Patience Premium

Igor Yelnik
ADG Capital Management

Discussing How the Concept of Ambiguity Aversion Leads to a Patience Premium

Introduction

We introduce the notion of a patience premium, which is based on the concept of ambiguity aversion and is an ambiguity premium. We identify three reasons for the existence of the patience premium:

- Certainty preferences: perceived confidence in the expected performance;
- Comparison with peers: desire to outperform the competition drives the focus towards short-term outcomes;
- Loss aversion: intolerance to negative performance leads to the use of sub-optimal trading strategies.

These reasons are driven by the behavior of market participants and are interconnected.

The phenomenon of the patience premium helps explain why the performance of investment strategies may benefit from having longer holding periods.

Is a Long Term View Good?

From our everyday experience, we know that performance uncertainty is often lower over the long-term than over the short-term. Simple intuition helps explain why this may be the case – even if we know exactly how a process will develop in general, i.e. we know the probability distribution, some random events or unexpected influences may lead to significant fluctuations along the way. For example, we can be comfortable with saying that the US stock market will almost certainly deliver a positive return over the next 50 years, but not so by whether it will be up tomorrow or over the next week.

This effect is immediately explained by the well-known fact that expected return is proportional

to time while its standard deviation is proportional to square root of time (under the assumption that the returns are independent and identically distributed). Hence, the ratio of accumulated returns to their standard deviation should increase for longer horizons.

For example, consider an investment with a mean expected annualised return of 10% and a standard deviation of 20%. In one year the ratio of expected return to standard deviation is: 10% / 20% = 0.5. In ten years, however, it will be: (10% * 10) / (20% * 10^0.5) ≈ 1.58, more than three times higher.

This implies that having a reliable forecast for the mean expected return the investor will be better off by investing over the long-term and absorbing volatility around that mean. In other words, the patient investor will earn a patience premium.

Now let us discuss reasons for its existence in more detail.

Certainty Preference: Finding a Rational Reason for Behaving Irrationally

The investor will never know with absolute certainty whether the expected return of a strategy will be positive. More generally speaking, the investor acts under ambiguity, as they can estimate but will not know for certain the probability distribution associated with an asset or an investment strategy. The problem of portfolio choice under ambiguity has been studied in the academic literature for a long time, see, for example, literature reviews in Tobelem-Foldvari (2010), Izhakain (2012) and Izhakain (2015).

In practice, investors' real-world utility functions are not only about return optimisation over the long term but are also influenced by other concerns and constraints specific to their situation. One example is that a typical investor's perception of losses and gains is asymmetric and they will often be judged over a time frame which is shorter than the one needed to statistically prove a concept. Following on from this, out of two equally volatile investments with equal estimates of expected returns and different levels of confidence in them, a typical investor will naturally choose the investment with a higher confidence. In other words, a higher certainty about expected return is preferred, all else being equal.

Thinking more generally, a higher certainty about the probability distribution is preferred (see, for example, Ellsberg (1961)). According to Easley and O'Hara (2009), this effect known as ambiguity aversion, causes limited market participation and impacts risk premia and in particular the equity premium. In a similar vein, lower participation in slower strategies causes the patience premium.

Put differently, out of two investors with equal ambiguity aversion, one with better knowledge of the expected distribution can afford greater patience, and out of two investors with equal knowledge of the expected distribution, the one with the greater patience should be able to collect the premium due to non-participation of the other investor. Shleifer, Vishny (1997) use the glamour/value anomaly as an example of high uncertainty that prevents many investors from taking advantage of it. They make an important general conclusion that market anomalies must have a high degree of uncertainty to persist over the long term.

Even though the literature mainly deals with the concept of ambiguity about the probability distribution in general, it is sufficient for our purposes to only focus on the ambiguity about the expected mean of the distribution, which we will refer to as uncertainty. The more general use of the term "ambiguity" allows one to account for preferences related to higher moments or joint distributions but we leave these generalisations to the reader.

Considering the uncertainty graphically in Exhibit 1, we show a volatility/mean return plane, which is traditionally used to show the trade-off between risk and expected gain, as a shear of a three-dimensional space in which the third axis shows uncertainty. Moreover, this shear is a very particular one as it assumes no uncertainty around estimated mean of the returns distribution.

If we fix volatility at a certain level, we can consider the mean return/uncertainty characteristics of strategies which will be seen as points on this plane which is orthogonal to the mean return/volatility plane.

The difference between volatility and uncertainty is fundamental for the understanding of risk; *volatility* shows variability of performance around the mean return while *uncertainty* indicates how trustworthy the estimate of the mean return is. An investor who knows the true expected return should only be interested in the volatility. However, in the real world this is rarely the case, which is why investors should take a much more meaningful look at the uncertainty.

Where does this fit in with our concept of the patience premium? It is well known that statistical significance of the mean return estimate depends on the sample size; the larger the sample, the lower the uncertainty, all else being equal. However, it is not the nominal sample size, e.g. the number of days in the sample, that

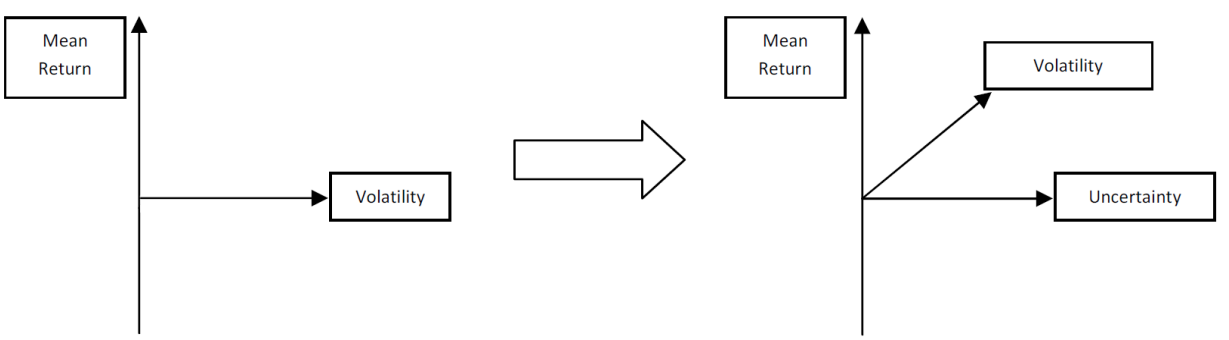

Exhibit 1: The Return / Volatility plane vs Return / Uncertainty / Volatility Space

matters. One needs to use a measure of the effective sample size that reflects the portfolio diversity over time. Faster strategies usually have a larger effective sample size versus slower strategies, all else being equal (see, for example, Gnedenko, Yelnik (2016)). As investors search for lower uncertainty, they are drawn towards faster strategies.

Unfortunately, most statistical tools do not tell us all about uncertainty and many of these tools or approaches assume that the markets are stationary while we know that they change over time.

Additionally, there's an argument to be made that (successful/ profitable) faster strategies play a role in markets being non-stationary. Such strategies are often based on lower capacity anomalies that are identified and traded on by a growing number of managers until they are no longer persistent and can no longer be exploited. At first, investors discover a market anomaly or a risk premium which appears to be statistically significant under the stationarity assumption, then they try to exploit it, and by exploiting it they eliminate the effect they are chasing.

Fast vs Slow: The Future Was Different in the Past

As more investors lean towards faster strategies, they select strategies with lower realised uncertainties and higher realised expected returns, i.e. the returns that were expected in the past. If the market remained stationary, they would be clear winners. However, as we touched on earlier, and as Yogi Berra famously said, "The future ain't what it used to be." As more people identify and chase the same effects, the expected uncertainty grows above the realised uncertainty and the expected return falls below the realised expected return.

Slower strategies are not that lucrative as far as their statistical confidence is concerned. Therefore, the degree of degradation of their expected returns shall not be as significant as that of faster strategies. Unfortunately, higher statistical confidence comes with the trade-off of faster expected performance degradation. In other words, the expected performance degradation should occur slower for strategies with a longer trade horizon for at least two reasons:

1. Higher uncertainty means there are fewer participants utilising slower strategies;
2. Slower strategies are typically less capacity constrained and thus need many more participants to be degraded compared with their faster peers.

A Noisy World: An Effect of Competition

"Half the money I spend on advertising is wasted; the trouble is I don't know which half."
John Wanamaker (1838-1922)

Almost all businesses operate in a competitive environment and this is no different for investors and the asset management industry. As in business, some investors will be more aggressive in exploring and trying to identify new sources of return. Some recent examples include social networks, machine learning and artificial intelligence. Investment companies have a strong financial incentive to be ahead of the curve and introduce novel

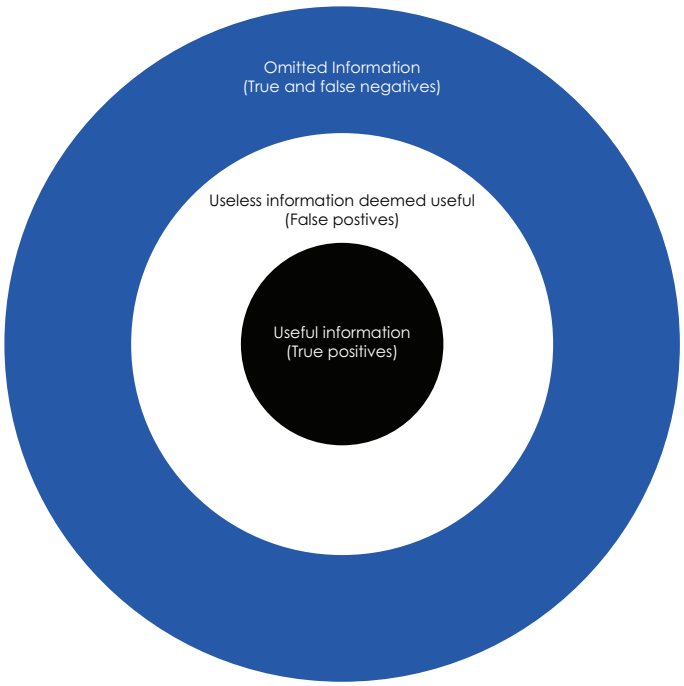

Exhibit 2: Information sets of an Investor

ideas before their competitors. Even if they are not the first to look in a new direction, they may at some point decide that the risk of not joining the crowd is too high. In either case driven by the competition they eventually start expanding the information set used in their decision making in an attempt to get better.

The pinnacle for an investor is to incorporate all useful information into their decisions. Since this task is practically insolvable each identifies a subset that is individual to them instead. This is shown in Exhibit 2 with the blue area indicating information omitted from the decision-making process (which is deemed noise) and the black and white areas indicating included information (deemed useful). The division between the union of black and white areas and the blue ring is subjective and unique to each investor; information omitted by one may be exploited by another.

Explaining exhibit 2 in more detail:

- The outer blue ring shows information omitted by an investor, including both information that is intentionally ignored and that is outside their knowledge, which will contain both true and false negatives or type II errors.
- The white ring shows information that is included (deemed useful) but is objectively useless for an investor so produces false positives or type I errors.
- The black circle represents information that is included and is useful for decision making and is what investors strive to increase.

As mentioned, the combination of the black circle and the white ring shows the full subset of information that the investor includes in their process, both useful (true positives) and useless (false positives). In practice, no investor knows with certainty which information in his subset is black or white so the border between them is blurred. Therefore, a more appropriate representation would be a grey circle scaling from black in the

centre (indicating a higher certainty) to white towards the borders (indicating a higher uncertainty).

As a result, an investor can only know for a fact that he expands the union of the black and the white and hope that the black will expand more. The presence of some unknown amount of false positives in any investor's information subset is inevitable.

For example, price information and some technical data like volume, open interest, etc. populate the black and white areas of a trend-following investor, though there is no way to tell whether a particular type of information lies in the black or white area.

Competition encourages investors to expand the subset of information that they use, moving the boundary of the blue ring since it will contain some useful information (false negatives or type II errors). This expansion of the subset leads the investors to take the risk that they identify more false positives or type I errors that may bring them no investment profits and not lead to higher positive expected returns.

Furthermore, for any given investor, there may exist a faster investor. That faster investor will try to use information the slower investor will omit. Therefore, the faster investor will use more information per unit time, which may increase statistical confidence in his realised expected returns and supposedly give him a competitive advantage. This drives many investors towards higher frequency. This in turn leads to the herding effect and to a lower predictability of usefulness of information or, in other words, to the lack of ability to tell whether information belongs in the black or the white area. As a result, the share of the false positives in the investors' used information subset may swell and the uncertainty of expected returns, increase. Speaking more formally, while the ambiguity about the realised distribution may decrease as the speed increases, the same may not necessarily be true about the expected distribution.

The reciprocal effect is rather weak and mainly driven by the loss of agility by investors growing too big for their markets in practice. The low, especially, the ultra-low frequency space has not been a magnet for investors. Even if it were, capacity of slower strategies is typically higher. As a result, slower strategies are characterised by a more stable predictability of usefulness of information due to less herding. The share of false positives in the slower investors' used information subset does not increase at a fast pace. (Note: It may be argued that in the general case, using low latency data does not necessarily force an investor to trade more frequently. For example, one may collect tick market data in real time, aggregate it and place one trade a month. However, in practice correlation between latency of data used and (a reciprocal of) frequency of trades is high enough for us to ignore the difference between the two for the purpose of this paper)

Pain Threshold: How Patient Can We Be and How Much Noise Can We Tolerate?

The Buddhist concept of dukha can be loosely translated as suffering or unsatisfactoriness. From the Buddhist point of view, dukha is an inherent part of life, which is difficult and imperfect. Dukha is not necessarily physical suffering such as pain, illness or dying but also ordinary, everyday difficulties. For example, the frustration of not getting what one wants may be considered dukha. Buddhism teaches that clinging to the pleasurable and aversion to the unpleasurable eventually results in dissatisfaction.

We don't have to be Buddhists to recognise some truths in the above. Applying this more directly to investing, if positive returns may be associated with the pleasurable and the negative returns with the unpleasurable, loss aversion should result in inferior returns, with all else being equal or in the absence of an informational advantage or harvesting shorter-term premia.

Investors are often tempted to act on noise even though it can't deliver positive expected returns. This happens because of the failure to recognise and acknowledge that one deals with noise. In other words, they are driven or persuaded by a desire to act, which is more likely to be based on instincts and feelings than on new useful information.

These instincts and feelings may be summarised as a concern that the investor's return forecast is not that reliable. Sometimes this concern leads to relying on information that ought to be classified as noise and thus belong in the blue ring but is instead erroneously classified as a useful one and thus appears in the white ring.

Moreover, the loss aversion and the certainty preference often force investors to take actions that result in negative expected returns. A classic example would be stop-loss rules applied in a strategy with negative autocorrelation of returns.

In other words, investors do not only avoid strategies with higher uncertainty, but in the foray to reduce uncertainty they reduce expected returns.

Eventually, there is a human investor behind all, even the most automated, investment strategies, be it a one-off bet against a currency, buy-and-hold exposure to the stock market or a high frequency strategy. The conviction of that investor or the extent of their certainty about the investment strategy is an integral part of the process. However, the majority of investors feel compelled to act too soon or train their strategies to act too soon as losses are unpleasurable. This adds a cognitive element to the otherwise financially and economically supported logic.

An investor who has implicit trust in longer-term return forecasts, patiently waits and lets the information he possesses work for him will avoid trading on noise and enlarging his personal white ring.

Concluding Remarks

The core reason for the existence of the patience premium is investors' aversion to ambiguity (or uncertainty). In other words, the patience premium is an ambiguity premium.

Due to certainty preferences, the average investor may have been drifting into an area of increasing information processing speed and shorter holding periods. This move has involved the use of information never used before. For such strategies, making assumptions about expected performance and its uncertainty may be even more challenging than for more traditionally exploited slower signals, in particular because of the ever-increasing competition in the area. This competition calls stationarity and hence forecastability of the return distribution into question.

Some of this high frequency information may be useful but other sets may not. Then the matter of expedience of shortening holding

periods will boil down to answering a question: does the benefit of using the extra set of useful information exceed the price to be paid for using other sets of information which are irrelevant at best?

Premia exploited in slower strategies should suffer less due to their high capacity nature and because the competition is much more severe in lower holding period strategies, which mainly target different effects and lack the patience of a slow and conservative approach.

However, patience has to build on confidence. Since statistical tools are less helpful in identifying slow investment strategies with positive expected return, discretion of a human investor plays an important role in forming the ultimate evaluation of certainty of such strategies.

Bibliography

Gnedenko, Boris and Yelnik, Igor, 2016. "Speed and Dimensions of Trading", *SSRN Electronic Journal*.

Easley, D. and O'Hara, M., 2009. "Ambiguity and Nonparticipation: The Role of Regulation," *Review of Financial Studies*, Vol. 22, No. 5, pp. 1817–1843.

Ellsberg, D., 1961. "Risk, ambiguity and the savage axiom," *Quarterly Journal of Economics*, 75, 643–669.

Izhakian, Yehuda, 2012. "Capital Asset Pricing under Ambiguity", https://archive.nyu.edu/bitstream/2451/31464/2/Izhakian-CapitalAssetPricing_Feb2012.pdf

Izhakian, Yehuda, 2015. "A Theoretical Foundation of Ambiguity Measurement", http://econ.ohio-state.edu/seminar/papers/151120_Izhakain.pdf

Shleifer, Andrei and Vishny, Robert, 1997. "The Limits of Arbitrage", *The Journal of Finance*, Vol. 52, No.1 (March 1997), pp.35-55.

Tobelem-Foldvari, Sandrine, 2010. "Robust Asset Allocation Under Model Ambiguity," *The London School of Economics and Political Sciences.*

Author 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).