For professional investors only

… Graphene Investments does not believe in pure Quant

Warning: Just because this article deals with controversial topics doesn't mean we are seeking controversy. We do not feel obliged to systematically align to the mainstream, one-track view when we consider that we have strong points in favor of a different stance, and our only objective here is to explain the reasoning behind certain choices we made when designing our project and our product line.

Many who have been following us over the years know it: we, at Graphene Investments, use Quant, like Quant, and trust Quant … but only for certain tasks. No decision is ever made without a deep review of every investment idea through judgmental analysis. Why don't we go further, and let the machine do more on our behalf?

While quantitative techniques have been around for a long time in portfolio management, they long remained focused on specific parts of the investment process, such as preliminary screenings or risk control. Aside from passive managers, few investment professionals would rely entirely on computers… and even fewer would advertise that reliance in their marketing, because customers were not prepared to trust a fully automated "black box" approach. In the last few years, though, things appear to have changed. Increased awareness about AI and its (probably overstated) future capabilities is making computing popular. The man in the street (or at least the younger one) is now prepared to entrust his life to a computer-driven vehicle, and his money to a robo-advisor. At Graphene Investments, we have always been using a balanced mix of quantitative and judgmental analysis, and we intend to continue to do so, because we believe pure Quant has major shortcomings that will persist. Since transparency is one of our core values, we tell you why…

Garbage in, garbage out

The first point, which limits the potential of pure Quant, is the quality, the depth and the comprehensiveness of available data. Not only do most databases contain a lot of mistakes or missing pieces but, even when they focus on information that is fully quantifiable, they only give a superficial view of that information. While the most commonly used numbers are widely available with a decent reliability, the problem grows exponentially when one tries to dig into more specific details, or cover companies beyond the pure "large cap" universe in mature markets.

It is usually admitted that quantitative models' main advantage is their ability to process, swiftly and reliably, a huge mass of data, and extract the most valuable information from it. A basic, cheap computer will do in seconds what a whole group of human beings would do in hours or days. However, the result will be useful only if the

original data was clean and reliable, and this is far from guaranteed in most publicly-available sources.

Feeding a database is not the most gratifying job one can dream of, and this work is usually entrusted to interns or junior staff, who cannot reasonably be expected to fully understand the meaning of the numbers they process, nor the importance of their accuracy for further investment steps. In recent years, input processes have even been increasingly automated through data collection software, with questionable results. Anyone who spent a few minutes on a database knows that it doesn't even take a corporate action, an accounting change or a corporate tax reform to cause misleading variations in stock data.

The increased use of "Off-balance sheet" items also contributes to make things trickier to interpret, and requires increased disclosures that can't always be brought down to code that a computer will understand. The fact that many of the latest major market crashes somewhat originated in off-balance sheet liabilities should be kept in mind.

In the absence of human beings with a deep market knowledge at any stage of the "pure Quant" process, one cannot expect the issues to be detected easily. It is possible to screen data for anomalies and inconsistencies, but not all mistakes can be pointed out this way, because a given number can be completely wrong without being obviously mistaken. For example, depending on field formats, 15% may also be shown as 0.15, and a 15% return-on-equity will easily, but wrongly, be confused with a 0.15% ROE. This example may look grossly exaggerated, but it actually highlights a very common type of issue in many sources.

To make things worse, a systematic approach makes the process much more sensitive and vulnerable to wrong data. People often think that, if the percentage of mistakes in the database can be kept low enough, the impact on the quant model's reliability will be statistically insignificant. What works for a survey unfortunately turns out to be dramatically wrong in asset management because the model will systematically point out and misuse the wrong numbers first. Let's assume for example that, within a universe of one thousand companies, only ten have a mistaken valuation, which makes them look ten times as cheap as they really are. A human portfolio manager has a 1% probability to bump into one of these as a Buy candidate at that precise moment, so that the errors will most likely remain unnoticed and impactless. Moreover, even if he happens to review one of these stocks, he will immediately notice the anomaly, and drop or adjust the case. On contrary, if stock selection is handled by a fully quantitative process, the chance is that all ten companies will end up in the portfolio, for being the cheapest in the whole universe. Similarly, if a stock is abnormally valued by the market (for reasons such as a high bankruptcy risk, or another major issue to which the usual metrics don't react properly), it will be picked by the system.

This is, by the way, the reason why investors with a knowledge of markets and securities often find the result "weird" or "scary" when they review a Quant-based portfolio with their own, judgmental criteria. The fact is, such characteristics don’t always lead to a disaster, but when they do (which statistically has to happen one day or another), it usually turns out to be extremely painful.

Not everything can be reflected in numbers

The next issue is that many of the criteria that a judgmental analyst will take into account when reviewing an investment idea are based on information that cannot be quantified, and that, therefore, will be ignored by a pure Quant approach. Admittedly, some of these criteria are irrelevant or too emotional to be reliable, and the machine will benefit from not using them. Some others, however, could be useful.

The most obvious case is that of discretionary opinions, such as the assessment of a corporate management's consistency in execution quality, or the analysis of the tone

during an investor day. While all such impressions are not worth following, there will be times when they lead to a strong "gut feeling" which may be a useful sign that the computer will not capture. In some cases, non-quantifiable data may become much more important than reported numbers. In a bank, for example, reputation is key and a negative rumor will kill the business and the stock's price even if numbers prove that everything is fine.

Even in a more rational area, some details will never be in a database because they cannot be easily formatted for that purpose. While reported data is usually in a relatively standard format, there is a lot more freedom for forecasts, whether they reflect companies' guidance or analysts' estimates. Each source has its own assumptions and scenarios behind the headline numbers it publishes. The details are not always disclosed and, even if they are, assumptions only highlight the key factors of the moment, which vary over time.

Users of earnings estimates databases are familiar with the problem. Due to mergers, accounting changes and many other events that distort numbers, it has become increasingly difficult to compare analysts views on future earnings, even on a same company and at a single point in time. The time at which they integrate the event in their forecasts may vary greatly from one contributor to another. Consensus forecast data providers tend to get around the problem by adding footnotes to explain every piece of data, and by excluding certain sources from their consensus calculations to focus on comparable numbers. The method, which has occasionally been criticized for lacking transparency, is useful but it contributes to make the data less understandable by a computer.

Even analysts' target prices, which technically are fully quantifiable, actually prove difficult to use. They don't mean anything without a time horizon and a general scenario regarding the investment backdrop, and these are not published in a formal manner. The only way numbers can be reconciled with the assumptions behind them is by retrieving small bits of the underlying landscape, through the interpretation of details spread across written reports, additional comments in oral presentations, and possibly the same analyst's coverage of other companies in the same sector.

Here again, the most advanced word recognition systems may theoretically be able to deal with the situation, but the fact is, they will do so with so many misinterpretations that the result will require human verifications anyway.

Artificial intelligence doesn't exist

Agreeing with the points exposed in the above two sections is another way of recognizing that human thinking is able to understand things better than even a sophisticated data processing. Those of us, who have tried to use the so-called "AI systems" that are supposed to help us in daily life (eg spelling assistants in word processing software, voice recognition systems in dictation devices or virtual assistants, or self-driving solutions in certain modern cars) know the situation well. As long as the conditions are fairly undemanding, the result may be OK, but when things go wrong, they go REALLY wrong!

Granted, these systems use simplistic forms of Artificial Intelligence, and should not be compared to the much more powerful solutions, which are currently being run in advanced research labs, and may eventually become common in corporations' IT. When this happens, we may have to revise our view, at least in part, and admit that progress has made pure Quant a safer, more dependable solution than it currently is, but we don't think we will be there any time soon.

For the moment, active developments in AI remain focused on algorithms from the end of the 80's, and consist in learning from big data sets. The exponential growth in the

amount of data generated as digital solutions spread through everyday life provides an obvious raw material. The increased affordability of computing power, which came as a side effect of investments made to develop GPUs (graphic chips) for video and gaming purposes, brings a cheap tool. Both together simply allow researchers to implement old solutions with better results. However, they remain limited to very specialized areas, such as image / text processing and games. At this stage, despite these impressive improvements, AI remains far from understanding the context, and even farther from learning by itself. Show one or two dog pictures to a child and he will recognize every dog he meets in his whole life. Try to do the same with a computer, and it will take thousands of pictures for the machine to start differentiating a dog from a cat. However interesting and promising AI may be, calling it Artificial Intelligence is probably very excessive, and something like Enhanced Processing would be more accurate, although not as glamorous.

To illustrate the difference it makes for asset management, we will take the example of this technology company, whose four or five successive CEOs repeatedly failed to convince investors over the last decade, to such an extent that the stock price would systematically shoot up when their replacement was announced. At its current stage of development, AI would soon have concluded that, at that company, a CEO resignation was a good thing, and meant that a relief rally was in the pipe for the security. When the latest CEO's departure was announced a few months ago, most computer-driven strategies probably reacted according to this pattern. The problem is, human portfolio managers did exactly the opposite this time, because they (unlike the computer) knew that the latest leader of the company was highly regarded, and that his departure for personal reasons would be a blow to the company's nascent turnaround.

There are plenty of examples like that, where a small, unquantifiable detail will occasionally derail a mechanism, which looked well established. The "pure Quant" believers will probably keep adding more sophistication to their processes, so they can manage more different situations. They may succeed, but the risk is, they might eventually end up with so many different inputs that it will be difficult to get a clear-cut output message.

Quant is not perfect, because it has technical shortcomings that, for the moment, look almost impossible to eliminate or circumvent. We believe that, for many of its proponents, the actual, true reason for using it is they view it as a low-cost solution (which, by the way, is wrong if one cares to reach the highest level of data quality and to run comprehensive development backtests).

Does this mean pure judgmental management is any better? We don't think so, because it has other weaknesses, such as its excessive reliance on intuition, and its lack of discipline and consistency when circumstances become stressful. Under those circumstances, the Quant's advantage is its cold, emotionless focus on facts. Moreover, human judgment cannot be backtested, whereas a Quant strategy can be tested and refined on historical series.

This is exactly why, for active management, we believe in an hybrid process, which takes the best of both worlds: the combination of two unperfect approaches can then lead to something much better than any of the two.

Important information: *The views expressed herein are for information purposes only. They should not be interpreted as a recommendation to adopt or modify an investment stance, or purchase or sell a financial instrument. They reflect Graphene Investments' analysis as of the specific date stated at the top of this document, based on information that was available at that time. Such information, and the resulting opinions and assumptions, are subject to change without notice. Graphene Investments does not guarantee their completeness and accuracy.*

Before making any investment decision, investors should carry out their own analysis, based on up-to-date information, to form a personal opinion about the suitability and risk of that investment.

AMF license #GP-16000022
10 rue La Boétie 75008 Paris (France)
T: +33.1.70.82.44.50
F: +33.1.70.82.44.49
E: contact@graphene-investments.com
W: www.graphene-investments.com