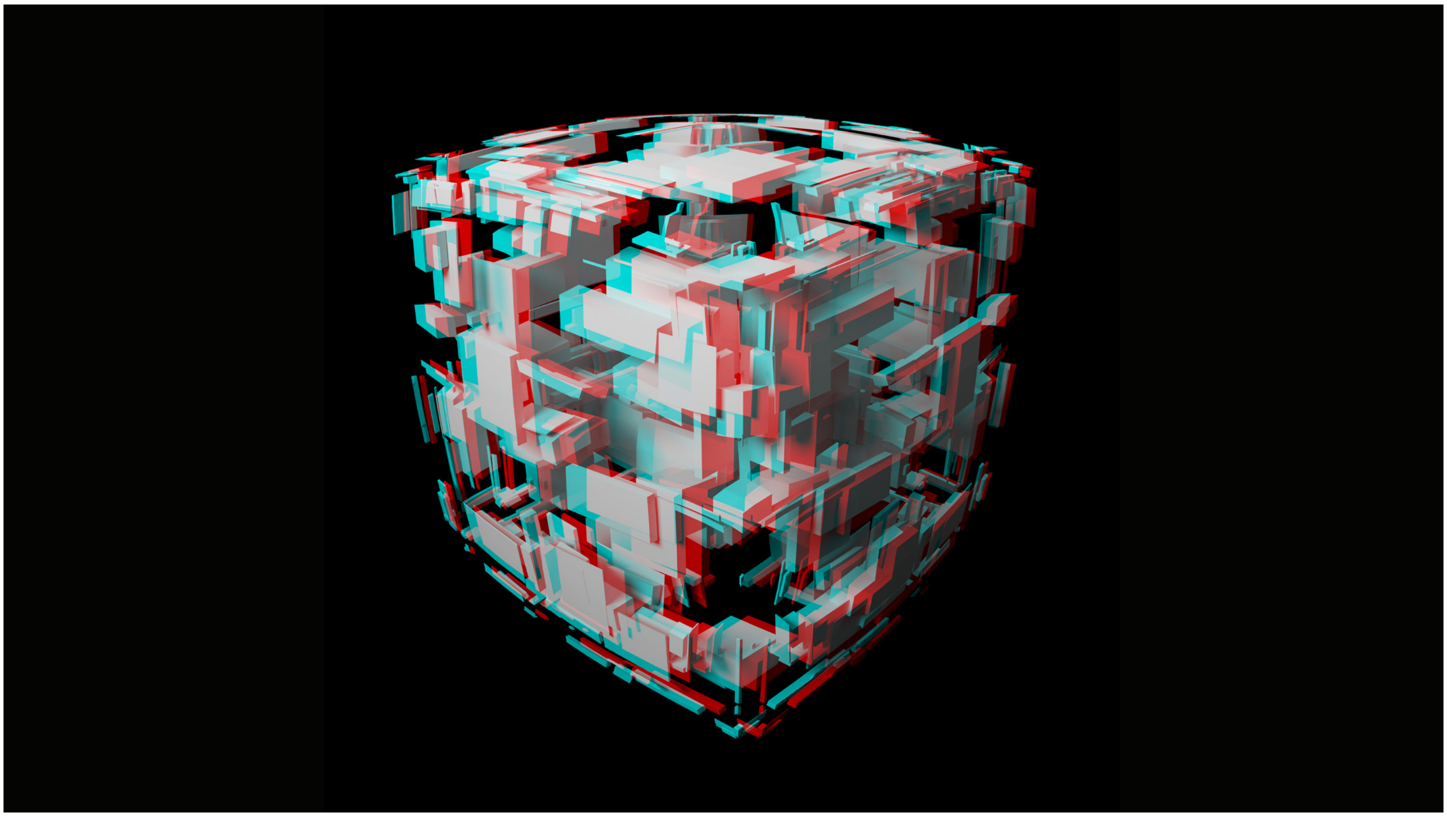

Hypercube in the Kitchen: Reading a Menu of Active Investment Strategies

Igor Yelnik
ADG Capital Management

Boris Gnedenko
ADG Capital Management

Introduction

The importance of skill in active investment management cannot be overestimated.[1] Investors' belief in their managers' skills is the only justification for the existence of the multi-trillion industry. Some skills are unique, i.e. only possessed by one manager. It is a set of independent unique skills that helps a fund manager to deliver long-term outperformance over his benchmark, be it an index or a peer universe. Moreover, unique skills that constitute the firm's investment edge are not easy to migrate from another firm. While more commoditized skills are readily available through the job market, core competences are likely to remain in scarce supply. Therefore, investment skill appears to be the most natural candidate for segregation of types of investment processes and of managers implementing them.

Indeed, classifying fund styles[2] based on skills ensures that respective segmentation is pretty stable: changing style in such coordinates is hard and expensive as it usually means acquiring new skills and only rarely abandoning those not required anymore. Different types of funds already possess well-established classifications.[3] Hedge funds are commonly classified by a strategy type. Though several competing classifications exist, they all closely resemble each other, differing predominantly by depth of granulation.[4]

The skill-based classification we propose below provides an extra dimension for diversification between active investment strategies. It is by no means a substitute for the traditional fund classifications. On the contrary, the two approaches are supposed to complement each other in a similar fashion as industry and style classifications work together in the equity space.

In some way, unique investment skills play the same role for hedge funds as factors do for traded assets. Much like returns of any asset may be attributed to returns of its basic ingredients – factors, one can attribute returns of an arbitrary hedge fund to a mix of its unique skills.

An investor may want to diversify her portfolio by allocating to managers possessing different unique skills. This creates a conceptual link to the Grinold-Kahn's fundamental law of asset management, only applied to skills. Thinking of unique investment skills akin to independent bets in the traditional formulation of the law, one can conjecture that the breadth of a unique skill set is a determinant of a fund manager's performance.

This article is deliberately non-technical and should be regarded as an invitation to further discussion on the subject. In particular, we do not go into quantitative aspects of the problem such as introducing a systematic methodology for measuring unique skills of hedge fund managers. Admittedly, inventing such a methodology for measuring uniqueness of skills is not a problem that only has one solution. While measuring the "unusualness" of a manager may be done by analyzing his correlations with peers or a R-squared from a regression on his benchmark, these and other similar approaches do not allow decomposition of the final product back into ingredients, i.e. individual skills, which are our focus in this paper. Instead of delving into technical details, our aim was to introduce the concept of skill based classification of fund managers and provide an intuitive justification for it.

Properties of a viable fund classification

Classifications built upon skills generally satisfy a number of properties pertinent to a good classification:

- **Stable**- Stability is guaranteed by the funds' need to maintain focus around their major investment edge.
- **Informative**- Funds that leverage on essentially different investment skills are supposed to have distinct performance: their decision-making processes should be rather uncorrelated.
- **Universal**- Unique investment skills developed in one asset class can often be transferred to other asset classes. Note that an absence of asset-class specific implementation skills will not present a hurdle for this universality because such skills are already commoditized to a high degree and should not be treated as unique.
- **Identifiable**- A strategy is identifiable as soon as its major edge is known. A clearly stated investment philosophy and description of investment process are examples of clues a potential investor may use to draw his conclusions about such positioning.
- **Exhaustive**- Each strategy can find its place within such classification.

Skill Scales

Combining several skill scales we would construct a viable skill-based classification. But first, given the fact that we aim at classifying fund managers based on their unique skills, we can list several examples of scales that are unsuitable for such a fund style classification:

- Absolute vs. Relative - based on a benchmark type;
- Long-only vs. Long-short - based on portfolio constraints;
- Leveraged vs. Unleveraged - based on amount of leverage used;
- Offshore vs. Onshore - based on a fund's jurisdiction.

None of the above dichotomies are based on unique skills: one does not need an essentially different investment edge to move along the spectrum of possible strategies for each of the above dimensions.

Skill-Based Classification of Investment Processes

As was already stressed, what we are going to classify is in essence the universe of distinct investment processes. We start from a description of an arbitrary investment process as an information processing system (IPS). An abstract IPS consists of three major parts: Input, Output and Processor in between, see Exhibit 1.[5]

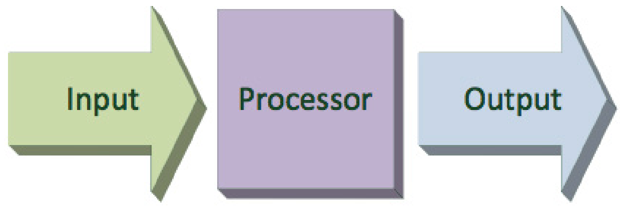

Exhibit 1: Information Processing System

Input is information received by Processor, Output is information produced by Processor based on Input. Processor itself can be imagined as a standard computer processor running a certain set of applications - decision making rules. Such a trivial representation of any informational processing including an investment process can be surprisingly beneficial for our purposes. To be more specific, we are going to associate appropriate scales of investment skills with each of the three IPS elements. These scales are depicted in Exhibit 2.

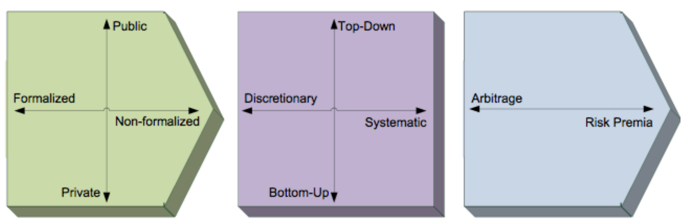

Exhibit 2: Investment skill dimensions projected on IPS parts

As is evident from Exhibit 2, we propose using five scales to classify investment processes: two for Input, two for Processor and one for Output. Below we discuss each of them in detail.

Output

We start with the Output scale because we believe that its role in classifying investment strategies is the most fundamental one. In fact, we would like to present two alternatives for the output scale. While one such scale shown in exhibit 3, is more important from the theoretical standpoint, the other is intended to be more useful in applications.

Alternative 1. Arbitrage vs. Risk Premia Performance Driver

Arbitrage   Hybrid   Risk Premia

Exhibit 3: Output: Performance driver scale

This dimension depicts the nature of major performance drivers of a fund.

The original CAPM only recognizes the market risk factor whose expected return is the market risk premium. According to a typical practical approach building upon this model, a part of a manager's return in excess of the market risk premium is considered to be driven by the manager's skill and is commonly referred to as alpha.

Loosely speaking, our performance driver scale can be seen as a variation of this alpha-beta dichotomy, only brought into a world of many systematic risk factors.

By systematic risk factors we mean a set of uncorrelated portfolios which serve investors as insurance against their bad times (times when consumption growth decreases or, equivalently, utility value of one extra dollar increases). Risk premia are expected returns of systematic risk factors.

We distinguish between two extremes: pure arbitrage (i.e. risk-free) and pure risk premia. An example of the former is a geographical arbitrage, i.e. arbitrage between prices of the same instrument quoted on different exchanges. Such strategies are critically dependent on the technological infrastructure, as they require ultra-fast market access and information processing. In contrast, strategies that only exploit risk premia can exist even in fully efficient rational markets since the existence of risk premia does not premise on any mispricing in assets. Put differently, pure arbitrage strategies exploit market informational inefficiencies on increasingly short time frames, while pure risk premia strategies aim at collecting profits that are left on the table after all available information has already been incorporated in prices. In reality, the performance of risk premia strategies is usually enhanced by persistent heterogeneities among market participants (heterogeneity in utility functions including investment horizons, presence of different types of costs and investment constraints), which can lead to stable market segmentation not easily arbitraged away.[6]

In general, the more one moves to the right along the performance driver scale, the more uncertainty is associated with performance. This is a reflection of the inevitable risks that one has to bear when collecting risk premia as opposed to arbitraging away market imperfections. Importantly, transition from arbitrage strategies to risk premia collection is rather smooth: even such "risk-free" strategies as a geographic arbitrage still bear some risk related to asynchronicity of order fills on two exchanges, FX movement or simply connectivity risk.

Performance driver	**Arbitrage**	**Risk Premia**
Use of market imperfections (inefficiency, irrationality, segmentation)	Based solely on the imperfections	Do not rely on imperfections
Time scale [7]	Short-term	Long-term
Risk	No risk	Systematic risk only
Uncertainty	Certainty: outcome is known	Uncertainty: probabilities are unknown
Critical skills	Recognition of arbitrage opportunities Speed	Factor capture Factor allocation Risk management

Exhibit 4: Characteristics of performance drivers

Exhibit 4 brings together various archetype features of the two performance drivers above.

Though determining exactly the major performance driver for a given strategy is not always easy, to say the least, the above discussion provides two indirect ways to approach this task:

> Risk premia strategies generally spend significantly longer time in trades as they do not aim at getting an informational advantage.[8] As a consequence, the Output scale also provides indirect information about a fund's investment horizon and capacity. The reverse is also true in most cases.
>
> For example, a high frequency trading (HFT) fund is more likely to exploit various degrees of arbitrage than risk premia;[9]
>
> Risk premia strategies usually have an upper limit on their risk-adjusted returns. Information ratios around 2-3 are extremely hard if even possible to accomplish in the long term. Pure arbitrage strategies, in contrast, can reach double-digit ex-post information ratios due to their near "risk-free" nature.[10] However, such strategies have relatively low capacity. Their expected performance is more uncertain because of the costly technology race among competitors. Therefore, such strategies' high historical risk-adjusted returns are less likely to be repeatable.

In theory, arbitrage strategies should be insensitive to bad times. Assuming that some arbitrage strategies are run alongside risk-premium type portfolios (e.g. equity portfolios) the overall allocation of capital to arbitrage strategies may suffer during the general market bad times, thus widening spreads and giving rise to richer arbitrage opportunities. Conversely, it may be expected that good general market conditions may lead to an increased competition between arbitragers and poorer arbitrage opportunities. However, real-life arbitrage strategies where profits are almost guaranteed may suffer during the transition periods of capital exhaustion as spreads widen and segmentation unimagined before becomes reality. Without pretending to explain when and why these bad times come, we only attempt to show that an inevitable exposure to a systemic risk provides another evidence of a risk premium component necessarily present in real-life arbitrage strategies.

A reader must have noticed that we did not mention a manager's alpha as a performance source in its own right. Jarrow and Protter (2013) provide theoretical justification for such omission. They show that in the absence of arbitrage positive alpha is always illusory: it is an artifact of a miss-specified factor model used to obtain alpha, or an incomplete information set.[11]

Hence, what people label as a market anomaly, in fact usually manifests some hidden systematic risk: the only way a positive alpha can be achieved is by exploiting (rare) arbitrage opportunities. This does not mean that active managers should only be compensated for arbitrage profits. Indeed, systematic risk factors are unobservable in the real world, and the respective risk premia are unknown. Moreover, the set of prevalent systematic factors can change with time and respective risk premia can also be dynamic. Therefore, identifying the most essential systematic risk factors and correctly estimating their current risk premia represents a special skill. This skill, crucial to risk premia strategies, is not covered by the notion of alpha but it would be reckless to underestimate its importance when choosing between active managers.

Needless to say, such a view is in sharp contrast with the conventional academic approach, where a risk model driving asset returns is assumed to be fully specified and known. In a common industry parlance, returns in excess of a well-defined and easy-to-capture set of risk factors are often referred to as alpha. In most cases, such a fixed risk model is misspecified, however, it may be very convenient for practical purposes. From a perspective of an investor whose starting point is such model, the difference between pure arbitrage and factors lying outside his model is blurred as both produce alpha, in his terms.

Alternative 2. Diversity of Risk Premia

Single Multiple

Exhibit 5: Output: Risk Premia diversity scale

Looking at the performance driver scale, we notice that the majority of existing funds would be positioned near its right edge. Indeed, as we have already discussed, arbitrage opportunities are rare, so large enough funds usually exploit various risk premia, even if they declare to deliver alpha. This leads to a natural desire for a specific classification of risk premia funds. This can easily be accomplished within our framework by replacing the performance driver scale in Output with a scale distinguishing between risk premia strategies only. Thus we introduce a premia diversity scale that differentiates between single and multiple risk premia exploited within a strategy.

The left edge of this scale is occupied by single premium funds, which invest in one clearly defined systematic risk factor. For example, most CTAs would be located here due to their overwhelming exposure to a momentum factor.[12] Funds that attempt to identify and trade as many systematic risk factors as possible, ideally spanning the whole factor space, would reside on the opposite edge of the scale. Funds trading a multiple, but apparently incomplete set of risk factors would then be located between the two extremes. Global macro funds, both discretionary and systematic, would tend to lie closer to the right edge of the scale.

A skill critical for single premium funds is efficient exploitation of the respective risk factor. In our CTA example this mostly reduces to identifying an optimal definition of such a blurred notion as trend. In this respect, enhanced index funds are likely to belong to this point as well, since they are supposed to capture market premia better than cap-weighted index funds do. In contrast, distinctive skills of a fund positioned on the opposite edge of the scale would be factor allocation and risk management.

Should We Consider Factor Timing as a Unique Skill?

One could wonder whether dynamically managing factor exposure(s) represents a separate skill. At first glance, it may seem that some strategies skillfully provide a time varying exposure to a single risk factor. Consider, for example, a dynamic beta (market timing) strategy that takes a full long position in an equity market except in times of elevated volatility when the position reduces. Whatever it says on the tin, in reality this strategy provides constant exposure to an equity market risk factor (by always maintaining a full long position in the market) and a low volatility factor (by taking a short position in the market from time to time).[13] Thus, a time varying exposure to a risk factor is often equivalent to a combination of constant exposures to several risk factors.

One may reduce a dynamic allocation to a static one by following Brandt and Santa-Clara (2006). Their approach makes use of signals that are supposed to predict risk premia, or expected returns of factor portfolios. They suggest that the set of factors can be expanded by adding pairwise products of original factor portfolios and the predictive signals, thus marrying "factors for assets" to "factors for factors" and reducing the task to solving a conventional static portfolio choice problem.

Hence, a dynamic allocation to risk factors (factor timing) is not necessarily a skill per se because in reality it may be managed as a static allocation to an expanded set of risk factors.

Of course, in reality a set of factors that a manager is able to capture and trade is never exhaustive. Respective risk premia may be dependent on other, more fundamental forces driving investors' risk preferences. For example, an equity market risk premium being time varying may be further decomposed on underlying risk premia. A manager could try to identify these forces and exploit them to dynamically manage factor exposures, or, equivalently, to allocate statically across an expanded set of factors. At the same time, for all practical intents and purposes, this factor expansion process must stop at some level acceptable to a manager or to an investor. Therefore, this level of detail, or the choice of risk premia that one considers the most fundamental, has a strong bearing on whether allocation to the respective risk factors should be treated as static or dynamic.

This leads us to an analogy between traded instruments and risk factors – expected returns of both may consist of more fundamental building blocks. If we accept this view then portfolios of risk factors may be managed based on the same principles as portfolios of assets.

To summarise, identification and selection of risk factors and managing exposures to them by taking positions in tradable financial instruments represents a special skill. Note, however, that even in conjunction with the other edges of our hypercube, the number of risk premia exploited does not provide enough information to guarantee that similarly classified strategies would always display high correlation. Two strategies may use an equal number of risk premia but those premia may be qualitatively different. For example, a single-premium strategy exploiting value will provide very different characteristics when compared with a strategy exploiting momentum.

For any practical use, this scale should be complemented with risk premia actually used in the strategy. We are deliberately avoiding classifying risk premia in this paper: at the current level of perception of various risk premia it would be difficult to offer a non-controversial, theoretically justified and complete classification that would stand a chance of wide adoption by the industry. Even the best researched concepts like an equity risk premium still cause discussions about their interpretation and decomposition. Classification of less researched premia would risk facing a more heated opposition that could divert attention from our point, namely that even though the skill based classification does not aim to provide a complete risk premia classification it helps investor to better understand investment strategies.

We believe that each investor may use his own classification of risk premia to be used alongside our hypercube. Better constructed classifications of risk premia may help the investor to stand out from his competition and represent his own competitive edge.

Input

Input determines the type(s) of information used in investment decision-making. It has two scales that differentiate between formalized vs non-formalized and private vs public information.

Formalized vs. Non-Formalized Information

Formalized Hybrid Non-formalized

Exhibit 6: Input: Information formalization scale

Examples of formalized information are historical prices of financial instruments, fundamental and macro data. Non-formalized information would be mostly represented by news stories presented in various formats. Of course, non-formalized information can often be converted into formalized in many ways, but we leave this job for Processor as it is a part of interpretation of the information.

A hidden assumption behind introducing this scale is that hard-to-formalize data such as news or sentiment data derived from web mining might contain information not fully present in prices or fundamental data. Looking at this phenomenon from a different perspective, we can describe the whole financial market itself as an IPS representing a full set of IPS's active in the market. Such a combined market IPS processes all new information available to market participants into prices of financial instruments. Updates in fundamentals usually arrive at discrete moments, while historical prices reflect previous output of the same IPS, which could decrease their value. What is left, and presents a continuous information flow, is non-formalized information contained in news.

Crucially, gathering and processing these two types of information requires essentially different skills. Formalized information is relatively cheap to access and interpret. However, exactly because of this reason the universe of market participants utilizing it is extremely competitive. On the other hand, non-formalized data is hard to comprehend and apply and if implemented on a large scale, it requires extensive text mining and processing skills.

Private vs. Public Information

Private Hybrid Public

Exhibit 7: Input: Information accessibility scale

The second Input scale distinguishes between private and public information. Public information, in our terms, is information which is acquired relatively cheaply and often comes down to data vendor subscription fees. In contrast, obtaining private information, i.e. information not readily available through public information channels, is often associated with significant ongoing expenses, be they explicit or implicit. Leaving aside insider information (whose use is generally illegal), examples of legitimate private information gathering include detailed analysis of underlying companies or economies or commodities, all the way through to maintaining ultra-fast fibre-optic lines and co-located servers.

As with the formalized/non-formalized scale, the reliance of the decision-making process upon private and public information requires different skills. Private information gathering is an expensive and often technologically advanced process, so we class it as a skill. In contrast, the ability to avoid such expenses, i.e. make investment decisions based on information already disseminated in the marketplace is a skill that we position on the other side of the scale.

Development of computer technologies has been continuously pushing the boundaries: the same information may be classified in a more formalized and public manner now than only a few years ago, and there is no end to this push in sight.

It is worth noting that the input axes answer two fundamentally different questions - how difficult is it to obtain information (public/private) and how difficult is it to make the information usable (formalized/non-formalized). One can have private formalized info (e.g. an exclusive weather forecast) or public non-formalized info (e.g. a central banker hinting at something in her speech). Two other quadrants are obvious. Undoubtedly, one can digitize newspaper articles or apply artificial intelligence to interpret the central banker's speech. However, that would still be transformation of information not a part of decision making. For example, if the artificial intelligence suggests that the central banker has hinted at the possibility of a monetary expansion next year, what should one do about it? The processor will produce an answer.

Another thing to remember is that we are dealing with edges of a hypercube which are not binary, in the general case. Therefore a digitized newspaper is a bit farther from a formalized vertex than a price time series. The central banker's speech is closer to the non-formalized vertex.

Still, if one wishes to have fewer axes in his classification, he can combine the two input axes into one complex/simple axis at the expense of losing some information.

Processor

Unique investment skills relevant for Processor also form a two-dimensional plane. One dimension partitions investment processes into bottom-up vs top-down. The other - into systematic vs. discretionary.

Bottom-Up vs. Top-Down Analysis

Bottom-Up Hybrid Top-Down

Exhibit 8: Process: Analysis scale

The bottom-up and top-down approaches, so familiar to investment professionals, are in fact two alternative approaches to information processing in abstract systems.

In the case of an investment process the bottom-up approach usually means concentrating on the analysis of information relevant to particular securities and largely ignoring the information related to the whole environment. A good example would be an equity market-neutral fund that maintains neutrality to a wide variety of factors, such as market, region, sector, interest rates, size and possibly others. Such fund may focus on analyzing companies' fundamentals and build its portfolio bottom-up because its neutrality would arguably insulate it against macro risk.

Adepts of the top-down approach, in contrast, usually start with the big picture reflected in macro data, and only then descend to more granular levels to form positions in specific securities. A typical discretionary global macro fund would have a view on the economy and select individual trading ideas that should not contravene with it.

Obviously, the two types of analysis require very different types of skills. In reality, however, the two are often combined in some proportions, so one can rarely see their pure realizations. But still, one of them, where the firm has more expertise, would be dominant.

Discretionary vs. Systematic Architecture

Discretionary Hybrid Systematic

Exhibit 9: Process: Architecture scale

Discretionary and systematic information processing architectures are self-descriptive. The former are based on discretionary decisions of portfolio managers and the latter are meant to be purely algorithmic. Each architecture type has its pros and cons. Discretionary processes are supposed to be far more adaptive to changing markets and are better suited to processing hardly quantifiable information. However, since their indispensable components are the black boxes of human brains, the whole process is on average less transparent and replicable. The latter means that it is harder to rely on past performance generated by discretionary managers. Not only because their portfolio managers are always at risk of losing their feel for the market, but also because such firms are more dependent on their key people. On the contrary, systematic managers are supposed to have more reliable processes, but at the same time their investment processes are in general less adaptive and not so suitable for processing qualitative information.

A fundamental but often overlooked distinction between discretionary and systematic architectures lies in a notion of trade. A discretionary manager's trade is a one-time activity in buying/selling financial instruments. A systematic manager's "trade" is a modification to the trading algorithm. Indeed, introducing occasional changes in their algorithms is in fact the only way that systematic managers can affect their performance. Buying/selling financial instruments is the algorithm's trade, not really the systematic manager's.

Hypercube in the Kitchen

The above classification (to be precise, there are two classifications corresponding to the two different output scales introduced) of investment processes based on the five scales can be represented as a 5-dimensional hypercube. Vertexes of this hypercube correspond to $2^5 = 32$ "pure" investment styles, where each pure style is associated with a set of five investment skills. All other points correspond to intermediate states. One can argue that some combinations of skills are more common in real life than others. For example, one would expect to see more funds that are [Formalized & Systematic] than [Non-formalised & Systematic], more [Systematic & Arbitrage] than [Discretionary & Arbitrage] or more [Formalized & Public & Systematic & Single risk premium] than [Formalized & Public & Systematic & Multiple risk premia] .

When cooking an investment strategy, a manager selects (or sometimes is pushed to) a point on each of the edges of the hypercube which eventually determine his position in the managers' universe or, in other words, his investment edge (the authors apologize for the pun). The menu of investment strategies visible to an investor may be split by these five categories and analyzed accordingly.

Note. It is natural to associate an interval [0,1] with each of the five hypercube dimensions and define the Euclidean distance between any two funds:

$$\|F_1, F_2\| = \sqrt{(a_2 - a_1)^2 + (b_2 - b_1)^2 + (c_2 - c_1)^2 + (d_2 - d_1)^2 + (e_2 - e_1)^2}$$

where $(a_i, b_i, c_i, d_i, e_i)\big|_{i=\overline{1,2}}$ are coordinates of the two funds on the five scales.

At the same time, the five hypercube dimensions are not born equal. When a customer comes to a restaurant, he first looks for a type of food that he wants (answering the question "what?"), and only afterwards chooses between competitive products based on how they were cooked or priced. In the same vein, while the first four scales answer the question "How?", the question "What?" is answered by the last one, nested in Output. Hence, one could consider modifying the above formula in the following way:

$$\left\|F_1, F_2\right\|_{\mathbf{w}} = \sqrt{w_1\left(a_2 - a_1\right)^2 + w_2\left(b_2 - b_1\right)^2 + w_3\left(c_2 - c_1\right)^2 + w_4\left(d_2 - d_1\right)^2 + w_5\left(e_2 - e_1\right)^2}$$

where $w = (w_1, w_2, w_3, w_4, w_5)$ is a vector of non-negative dimension weights.

The basis proposed in this paper is not the only skill based basis imaginable. We mentioned some alternative possibilities, and did not mention even more that we had thought of but rejected. The reader may and is indeed very welcome to come up with his own hypercube that better fits into his kitchen and helps produce a better investment product.

Another important question left unanswered is how exactly should one determine the coordinates on all scales? While the answer may be obvious when a strategy clearly lies at an edge (like a price-based, i.e. a formalized-input strategy) it is trickier in the general case. The scales we propose are quantifiable but we have not approached the task of introducing specific measures for each of them. We believe that at this initial stage, discretionary approaches may work reasonably well while the quantification may present a subject for future research. The coordinates and the weights of scales may represent a unique investor's view and help him interpret the menu of investment strategies in his own way.

Concluding Remarks

Summing it up, this paper proposes a new classification of active investment styles based on characteristics of underlying investment processes. The latter are usually determined by unique investment skills that firms possess. Hence we distinguish investment processes by suggesting five skill scales (dimensions), such that an arbitrary investment process can be characterized by five coordinates corresponding to its positions on these scales. Thus, the direct geometrical analogy with a 5-dimensional hypercube.

Importantly, the hypercube dimensions are well-structured as they correspond to different functional parts of an investment process represented in the form of an abstract information processing system consisting of Input, Processor and Output. With regards to Input, we distinguish between Formalized vs Non-formalized information and Private vs Public access to it. For Processor the two scales are Systematic vs Discretionary architectures and Top-down vs Bottom-up analyses. Within Output we suggest segregating Arbitrage and Risk premia performance drivers, but for practical purposes we prefer to distinguish between the use of Single and Multiple risk premia. Apparently, our choice of skill dimensions within each part is rather subjective, but we believe it serves as a good starting point for further discussion.

Each scale represents an interplay of two opposing skill sets, which are hard or expensive to combine within one investment process. Among the scales above, the most fundamental is the performance driver scale, which distinguishes between arbitrage and risk premia as the only two sources of active returns that exist. Crucially, not only arbitrage, but also risk premia strategies require special skills. Since we live in a world where true systematic factors are unobservable and their risk premia are unknown, extracting such factors, estimating their expected returns and managing exposures to them is a skill critical for all non-arbitrage investment strategies. This skill is not covered by the notion of alpha and is, in fact, orthogonal to it. The classification constructed has a direct application in asset allocation and risk management, especially for funds of hedge funds and pension funds. It could also serve as a basis for a new family of hedge fund indices. Though it is not clear how to measure quantitatively an exact location of a fund on each of the scales, their qualitative estimates made by the investor should not represent a problem.

It is interesting to speculate how the active investment industry is going to develop in the years to come. We would expect an increase in specialization, i.e. investment styles of successful hedge funds gradually drifting towards the hypercube vertexes. This is a manifestation of a natural trend towards separation of skills, the one we already witnessed during times when the idea of alpha-beta separation was so popular. We suggest that some form of skill separation similar to the one provided by the hypercube above will eventually replace the increasingly obsolete alpha-beta separation paradigm.

Endnotes

1. We define active investment as any type of investment whose value materially depends on the investment manager's decision-making. This definition is very general and covers a broad class of investment vehicles: hedge funds, mutual funds and ETFs with active investment policies, personal and professionally managed investment accounts. Since managing real estate or private businesses also means participating in investment decision-making, real estate funds, private equity and venture funds do fall into the category of active investments according to the definition above. Moreover, since the value of public and private companies is critically dependent on management decisions and corporate governance, they present examples of active investments as well. However, we won't reach that far and will concentrate on vehicles that invest in financial instruments, where a logical reasoning that we adopt seems the most fruitful. Further such vehicles are denoted in this paper as "active funds" or simply "funds", even if they are implemented via different legal structures such as institutional managed accounts. Hedge funds in this respect represent by far the purest form of active investments since their internal decision making (investment process) plays the most critical role in their performance and survival.

2. We use terms "active investment style", "fund style" and "fund manager style" interchangeably, always referring to an investment process that stands behind the scenes. Of course, some fund managers, especially the largest ones, implement dozens of investment processes simultaneously, so we treat them in our classification as baskets of different active investment styles.

3. Cazalet and Zheng (2014) compare hedge fund classifications employed by different data vendors and propose one of their own based on a role in an investor's portfolio. See also Fung and Hsieh (1997) who first introduced style analysis into the hedge fund world as well as Connor and Woo (2004) for a description of major hedge fund strategies and an overview of several classification principles.

4. For instance, the Credit Suisse index family uses the TASS database classification, breaking the hedge fund universe into 10 major groups (Convertible Arbitrage, Dedicated Short Bias, Emerging Markets, Equity Market Neutral, Event Driven, Fixed Income Arbitrage, Global Macro, Long/Short Equity, Managed Futures, Multi-Strategy).

5. Here we omit Storage (Memory), a part of a system responsible for storing information, because it does not play an essential role in our model.

6. As Cochrane (2007) puts it: "If a piece of information is not correctly incorporated in market prices, we only need a few arbitrageurs or "marginal investors" to trade. They don't have to take very large positions or bear much risk. In fact, the no-trade theorem studies the puzzle that in theory their private information should be revealed in prices with no trading at all! However, if some "systematic" factor (momentum, carry trade, put option writing) has an unwarranted risk premium, the only cure is for that risk to be more widely shared. The average investor must change his demands. This is much harder, so markets can maintain "segmented" risk premia for a long time, even while trading within each market quickly removes any informational "inefficiencies"."

7. The time scale uses relative rather than absolute measures of time. We are only interested, loosely speaking, in the strategy's trade cycle vs a frequency of news. For example, in the real estate market days or even weeks is almost unimaginably short term while in the geographical equity arbitrage moving ahead in a few seconds may be unaffordably long term.

8. As above, we are using the relative notion of time. For instance, to collect the earnings announcement premium a fund manager can buy several days before and sell immediately after the respective event (for further details see Barber et al. (2013)). Even though such a holding period may not seem too long at first sight, in comparison with the relevant news flow it is.

9. As a counter-example, many HFT funds play the role of liquidity providers. In other words, they accept a transfer of risk from a liquidity taker who is likely possessing an informational advantage. Therefore, one can speculate that such funds collect a risk premium associated with their market-making activity.

10. The "peso problem" teaches us that historical track record alone cannot serve as a sufficient evidence of skill. Often high historical information ratio is originated by an extremely skewed distribution of returns, where negative events are very rare but are disastrous for the strategy. The fact that such an event did not occur in the past can create an illusion of an exceptional performance. A good example of such a strategy is writing far out-of-the-money options.

11. The reasoning is based on the fundamental theorem of asset pricing which suggests an equivalence between absence of arbitrage and an existence of a strictly positive stochastic discount factor (SDF), which has an intuitive meaning of an index of bad times. In other words, the absence of arbitrage opportunities is equivalent to the existence of a universal way to discount future cash flows of all assets that makes the present value of each asset equal to its current price. An SDF is a set of random variables determining such discounting for each moment in the future. Hence, it incorporates the risk preferences of all market participants. For instance, if each investor's consumption growth depends on market return only, then SDF is a linear function of future market returns and CAPM holds (see Cochrane and Christopher L. Culp (2003)). The SDF plays so fundamental role in modern finance, that it is also known under many other names such as marginal rate of substitution, state price density, pricing kernel, change of measure and risk-neutral density. Back (2010) showed that when an SDF exists, assets' expected returns are fully determined by their covariances with it. Since an SDF does not in general represent a traded asset or a portfolio, a finite number of uncorrelated factor portfolios are considered instead. One may say that factor portfolios represent different dimensions of an unobservable SDF. Importantly, Jarrow and Protter (2013) show that an expected return of any asset is fully determined by its covariances with these systematic factors. Expected returns of such factor portfolios are called systematic risk premia. Hence, an expected return of any traded instrument or portfolio can be decomposed into a sum of systematic risk premia.

12. See Baltas and Kosowski (2012).

13. This exemplifies, by the way, that exploiting multiple risk premia does not necessarily mean trading many assets.

14. As an extreme example imagine that one is able to count all bottles of Coca-Cola being sold in every store worldwide in real time. Having such an infrastructure would allow a fund manager to see Coca-Cola sales figures well before they are announced by the company in a quarterly report. Another example would be a commodity manager having a network of their own meteorological stations and/or satellites.

15. The latter example highlights a subtle difference between private and public information with respect to market data available from exchanges. A speed advantage over commercial data vendors measured in milliseconds allows an HFT fund manager with an appropriate infrastructure to access what is effectively private information. The fact that such information will stay private for an extremely short time period only does not preclude its potential for profits.

16. The best general definition of top-down and bottom-up approaches that we came across is the one given in Wikipedia (http://en.wikipedia.org/wiki/Top-down_and_bottom-up_design):

"A top-down approach... is essentially the breaking down of a system to gain insight into its compositional sub-systems. In a top-down approach an overview of the system is formulated, specifying but not detailing any first-level subsystems. Each subsystem is then refined in yet greater detail, sometimes in many additional subsystem levels, until the entire specification is reduced to base elements. ...Top down approach starts with the big picture. It breaks down from there into smaller segments.

A bottom-up approach is the piecing together of systems to give rise to more complex systems, thus making the original systems sub-systems of the emergent system. ...In a bottom-up approach the individual base elements of the system are first specified in great detail. These elements are then linked together to form larger subsystems, which then in turn are linked, sometimes in many levels, until a complete top-level system is formed."

17. The term "systematic" used in this section has nothing in common with systematic risk factors discussed earlier.

18. Interestingly, the systematic manager's activity of changing trading algorithms (including conscious decisions of leaving them unchanged, as a special case) inevitably introduces a discretionary component into the whole decision-making process. Therefore, a pure systematic process may hardly exist in real life because it is a process that is guaranteed to stay unchanged and be allocated a certain amount of risk regardless of any exogenous events.

19. The latter pair would describe a typical CTA and a systematic macro, e.g. GTAA strategies. These two types of strategies may also take different positions along the Analysis scale where we would place the CTA closer to the bottom-up side and the systematic macro closer to the top-down side.

20. For example, it may seem natural to consider adding another scale to Processor, fast vs slow, to reflect its information processing speed. However, while being intuitively appealing this concept is closely related to the discretionary vs systematic dimension, which is why we decided to keep it out of our classification.

References

Back, Kerry, 2010, *Asset Pricing and Portfolio Choice Theory* (Oxford University Press).

Baltas, AN, and Robert Kosowski, 2012, "Momentum Strategies in Futures Markets and Trend-following Funds," Paris December 2012 Finance Meeting EUROFIDAI-AFFI, 47.

Barber, Brad M., Emmanuel T. De George, Reuven Lehavy, and Brett Trueman, 2013, "The Earnings Announcement Premium Around the Globe," *Journal of Financial Economics* 108, 118–138.

Brandt, Michael W, and Pedro Santa-Clara, 2006, "Dynamic Portfolio Selection by Augmenting the Asset Space," *The Journal of Finance* LXI, 2187–2217.

Cazalet, Zélia, and Ban Zheng, 2014, "Hedge Funds in Strategic Asset Allocation," *The LYXOR white paper series* 11.

Cochrane, John H., 2007, "Efficient Markets Today," Talk at the Conference on Chicago Economics in November 2007.

Cochrane, John H., and Christopher L. Culp, 2003, "Equilibrium Asset Pricing and Discount Factors: Overview and Implications for Derivatives Valuation and Risk Management," in Peter Field ed.: *Modern Risk Management: A History* (London: Risk Books).

Connor, Gregory, and Mason Woo, 2004, "An Introduction to Hedge Funds," London School of Economics.

Fung, W., and David A. Hsieh, 1997, "Empirical Characteristics of Dynamic Trading Strategies: the Case of Hedge Funds," *Review of Financial Studies* 10, 275–302.

Jarrow, Robert, and Philip Protter, 2013, "Positive Alphas, Abnormal Performance, and Illusory Arbitrage," *Mathematical Finance* 23, 39–56.

Authors' Bio

Igor Yelnik
ADG Capital Management LLP

Igor Yelnik joined ADG Capital Management LLP in 2013. Prior to that he spent 9 years at IPM Informed Portfolio Management AB (Sweden) where he was a Partner and Head of Portfolio Management and Research. Prior to this, Igor co-founded St.Petersburg Capital, an asset management firm that specialised in the Russian securities market, and later Unibase Invest, a managed futures business based in Tel Aviv. Igor graduated from Leningrad Polytechnic Institute in 1986 where he obtained a Master's degree in Computer Science (diploma with distinction).

Boris Gnedenko, *PhD*
ADG Capital Management LLP

Boris Gnedenko joined ADG Capital Management LLP in 2013 as a portfolio manager of a systematic global macro investment strategy. Prior to joining ADG Boris spent 4 years as Head of Investment Management at Sberbank (Russia). Boris also founded the SmartFolio project, which resulted in the creation of a widely used investment software application solving advanced portfolio optimization problems. Boris graduated from Moscow State University in 1997 with a master's degree in mathematics (diploma with distinction). In 2006 he obtained a PhD in mathematical finance from the Central Economics and Mathematics Institute of the Russian Academy of Sciences