The human-centred, next-generation quantitative approach to constructing portfolios

December 23, 2019

CONTENTS

1 EXECUTIVE SUMMARY

InvestSuite helps banks, brokers, wealth managers and other financial institutions serve their clients better with a suite of human-centred, next-generation investment solutions. At the core of this suite of products lies a portfolio construction methodology, which is based on a measure of risk we have developed in-house: InvestSuite Value-at-Risk (iVaR). In contrast to traditional risk measures, iVaR captures the factors that investors actually perceive as investment risk: the frequency, the magnitude and the duration of losses. Minimising the iVaR value in an optimisation problem allows us to construct portfolios that are more comfortable and less stressful for the average investor. Alongside this, iVaR also addresses the mathematical shortcomings of traditional risk measures. As an overall summary, the properties of our in-house-developed portfolio construction methodology result in a human-centred, next-generation quantitative value proposition.

In the following sections, we explain why traditional measures of risk fail to embrace the intuitive perception of investment risk, and the mathematical shortcomings of such approaches. We then describe how our in-house measure of risk addresses these problems, and finally, we prove the unique value proposition of iVaR in practice through backtesting.

2 TRADITIONAL MEASURES OF RISK

2.1 VOLATILITY

Nassim Nicholas Taleb, most commonly known for his book 'The Black Swan', raises the question of whether, when exploring a large city that one has never been in before, one would prefer to take an inaccurate map or no map at all. With regard to financial modelling, many investment practitioners and academics prefer an inaccurate map to no map at all. In this context, the inaccurate map refers to Modern Portfolio Theory (MPT). In 1952, Harry Markowitz introduced MPT as an approach to constructing investment portfolios, receiving a Nobel prize for this pioneering work. MPT comes down to finding a portfolio that maximises the expected return given a certain amount of risk, or equivalently, minimising the amount of risk for a given expected return. Markowitz defined the amount of risk as volatility, which is also referred to as the standard deviation, a value that is measured as the dispersion relative to the mean. In other words, if a given financial security has a larger price range where the data points are further removed from the mean, this results in higher volatility. This is of course not what investors intuitively perceive as investment risk.

"A lot has happened since I published that article in 1952."

Harry Markowitz

2.1.1 MATHEMATICAL SHORTCOMINGS

The main problem with volatility as a measure of risk is that it treats sudden positive returns as being just as risky as sudden negative returns, and on the other hand, presents consistently negative returns as not being risky at all. Table 1 visualises the yearly returns of three different portfolio strategies, where the last row of the table represents the ten-year volatility for each of these portfolios.

We can see that the yearly returns from Portfolio 1 are exactly the same as those from Portfolio 2, but with a negative sign. The consistent negative returns from Portfolio 1 are of course not desirable, and an investor would always opt to be invested in Portfolio 2. Nevertheless, Portfolio 1 and Portfolio 2 are considered to be equally risky in terms of volatility.
In addition, while Portfolio 2 and Portfolio 3 deliver relatively similar positive returns, the volatility of Portfolio 3 is 15x larger than that of Portfolio 2. This jump in volatility for Portfolio 3 is caused by three years of extraordinary positive performance. While Portfolio 2 gives very desirable returns, an investor would always prefer to be invested in Portfolio 3; however, the volatility measure punishes portfolios that are characterised by high variability in positive returns.

Yearly returns	Portfolio 1	Portfolio 2	Portfolio 3
Year 1	- 5 %	5 %	5 %
Year 2	- 9 %	9 %	9 %
Year 3	- 6 %	6 %	6 %
Year 4	- 9 %	9 %	40 %
Year 5	- 8 %	8 %	8 %
Year 6	- 4 %	4 %	4 %
Year 7	- 8 %	8 %	65 %
Year 8	- 7 %	7 %	9 %
Year 9	- 6 %	6 %	8 %
Year 10	- 5 %	5 %	80 %
Volatility	**1,77 %**	**1,77 %**	**28,1 %**

Table 1: ten-year volatility for three different portfolio strategies.

Markowitz later acknowledged the shortcomings of volatility and proposed a new measure: semi-variance. This measure of risk is calculated in a similar manner to variance, but only takes into account those observations that fall below the mean. Although this risk measure does not penalise portfolios with high variability in positive returns, it would still present a portfolio that consistently loses 10% each year as risk-free.

2.2 VALUE-AT-RISK

In the 1980s, JP Morgan formalised the measure of Value-at-Risk (VaR), highlighting the idea that investors want to avoid extreme and negative events in general at the loss side of the return distribution. For a given confidence level $\alpha * 100\%$, the VaR represents the lowest amount of money one would expect to lose within a given time period, when the loss is higher than $\alpha * 100\%$ of the losses. For example, for an investment portfolio with a one-year VaR of 1 million USD with confidence level 0.95, there is a $(1-0.95) * 100\%$ (=5%) probability that the portfolio will lose more than 1 million USD over a one-year period. This is visualised in Figure 1. Although VaR provides more useful information in terms of what concerns investors compared to volatility, it still fails to capture what humans intuitively perceive as investment risk.

2.2.1 MATHEMATICAL SHORTCOMINGS

Although the VaR measure was adopted by the Basel Committee on banking Supervision in the Basel Accords, there are four concerning issues in terms of using VaR as a measure of risk. First, when using VaR as a risk measure, diversification can increase the total portfolio risk. The example below further clarifies this issue.

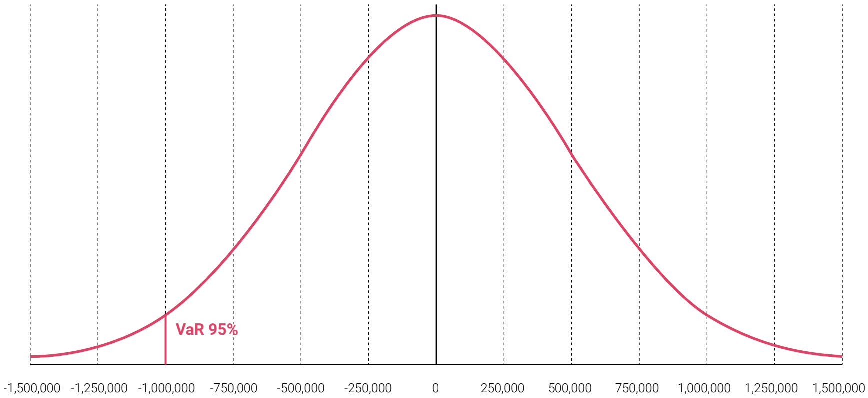

Figure 1: VaR 95% for a fictive investment portfolio.

Consider investor A and investor B. Investor A holds one bond worth 100 USD with a 1.5% probability of default over a time horizon of one year. Investor B holds two bonds, each worth 50 USD, both also having a 1.5% probability of default over a time horizon of one year. With a confidence level of 98%, $VaR_{0.98}$ = 0 USD for investor A, because in 98% of cases, the investor will not lose more than 0 USD. The $VaR_{0.98}$ for investor B is equal to 50 USD, because in 98% of cases, the investor will not lose more than 50 USD.

The fact that diversification can increase portfolio risk under VaR can also be derived from investigating the mathematical properties of the measure itself. Follmer and Schied [2002, 2010, 2011] recognise that the essence of diversification is captured in the sub-additivity axiom. Since the VaR measure fails to fulfill this axiom, diversification can result in higher overall risk in the context of VaR. It is important to note that this sub-additivity property is a requirement to be classified as a coherent[1] measure of risk. Consequently, VaR is also not considered coherent.

A second flaw of VaR was highlighted by the worldwide financial crisis in 2008, concerning the fact that models based on VaR assume a market with normal conditions. Consequently, this measure is of very little relevance when, for example, a crisis hits the financial markets.

[1] See Appendix 1 for an overview on the coherent risk properties.

Figure 2 visualises a VaR estimate for an AAA-rated security. VaR tells us that there is only a 1% chance of losing more than 31,815 USD on a 1,000,000 USD portfolio in any one month. However, in September 2008, the actual loss of a similar portfolio was 50,000 USD, which should only happen once every 700 years. In November 2008, the actual loss amounted to 180,000 USD, which should occur only once every two duodecillion[2] months. In other words, the VaR measure cannot be relied upon in non-normal contexts.

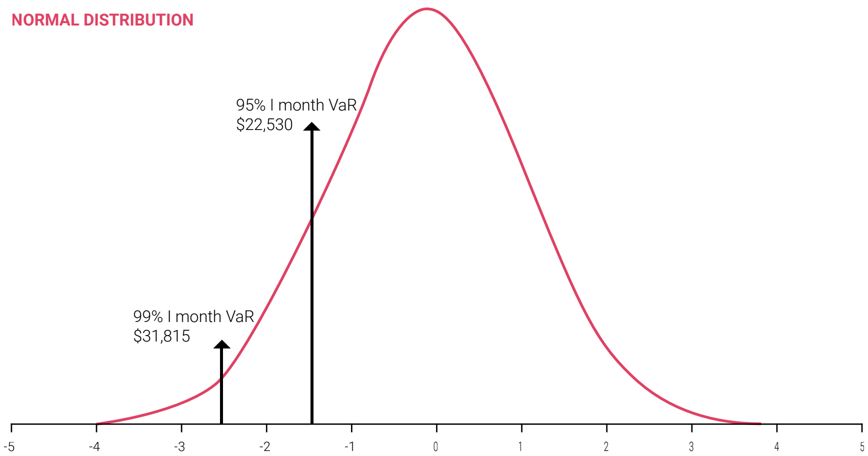

Figure 2: P&L distribution of a 1 million USD portfolio of 'AAA' CMBS.

Third, VaR suffers from the fact that a value for the confidence level α needs to be determined. In portfolio optimisation this becomes a hyper-parameter problem, which makes the optimisation less robust. This is because the choice of the confidence level α can have a significant impact on the result. It should be noted, however, that the selection of α is often done in quite an arbitrary fashion.

Fourth, since VaR represents the lowest amount of money one should expect to lose within a given time period, but only when the loss is higher than $\alpha * 100\%$ of the losses, the measure completely ignores the losses beyond the threshold value. This is worrying, since there always exists a chance that an investor will lose more than this threshold value. In other words, VaR does not give an indication of this 'Tail Risk'.

2.3 CONDITIONAL VALUE-AT-RISK

Compared to VaR, Conditional Value-at-Risk (CVaR) is a preferable alternative as a measure of risk. CVaR addresses the issue of 'Tail Risk' and even qualifies as a coherent measure of risk. For a given confidence level α, CVaR measures the mean of the worst $(1-\alpha) * 100\%$ losses. In other words,

---

[2] Duodecillion is equal to 10^{39}. The universe is 'only' 13.9 billion ($13.9*10^9$) years old.

it calculates the tail mean of the loss distribution. For example, take the investment portfolio visualised in Figure 3. With a confidence level of 0.95, the $CVaR_{0.95}$ is equal to the mean of the (1-0.95) * 100% (=5%) worst losses; in other words, it is equal to the mean of the losses that are larger than 1,000,000 USD, which is the VaR, or the threshold value.

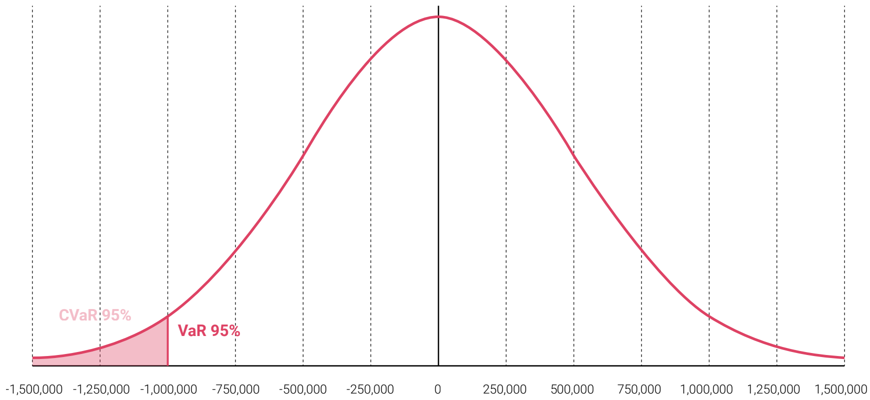

Figure 3: 95% VaR and 95% CVaR for a fictive investment portfolio.

Although the calculation and interpretation of CVaR are quite similar to those of VaR, there is an important difference between these measures of risk. While VaR does not meet the sub-additivity axiom, CVaR satisfies all of the required axioms to qualify as a coherent measure of risk. Since CVaR is coherent, it has received a great deal of attention in academic literature and in practice. The Basel Committee on Banking Supervision recognises the issues with VaR and is in the process of shifting from VaR to CVaR in the renewed Basel Accords. By taking into account all losses beyond the threshold value, CVaR represents another improvement in terms of capturing the information that investors are concerned with. Nevertheless, it still ultimately fails to capture what humans intuitively perceive as investment risk.

2.3.1 MATHEMATICAL SHORTCOMINGS

As a first example, investors are not only worried about extremely large losses, but also about smaller losses, which are not captured by CVaR. Second, not only the magnitude of losses matters to investors, but also the time it takes to recover from those losses, also referred to as the duration or time spent under water. Since CVaR fails to take into account the autocorrelation of portfolio returns, the duration of losses is not captured if this measure of risk is used. Third, CVaR suffers from a similar issue as VaR, in that a value for the confidence level α needs to be determined. In portfolio optimisation this again becomes a hyper-parameter problem, making the optimisation less robust.

2.4 CONCLUSION

The development of the VaR and CVaR measures surely represents steps in the right direction. Nevertheless, these measures do not capture what investors intuitively perceive as risk. As is recognised by Marc Odo (2017), most investors are likely to define the risk of an investment as 'simply not losing money', rather than thinking about the variance (var), VaR or CVaR. In financial risk measurement, the concept of 'simply not losing money' translates to limiting drawdowns and striving for monotonic growth. In contrast to the more traditional approaches, drawdown-based risk measures have not received as much attention in the academic literature to date. Although some drawdown-based risk measures have been developed in recent years, none of these exactly capture what humans intuitively perceive as investment risk: the frequency, the magnitude and the duration of losses (drawdowns). Because this is exactly what the in-house developed InvestSuite Value-at-Risk (iVaR) measures, we consider it to represent the most human-centred approach available for measuring investment risk. Figure 4 compares the traditional risk measures of volatility, VaR and CVaR with iVaR based on five criteria in order to demonstrate the superiority of iVaR.

	Takes into account that investment returns are non-normally distributed	Penalizes only losses, not gains	Considers the entire return distribution (i.e. all possible losses)	Takes into account the time it takes to make up for losses	Captures what people intuitively perceive as investment risk
INVESTSUITE VaR (iVaR)	✓	✓	✓	✓	✓
CONDITIONAL VaR (CVaR)	✓	✓	✓		
VALUE-AT-RISK (VaR)	✓	✓			
VOLATILITY					

Figure 4: Comparison between traditional risk measures and iVaR.

3 INVESTSUITE'S APPROACH TO RISK MEASUREMENT

3.1 EMBRACING THE HUMAN PERCEPTION OF RISK

InvestSuite's approach to measuring risk arose from the desire to capture what people intuitively perceive as investment risk. In the end, investors all want the same thing – an account that offers the steady growth of a savings account with a very low probability of losing money, but with the returns of the stock market. We cannot guarantee this, but we can optimise for it, which is precisely why we have developed our own measure of risk: iVaR[3]. By minimising risk against this measure, we are optimising for the most comfortable growth journey in an investment portfolio. Figure 5 clarifies what iVaR is about. In a very intuitive manner, it can be seen as the area between the horizontal lines from the running maximums of the portfolio value and the actual curve representing the portfolio value. The iVaR value is calculated as the average of all these drawdown areas. An investor would obviously want the iVaR value to be as low as possible, because that would represent the minimisation of the frequency of drawdowns, their magnitude, and the time taken to recover from these events.

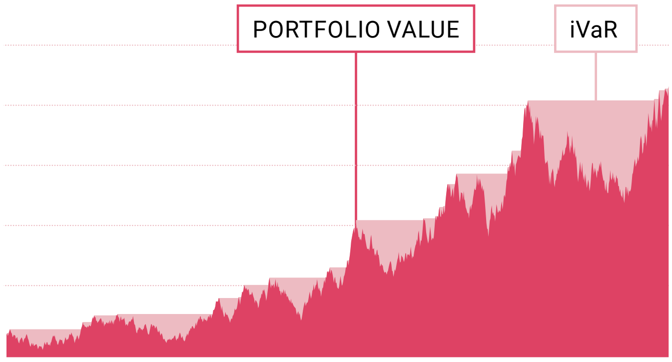

Figure 5: The in-house-developed InvestSuite Value-at-Risk measure.

[3] In our portfolio construction methodology, iVaR can be interpreted as the average expected drawdown (%) in a portfolio.

Because iVaR is so intuitive, the calculations involved are very straightforward. Ex post, iVaR can be derived as follows. We assume an investor with wealth V_t that is observed at regular intervals (times) t = 1, 2, ..., T. On each date, we define the rolling maximum of the investor's wealth V_t as:

$$M_t = \max_{s \leq t} V_s \tag{1}$$

The rolling maximum is thus the highest portfolio value up to a certain date. The corresponding drawdown on date t is then given by:

$$D_t = M_t - V_t. \tag{2}$$

The drawdown on a certain day is thus the difference between the rolling maximum and the portfolio value on that specific day. On date T, for a given time period [1,T], the (ex-post) Accumulated Drawdown is then given by:

$$AD_T = \sum_{i=1}^{T} (D_i). \tag{3}$$

The Accumulated Drawdown within a given time period is simply the sum of the different drawdowns within that time period. The iVaR value is then calculated as the mean of the accumulated drawdowns:

$$iVaR = \frac{\sum_{i=1}^{T} (D_i)}{T}. \tag{4}$$

3.2 BACKTESTING

From a more technical point of view, it is important to note that iVaR allows for backtesting, is a robust measure of risk, and is computationally simple. These three criteria are significant when considering the use of a risk measure in an optimisation problem. Since iVaR allows for backtesting, we are able to scientifically validate our human-centred, next-generation quant portfolio construction methodology and prove that it works extremely well in practice. The objective of our portfolio construction framework is to minimise the frequency, the magnitude and the duration of drawdowns; in other words, the goal is to minimise the iVaR value. Applying iVaR in portfolio construction should thus lead to portfolios that have less frequent and less severe drawdowns, and that recover from these drawdowns more quickly, compared to traditional portfolio construction frameworks. The backtest example showcased below applies to the EuroStoxx 50 universe[4] for the period 2004–2019. Figure 6 visualises the core concept of our portfolio construction framework. It plots the frequency of the % drawdowns in both the EuroStoxx 50 Total Return Index (TRI) and the iVaR portfolio construction framework applied to the point-in-time[5] EuroStoxx 50 universe.

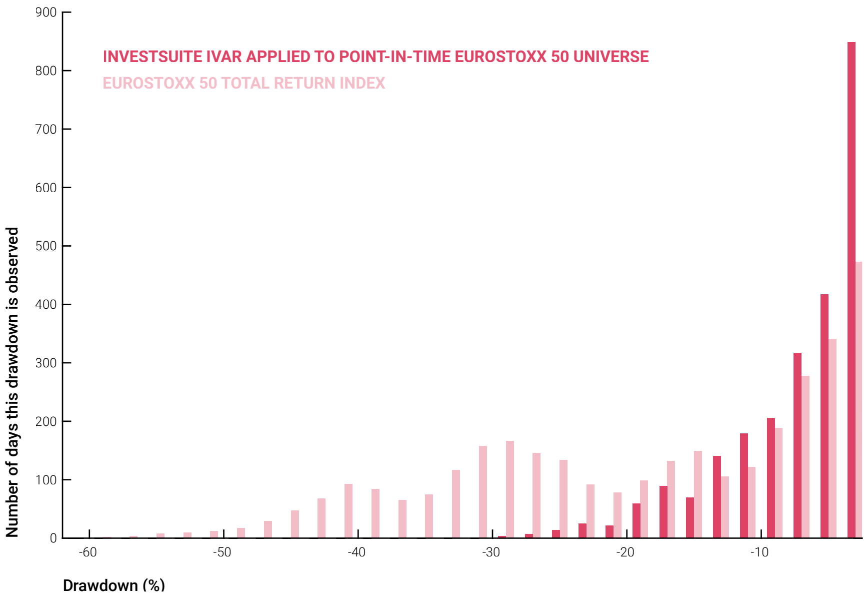

Figure 6: Frequency of % drawdowns.

It can be observed that the maximum drawdown for the InvestSuite portfolio construction framework is about 28%, while the value is around 60% for the EuroStoxx 50 TRI. Even more importantly, drawdowns larger than 15% are consistently more frequent for the EuroStoxx 50 TRI. Because the Invest-Suite portfolio construction framework aims to minimise the frequency, the

---

[4] If you are interested in applying our backtesting approach to your own universe, please do not hesitate to contact us.
[5] To backtest accurately on the EuroStoxx 50 universe, we rebalance the portfolio monthly to update the constituents of the Stock Index universe. At any point in time, there is 100% equity exposure.

magnitude and the duration of drawdowns, it is a natural consequence that smaller % drawdowns are more frequent for the iVaR portfolio construction approach than for the EuroStoxx 50 TRI. However, the effect of this concentration in smaller % drawdowns has no severe consequences, because larger % drawdowns are consistently reduced in frequency and duration. This is demonstrated in Figure 7, which plots the portfolio values for both the EuroStoxx 50 TRI and the iVaR portfolio construction framework, applied to the point-in-time EuroStoxx 50 universe. Focusing on the periods where the general stock market suffered, particularly the years 2008 and 2016, the value proposition of the InvestSuite methodology becomes clearer. From Figure 8, we can see that the annual return in 2008 for the InvestSuite framework is around -6%, while a disappointing -42% is shown for the EuroStoxx 50 TRI. In 2018, the annual return for the InvestSuite framework is positive, at 4%, while the EuroStoxx 50 TRI drops by more than 10%. Furthermore, Figure 7 visualises that the duration of the drawdowns is much shorter for the InvestSuite portfolio construction framework than for the EuroStoxx 50 TRI: for the latter, the losses suffered in 2008 are only recovered from early 2015. The demonstration of such significantly superior performance from our portfolio construction framework compared to the EuroStoxx 50 TRI is merely a fortunate consequence of our methodology, because our framework does not explicitly optimise for outperformance.

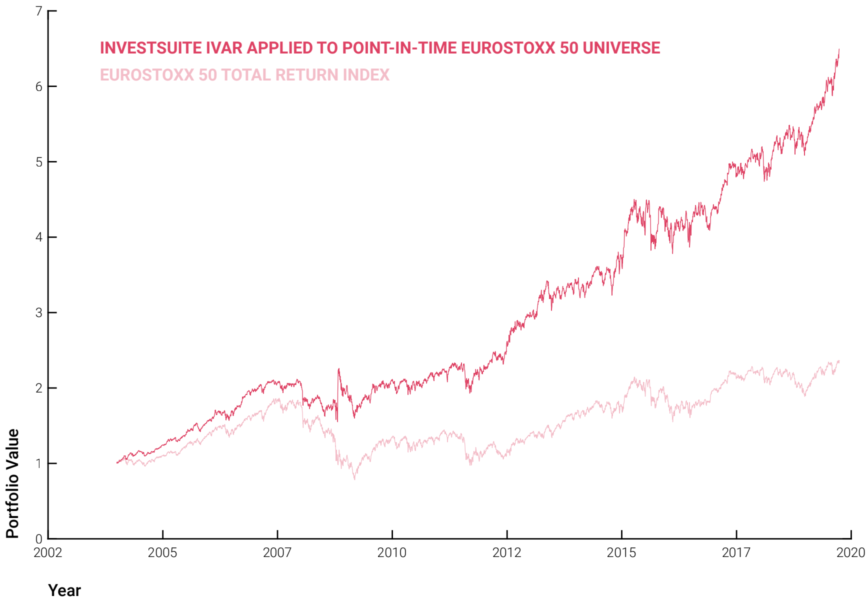

Figure 7: Portfolio values visualised over time.

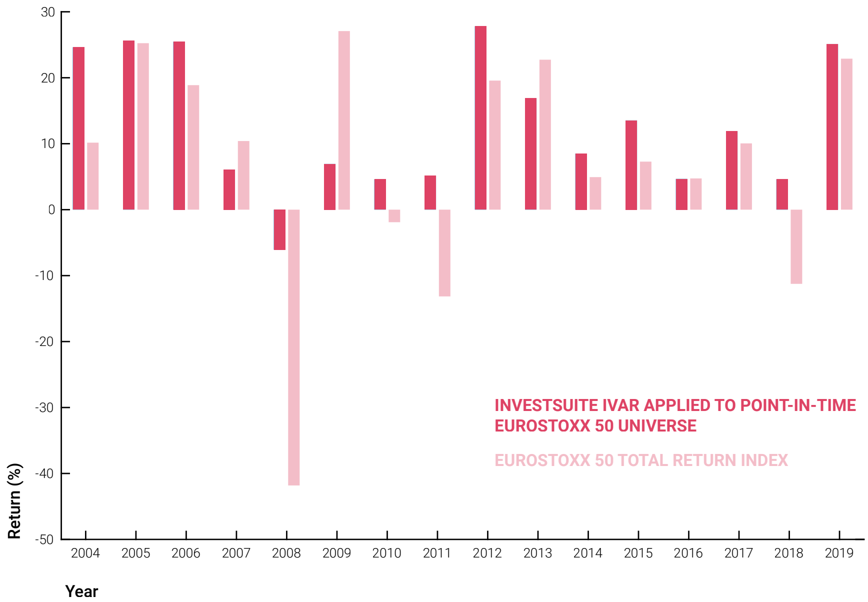

Figure 8: Annual % returns.

3.3 IMPLICATIONS

These backtesting results prove the unique value proposition of our risk measure in relation to a portfolio construction problem. When investors want to minimise the risk of their investment portfolio, this essentially means that they want to minimise the stress involved with investing. This stress is not a result of experiencing high volatility or high theoretical VaR or CVaR values in a portfolio, but rather simply a result of losing money. Our iVaR portfolio construction framework aims to minimise this stress, which has important advantages for both investors and for his or her portfolio managers/advisors. When an investor experiences large drawdowns in their investment portfolio, emotions of stress and fear blur rational investment decision-making processes, which often results in panic selling, the cost of which is significant. According to a study by JPMorgan Asset Management, for the period 1998–2017, the average investor had an annualised total return (CAGR) of only 2.6%, whereas buying and holding the S&P500 index for the same period resulted in a CAGR of 7.2%. This significant underperformance is caused by bad market timing and panic selling, as visualised in Figure 9. Excluding the ten best days for the S&P500 in terms of returns over the whole period results in a CAGR of ‘only’ 3.5%. Since recoveries from stock market declines often provide huge gains, the best days in terms of return often come straight after such a market decline. However, many investors panic during these severe market declines and sell some of their positions. Consequently, they are likely to miss out on huge potential gains during the recovery, meaning that the cost of this panic-selling tendency is significant. Because iVaR captures what people intuitively perceive as investment risk, the constructed portfolios will approximate a smooth ride with less severe drawdowns, thus making investors less vulnerable to panic selling.

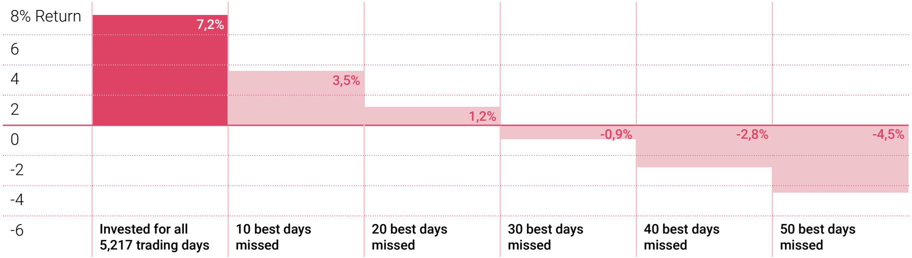

Figure 9: % annualised total return for the S&P500.
Source: Morningstar.

Portfolio managers and advisors also benefit from using iVaR in portfolio construction. As visualised in Figure 7, drawdowns larger than 15% are consistently less frequent for the iVaR portfolio construction framework compared to the EuroStoxx 50 TRI. The maximum drawdown for the iVaR portfolio construction framework is about 28%, while the maximum figure is around 60% for the EuroStoxx 50 TRI. According to Chekhlov[6] et al. (2005), a 50% drawdown is unlikely to be tolerated in any average account, and an account may be closed if drawdowns breach 20% or have lasted for over two years. The iVaR portfolio construction framework thus results in fewer difficult discussions for the advisor. This is because the frequency, the magnitude and the duration of drawdowns are the factors that cause investors to call their advisors, not an increase in volatility or a high theoretical VaR or CVaR. Portfolios where the iVaR measure is minimised naturally minimise these three aspects of drawdowns, thus resulting in portfolios that are more comfortable for investors.

[6] Chekhlov, A., Uryasev S., and M. Zabarankin (2005), "Drawdown Measure in Portfolio Optimization", International Journal of Theoretical and Applied Finance, Vol. 8, No. 1, pp. 13-58.

4 APPENDIX

4.1 COHERENT MEASURES OF RISK

A coherent risk measure ρ, as defined by Artzner & al. (1999), must respect the following properties:

1. Normalisation: $\rho(0) = 0$. The risk of holding no assets is zero.

2. Monotonicity: If X1 and X2 are two portfolios and $X1 \leq X2$ almost surely, then $\rho(X1) \geq \rho(X2)$.
   The risk of a better portfolio is always lower.

3. Positive homogeneity: If $\lambda > 0$, then $\rho(\lambda X) = \rho\lambda (X)$.
   This means that, if an investor doubles their portfolio, they double their risk.

4. Translation invariance: If m is a constant, $\rho(X + m) = \rho(X) - m$.

5. Sub-additivity: If X_1 and X_2 are two portfolios, $\rho(X_1 + X_2) < \rho(X_1) + \rho(X_2)$. The risk of two portfolios together cannot be higher than that of adding the risks of each portfolio. This is the concept of diversification.

6. Convexity: The notions of sub-additivity and positive homogeneity can be replaced by the notion of convexity: for $\lambda \in [0,1]$, $\rho(\lambda X_1 + (1-\lambda)X_2) < \lambda \rho (X_1) + (1-\lambda)\rho(X_2)$.

Discover our suite of products:

Robo Advisor
A low-cost, customisable digital wealth management tool that can deliver hyper personalised portfolios for Goal Based Investing

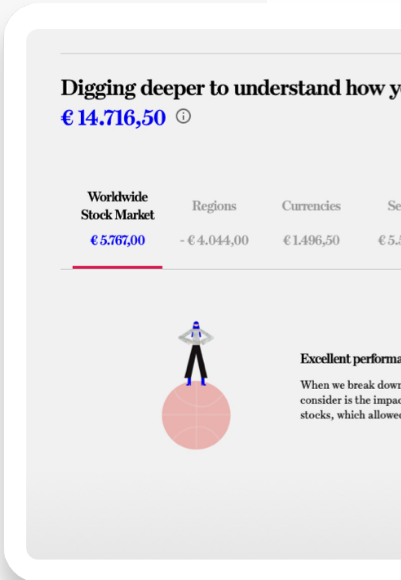

StoryTeller
A worldwide first new way to tell the story on portfolio performance

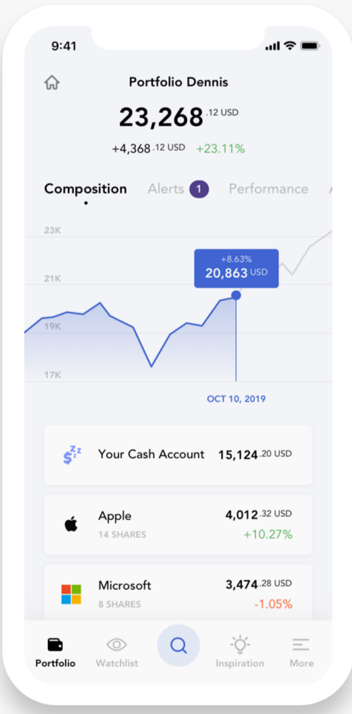

Self Investor
A best-in-class powerful white-label execution-only platform for easy investing

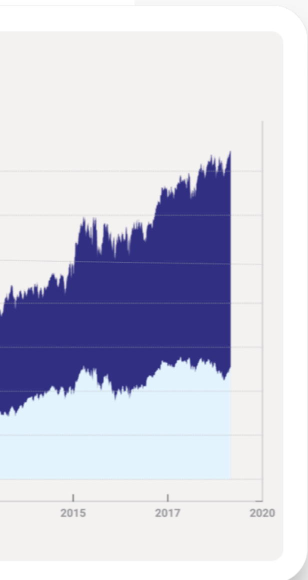

Portfolio Optimizer
The next-generation quant tools that provide cost-effective solutions for more efficient portfolio management

InvestSuite