff, trustees, or Investment Committee with the understanding that these individuals will treat it confidentially. Additionally, information from this report may be disclosed if disclosure is required by law or court order, but members are required to provide notice to CA reasonably in advance of such disclosure. This report is provided for informational purposes only. It is not intended to constitute an offer of securities of any of the issuers that are described in the report. This report is provided only to persons that CA believes to be "Accredited Investors" as that term is defined in Regulation D under the Securities Act of 1933. When applicable, investors should completely review all Fund offering materials before considering an investment. No part of this report is intended as a recommendation of any firm or any security. Factual information contained herein about investment firms and their returns which has not been independently verified has generally been collected from the firms themselves through the mail. CA can neither assure nor accept responsibility for accuracy, but substantial legal liability may apply to misrepresentations of results delivered through the mail. The CA manager universe statistics, including medians, are derived from CA's proprietary database covering investment managers. These universe statistics and rankings exclude managers that exclude cash from their reported total returns, and for calculations including any years from 1998 to the present, those managers with less than \$50 million in product assets. Returns for inactive (discontinued) managers are included if performance is available for the entire period measured. Performance results are generally gross of investment management fees. CA does not necessarily endorse or recommend the managers in this universe.

Cambridge Associates LLC is a Massachusetts limited liability company headquartered in Boston, MA with branch offices in Arlington, VA; Dallas, TX; and Menlo Park, CA. Cambridge Associates Limited is a Massachusetts limited liability company headquartered in Boston, MA and registered in England and Wales (No. FC022523, Branch No. BR005540). Cambridge Associates Limited also is registered to conduct business in Sydney, Australia (ARBD 109 366 654). Cambridge Associates Asia Pte Ltd is a Singapore corporation (Registration No. 200101063G).

C O N T E N T S

Abstract	1
Summary	4
Exhibits	
1 Five-Year Rolling Information Ratio Versus S&P 500	12
2 Annual Value Added Performance Versus S&P 500	13
Appendix	
A Note on Taxable Portfolios	15

ABSTRACT

1. Quantitative (quant) investment practices add important discipline and breadth of analysis to the investment process. Whether used as an aid in a process that is predominantly qualitative and manual, or as the whole body and frame of a mathematical, model-driven process where final decisions are determined by the computer program, quant methods add value by increasing the efficiency of the investment process.
2. Discipline and breadth are important advantages. Automation alone affords faster analyses and lower staffing costs, which managers can pass immediately on to clients in the form of lower fees. Additionally, analytical processes that are even a little bit successful (say, winning 52% of the time) end up adding meaningful value over time because the computer does not falter (or get bored or distracted by something else) in its analysis of hundreds or thousands of (often neglected) investment opportunities. When a consistent definition of risk is also automated into the portfolio construction process, the result is risk-adjusted returns that are even better. Quant processes are also value-adding in taxable portfolios because the potential tax cost is easily (and cheaply) incorporated into an optimization routine that balances costs with risks and returns.
3. Quant processes are rightfully associated with diversified, core portfolios where the advantages of speed and breadth are obvious. The investment disciplines, however, are not rocket science—they are based on the same basic principles that “fundamental” managers use. For this reason (and perhaps contrary to popular belief), quant processes have the most value-adding potential in less efficient markets where the gap between fundamental values and price setting is greatest.
4. Optimization problems are defined by their input variables, constraints, and objectives. Quant models are mathematical descriptions of how the markets work. Programmers set explicit limits on allowable risks and then ask the computer to calculate the various possible portfolios representing the optimal trade-offs among risks, costs, and returns. Quants bring their own definitions and tolerance for risk to the table. Their philosophies and objectives range from prioritizing a stable, low tracking error on the portfolio to maximizing the information ratio on individual investments—and everything in between.
5. For quants, factor definitions, factor weights, and time horizons are essential to the value-adding process. Almost all firms use a balance of security valuation, company fundamentals, and market technical analysis, but the particular (and personal) choices of each firm directly affect the portfolio performance.
6. Since investment markets are inherently unpredictable, no quant model can anticipate the unknown unknowns. For the same reason, the best quant processes incorporate a learning function so that they adapt as old factors are arbitrated away and new ones are invented.

7. The explicit work that quant managers go through in order to tell the computers what to do enables them to explain clearly to clients what portfolio changes and results to expect as the investment markets evolve.
8. Like all managers, quants want to please their clients. It is therefore important for investors to express their own risk tolerances and objectives as explicitly as possible, so that those goals can be programmed in and prioritized.
9. For fundamental managers, online data screens are a minimal requirement to keep up with the vast and detailed information flow. Firms that add computer-aided ranking and analytical routines to their daily work avoid the investment traps created by personal biases and tendencies to fall in love with old ideas. Fundamentalists are also starting to use portfolio attribution reports as well as BARRA-style risk analytics to impose discipline in the portfolio construction process.

SUMMARY

Introduction

Despite its relatively long history and relatively strong performance record, many investors still suffer from an aversion to quantitative (quant) investment management (Exhibits 1 and 2). In fact, much of the most creative innovations of recent years have come from managers employing increasingly sophisticated quant analytics and strategies. Moreover, since *all* equity managers now use quant analyses to a greater or lesser degree, it has become increasingly difficult to distinguish “quants” from “fundamentalists.” On one hand, quants are crunching more fundamental financial data faster (and often in greater depth) than the analysts who manually read the printed reports; on the other, fundamentalists are using screens and portfolio analytics to automate their daily routines. The first section of this paper will note the ways in which automation has seeped into all managers’ processes. Then we will focus on the quants who program their computers to make investment decisions. Our hope is to demystify their investment methods and thus induce investors to include them more often in manager searches.

Screens, Disciplined Processes and Portfolio Analytics

Screens

Screens are ubiquitous and used by virtually all equity managers to find and prioritize ideas. Factset, Baseline, Advent, and Thomson—all the programs portfolio managers use to keep track of client holdings—can be employed to screen masses of market and security-specific data. In addition, managers have access to multiple sources of online portfolio and market analytics, such as Bloomberg, Zacks, ValueLine, Reuters, Ford’s, and Ned Davis.

Just how much value can be added by the data feeds is questionable. Since everyone now has access to the same information at the same moment, the difference between winners and losers comes from how they interpret and how quickly they act on that information. We cannot say whether unskilled managers stay in business longer because the large information set creates more opportunities for luck, or shorter because they are fatally distracted by the noise. However, we do know that as the volume of information available for screening and off-the-shelf analysis increases daily, as does the speed with which the markets assimilate both the news and the ways it might be interpreted and acted upon, market efficiency rises.

Disciplined Processes

Many equity firms construct their investment process on the basis of a specific set of variables they derive from screens. They monitor idea flow, market average statistics, how the ideas and statistics both change, and then execute mechanically according to a pre-established preference for various investment opportunities. Despite their simplistic appearance (like a paint-by-numbers set), these processes avoid the traps created by managers’ biases and tendencies to fall in love with old ideas.

Portfolio Analytics

Increasingly, managers are using performance attribution reports and portfolio risk analytics as a feedback loop in their (otherwise manual) portfolio construction. The benefit is the revelation (before it is too late) of any unintended bets. However, we are skeptical when managers say their purpose is to ensure that they keep their sector or factor bets low in order to get their value added from stock selection. The exercise of maintaining appearances in no way assures that the remaining stock selection generates outperformance. In these cases, we ask more questions about alpha generation.

Factors, Weights and Time Horizons

Factor Selection

In the early days of computer-driven investing, valuation factors were the only game in town, leading virtually all quants to have a value bias. When that proved a liability in the late 1990s, everyone rewrote their factors. Nowadays, we often see factors labeled into valuation, growth and/or market sentiment segments of the model. Only some quants will reveal their factor definitions; most regard these as proprietary (i.e., easily copied) alpha sources and will only describe generic factor groups. Some firms select different factors for each sector, while others do not. The groups often include (1) absolute or relative assessments of security prices; (2) company financial statement data; (3) security trading statistics; and/or (4) sell-side analyst opinions. However, there are no standards for describing any given firm's stable of factors, and we often hear about factors that do not fit any of these types. For example, some firms have started telling us lately about corporate governance factors (all of which would have de-selected Enron).

Data Mining. Since data mining is a pejorative term, no manager will admit they deliberately select those factors that performed best in the past. On the contrary, most managers make a point of stating that they require alpha generation factors to make theoretical, academic, and/or economic sense. As outsiders who are not privy to the data and who have no intention of interfering with managers' research and development, it takes us a few years of observation to suspect that a firm's quant research team is floundering around for new alpha sources and rewriting the models from scratch every time they get concerned about underperformance. The truth is that it makes sense to use the factors that have won in the past. Who wants to overweight underperforming factors? Try selling that approach to prospective clients! In practice, the development of new, value-adding factors requires striking a balance between those that *have in fact* worked in the past and those that *should* work in the future.

In attempting to strike this balance, firms tend to follow one of two courses. The introverted lean toward the personal investment preferences and biases of the firm's leaders, essentially codifying a preferred investment philosophy in mathematical models. In these cases, the factor definitions are defined as hypotheses first and then tested and refined to reveal what operational vulnerabilities need to be avoided. The extroverted, by contrast, cast about for factor ideas from many sources, intent on harnessing the latest (perhaps undiscovered) alpha ideas. Common sources are quant investment journals and brainstorming

sessions among portfolio managers (although we once interviewed a manager who got an idea from reading an article in *Science* magazine that had nothing to do money). At the extremes, introverted developers focus on a few factors with significant explanatory power and edit the small details incrementally, while extroverted developers try lots of ideas and test them quickly, always seeking any possible new basis points of alpha. As with individual personalities, of course, no firm is entirely introverted or extroverted, only disposed toward one pole or the other. And in both instances, capable managers monitor the ongoing performance power of their factors, watching carefully for signs of decay.

Factor Weighting

Factor weights are as important as factor definitions. What are the factor weights? How are they determined? How do they change over time? The answers to these questions indicate what performance biases clients might expect. Some firms equal-weight key factors. Others overweight those that have produced the highest returns historically. Still others dynamically overweight factors that outperformed in the most recent years or recent weeks. At least one firm weights factors in a top-down manner based on the portfolio managers' assessment of the market environment. Even when the final decisions are based on decades of research data, the ultimate choice of a weighting methodology is just as subjective as a traditional managers' choice of a stock.

The Achilles heel of most quant managers is their tendency to change their portfolio exposures a little too late because of the difficulty in determining whether the performance power of a particular factor is temporarily eclipsed or terminally depreciating. If the market environment is turbulent with increasing price volatility and perhaps quick changes in direction, even the short-term, dynamically weighted processes can get whipsawed when earnings revisions are late and price momentum flip-flops. Investors should therefore attempt to understand how quickly a firm's factor-weighting methodology can recover from a period of underperformance.

Time Horizons

A firm's weighting methodology should be closely related to the time horizons required by their alpha factors. Price reversal, for example, is useful as a return factor for only a few weeks; on the other hand, business innovation and reform, or disparities between asset and book values, are long-horizon factors—so long that we associate these with managers heavily reliant on individual-company-level research. Essentially, managers will better avoid losses if they can remove unsuccessful short-term bets quickly and let still-incubating, long-term bets hold. Similarly, investors can select quant processes more appropriately if they are clear about the duration of their own risk tolerance.

Portfolio Optimizers

Key to optimization problems are the definition of the input variables, constraints, and objectives. Programmers create mathematical descriptions of how they think markets work, then set limits of allowable risks, and calculate the various possible solutions that make the optimal trade-off between risks, costs, and returns. That is, they come up with a suggested portfolio containing the optimal mix of security positions.

Originally, optimization techniques were used by mathematically inclined managers as a way to design portfolios by exploiting the computer's discipline, objectivity, and speed. However, in the early days, they spent most of their time cleaning data inputs that were rife with miniscule, security-specific errors. In short, garbage in, garbage out.

Since those early days, the publicly available datasets are generally cleaner. Quant managers tell us they still have to be vigilant about flaws, but the error rates are so much improved that most users find the basic data feeds reliable enough. In addition, optimization programs have become available to users as off-the-shelf software packages. The first and still one of the best-known risk measurement and optimization vendors, BARRA, opened shop in 1975. Today, research services and academics increasingly use these and their own quant models just to observe the market's behavior; as a result, the concept of factors, like earnings surprise and price momentum, has become widespread because these individual market/model relationships are so easily measured and discussed. Finally, fundamental managers—that is, those who pick securities and design portfolios manually, based on their personal understanding of the marketplace—increasingly use both factor models and tools like BARRA to define their investment processes and to test their portfolio ideas.

Thus, when it comes to screening and portfolio optimization, there is less difference than there used to be between quant and fundamental managers. However, we would suggest that quant managers still have the advantage of breadth. Computers can scan and compute over the whole universe, applying the same analyses simultaneously to thousands of securities, whereas fundamental managers still have to pore over company-specific information on the names their screens have identified. And breadth is an important variable in investment theory: Grinold & Kahn's Fundamental Law of Active Management stipulates that investment return is the product of predictive skill and breadth of opportunity. With enough opportunities, any skill level even marginally better than 50/50 will result in consistent profits. The counter-argument (which we seldom hear anyone make) is that fundamental researchers do not need a broad opportunity set because their predictive skill is heightened by their analytical depth. We agree that satellite managers capable of delivering higher—but less consistent—returns, still have an important role to play in a portfolio; our skepticism is therefore concentrated on core managers whose process lacks breadth.

Risk Constraints, Objectives and the Humanity of Quants

Quant managers' approaches to risk are not all the same. Risks can be defined by the same factors as returns, or by a different set of factors. Some firms use BARRA or Northfield or Investment Technology

Group, Inc. entirely, accepting the risk factors defined in those programs, and then fitting their return forecasts into the off-the-shelf procedures. Some use only the covariance matrices from these programs (i.e., the database of historical correlations between securities, industries, and factors), allowing them to reconfigure the off-the-shelf information set to their own purposes. Some use only their own databases and programs. Some separate alpha forecasts from risk and trading cost factors and then use optimizers to balance the two. Others incorporate risks and costs seamlessly into net-alpha forecasts and build the portfolio separately, perhaps even without an optimizer. None of these approaches is inherently good or bad—the color of the cat is irrelevant since all we need to know is whether it catches mice.

The important question regarding risk constraint is “how much?” Managers tend to say that they do not have to constrain all their risk parameters because the limits they place in one area means that the other limits are seldom reached. For example, if a manager keeps capitalization and all the BARRA risk factors close to benchmark neutral, the computer may seldom suggest that any security should be more than the +/- 3% relative to its benchmark weight. Having said that, however, there appears to be no consensus among managers about which risk parameters are most important. Some prefer to control security weights instead of factor exposures, others the opposite.

In differentiating among quant managers, investors should focus principally on what each manager is trying to achieve and only secondarily on how they intend to achieve it. The designers’ philosophies and priorities are literally encoded into the computer program and then mechanically executed without hesitation or further consideration. Optimizers combine variables and constraints in the service of an objective, which may differ considerably from firm to firm. For example, we know five experienced and well-established quant firms that prioritize foremost (1) the information ratio of each portfolio; (2) firm-wide total excess return; (3) low turnover; (4) predicted tracking error; or (5) fast quant development to capture new market opportunities. We have asked each firm why they selected these particular objectives, and in each case we were told “it’s what we believe in” and “it’s what our clients care about.” In other words, at the end of the day, the programmer is more important than the program—what is inside the black box is what someone has put there.

Caveats

Quant investing requires forecasts that in turn require a predefined understanding of how markets work. Since markets are organic rather than mechanistic—in other words, their shapes and patterns change as they grow (or shrink)—attempts to capture their essential characteristics in mathematical models that will enhance forecasting capabilities are fraught with risk. The difficulty in identifying reliable constants to feed into computer models is called “estimation error.” (But it is perhaps worth noting that the crystal ball that preceded the computer was equally—or more—subject to such risk.)

One obvious source of estimation error is the naïve assumption that the full range of potential risk, returns, and correlations is defined by past statistical patterns. Quant models—and those who feed them—have difficulty accounting for extreme events (which occur more often than statistical models predict) and

for patterns that have never been seen before (e.g., Mr. Greenspan’s yield-curve “conundrum”). As any insurance agent knows, the most important risks are the unexpected—and are qualitatively very different from statistical probabilities. With his customary flair, Peter Bernstein has vividly captured this distinction by differentiating between the risk that is captured in a probability distribution and the unquantifiable “wildness” that lurks beyond. Similarly, the legendary risk officer at NASA asks, “what unknown unknowns are still out there?” Answering that question, of course, is where the creativity and experience of the people who design the models comes in.

An additional source of imperfection stems from our own habit of latching on to one particular quant concept and beating it into the ground, losing sight of its limitations and thereby inadvertently signaling inaccurate messages to our managers about our objectives. Predicted tracking error, for example, is not on its own particularly instructive—despite being characterized in many quant manager reports as “total active risk.” The information ratio is a far more useful metric, since it shows the outperformance (value added) created for each unit of tracking error. If Manager A consistently generates a large, positive information ratio, but with greater tracking error than Manager B, whose information ratio is substantially less, a numerate investor should prefer A to B.

The fact that markets learn while successful investment methods are copied and arbitrated away necessitates a constant search for new, effective models. Moreover, as the learning process accelerates, quants have to imbed feedback loops into their development processes so they can respond to factor decay more quickly. Fortunately, although this problem is chronic, it is not terminal. While the windows of old investment opportunities shorten, they still persist long enough for new opportunities to be found. Investors who know their managers well and are able to listen and participate in the discussions of new developments can stay up to date and ensure that their ongoing interests are being considered.

Looking Forward: New Developments

Quant strategies, particularly those based on factor screening, are increasingly evident in the retail mainstream. ValueLine, which is essentially factor based, has of course occupied this space for decades, but now investors can select exchange-traded funds based on single factors, like dividend yield. Products based on risk (rather than simply return) factors are surely not far behind. The recent development of new “indices” (are they truly indices?) like Rob Arnott’s fundamentally-weighted constructs has encouraged quant firms to advertise their ability to create similar products in accordance with client specification, either by setting up the factors, weights, and rebalancing procedures directly or by mimicking the same objective with stratified sampling and other well-tested indexing techniques. This is all part of a broader industry trend to start thinking of portfolio construction in terms of *sources* of risk and return, rather than in asset allocation terms, and includes a steady increase in the quality and quantity of discussion about factor effects in market return streams, beyond sector and capitalization effects.

Quant managers are also heavily involved in the development of products that expand the investment opportunity set to include short selling. Since we discuss portable alpha and 130/30 (or similar)

portfolios in other reports, we will limit our comments here to noting that it is very much an open question whether most traditional quant firms (as opposed to quant hedge funds) have the capabilities on hand to implement and manage short-selling effectively, or whether a significant increase in the dollar volume of short-selling is likely to have unforeseen consequences (e.g., liquidity crises) if indeed these long-short constructs attract significant investment (as they seem already to be doing).


Conclusion

With advances in computing power and sophistication, and instant access to reams of increasingly reliable market and security-specific data, quant tools enable managers to review a broader opportunity set at a faster pace and make instant decisions while enforcing rational portfolio construction disciplines. All this contributes, of course, to greater market efficiency, which means that (perhaps contrary to popular belief) highly quant approaches actually work best—have the most value-adding potential—in less efficient markets where the gap between fundamental value and price setting is greatest.

However, most investors may find that core U.S. equity allocations are the best place to start developing experience with quant products, although experienced quant firms also offer small-cap and international mandates and are quickly gaining capabilities at higher risk and return levels. As with fundamental managers, the most important traits to look for are experience and aptitude for the task at hand and the ability to communicate and substantiate investment expectations.

EXHIBITS


Exhibit 1
FIVE-YEAR ROLLING INFORMATION RATIO VERSUS S&P 500
1993-2006



Source: Cambridge Associates LLC Investment Manager Database.

Notes: Calculations are based on quarterly data. Manger returns are gross of fees. Percentile rankings are based on a scale of 0-100 where 0 represents the highest value and 100 the lowest.

Exhibit 2
ANNUAL VALUE ADDED PERFORMANCE VERSUS S&P 500
1988-2006



Source: Cambridge Associates LLC Investment Manager Database.

Notes: Calculations are based on quarterly data. Manger returns are gross of fees. Percentile rankings are based on a scale of 0-100 where 0 represents the highest value and 100 the lowest.

APPENDIX

Appendix

A NOTE ON TAXABLE PORTFOLIOS

For every security in a portfolio, the potential taxes payable if the position were to be sold today can be calculated accurately based on the type of security, the applicable tax rate, the date of purchase, and the cost basis. The potential tax payment, then, is simply incorporated into the portfolio optimization as a cost that needs to be controlled, along with transaction costs, risk constraints, and return objectives. Investment Technology Group's optimizer seems to be the current one favored by managers off the shelf. It keeps tax-lot specific data, automatically avoids taxable gains, and highlights opportunities for active tax-loss harvesting. Taxable quant managers all emphasize their client service skills in order to start the client relationships on good terms. In truth, this is not just because it is good business practice, but also because the clients' deposits and withdrawals require purchases and sales from the portfolio that can significantly facilitate or hinder the manager's ability to minimize the tax bill.

Taxable investors should be aware that potential tax costs are part of a larger equation that includes transaction costs, total risks, and investment returns. If the manager can show that a small tax bill accompanies a large investment profit, we may not begrudge the taxman his payment as much as otherwise.

Conversations with taxable index managers on this topic focus on the tax cost as a percentage of the benchmark return versus the tracking error against the benchmark. For an index manager, of course, the tax cost is recovered by incurring small benchmark-relative risks: using substitute securities in lieu of the precise weights of the benchmark holdings, or trading the benchmark holdings on a different day when the investment profit is higher than the bench and/or the tax cost is lower. Taxable indexers have told us that 100 basis points (bps) to 200 bps of tracking error is essentially enough to let them maximize the tax efficiency of the portfolio.

Active quant managers, however, both at enhanced index and higher risk levels, want to outperform the benchmark, not just to match it as the indexers do. Interestingly, when we ask active quants how much tracking error and investment returns they have to give up to maximize tax efficiency, they like to raise the concept of tax *lambda* (λ). Rather like the tension knob on a sewing machine, *lambda* is a sort of dial that the portfolio manager uses to mathematically specify the trade-off preference for minimal (investment plus future tax) risk or maximal (investment minus tax cost) return. Setting *lambda* toward lower risk affects the magnitude of possible return outcomes, and vice versa. It is a concept that managers use throughout their quant models but it comes up specifically in taxable portfolios because managers can address clients' strong feelings about taxes by tactically adjusting the trade-off to accommodate client preferences or their own judgment of when incurring a larger tax bill might in fact make investment sense in after-tax terms.