

STANLIB

A “Wee Dram”
and the Big Question:

Optimal offshore
allocations under
new Regulation 28
limits

STANLIB Multi-Strategy

September 2022

Key Points

- Every investor's optimal offshore allocation should be different, but our work points to two crucial factors to consider:
 1. their willingness, or ability, to hedge some of the offshore currency risk
 2. their required real return, which is synonymous with risk appetite.
- If an investor can hedge currency risk, our analysis would suggest they should utilise most of their offshore allocation and hedge some of the offshore currency risk back to rands.
- The proportion of an investor's offshore allocation that they should hedge is the inverse of their risk appetite. The average investor should hedge about half, conservative investors should hedge more, more aggressive investors less.
- The case for offshore allocation is much less compelling for investors who are unwilling or unable to hedge the currency. Offshore allocation for these investors should simply rise in line with their risk appetite.

Our results are **framed through the lens of the new Regulation 28 limits**, but they apply to all South African multi-asset investors.

We apply a similar process and thinking to portfolios with different asset class universes, constraints, time horizons and return objectives.



David McNay

Senior Portfolio Manager
STANLIB Multi-Strategy

A “Wee Dram” and the Big Question

In the first week of March, I was getting married on a remote part of, a remote island, in the Highlands of Scotland, an area renowned for turbulent seas and outstanding natural beauty. Two days before the wedding, I was holding a “wee dram” in one hand, and a laptop in the other, trying to eke out enough signal to call SA to check on the storm, and keep a hand on the tiller, in the wake of a big change to Regulation 28.

For international readers, or domestic investors who somehow missed the news, on 23 February 2022 SA’s Minister of Finance announced a relaxation of exchange controls, extending what was ostensibly a 30%¹ cap to a cleaner 45% limit. This is, in no uncertain terms, a very big deal. In our view², offshore allocation is likely to be the **single most important determinant of the returns achieved by many South African investors in the coming decade.**

We were contacted by many of our clients asking for our thoughts. We were reluctant to give a view before we had done the work; having had time to crunch the numbers we are ready to set out our conclusions in this paper. We should first say that, as important as the new regulations are, we believe that investors have time to adapt to this new reality, since its impact will take years to play out.

For marketing purposes, our lives would be much simpler if there was a magic number for offshore exposure, say 38%, but such a simplistic approach belies the complexity of the question and would not be worthy of our audience. The truth is that the optimal level for any investor will reflect their risk appetite, investment universe, time horizon, and various other factors. But we can still establish the framework for investors to understand their likely best bet. Such a framework cannot properly be described on the back of a beer mat, so this paper we aim to:

- Understand the history: we look at historic market data to establish how one should ideally³ have managed offshore allocation and rand hedging over the last two decades.
- Take a view: based on our expectation of prospective returns from difference asset classes, we establish how optimal asset allocation looks for investors with different target returns and appetites for risk.
- Discuss foreign exchange: we look at the mechanics of hedging currency risk and how South African investors should include it in their toolkit.
- **FAQ:** we briefly answer some common questions we have been asked over the last few months.

¹ Plus, an additional 10% in Africa and some caveats on firm-wide limits.

² My view is probably biased by being a professional asset allocator and having spent many years researching and developing longer-term solutions.

³ If the current Regulation 28 limits had been in place for the last 20 years, which obviously was not the case.

Asset Allocation through the Rear-View Mirror

*“The four most dangerous words in investing are ‘**this time is different**’”*
Sir John Templeton

The STANLIB Multi-Strategy group’s asset allocation process is unashamedly forward-looking, but we recognise that there is nothing new under the sun: history is a great place to start building a framework to understand what might happen in future. We are not hostages to the past, though. Our attitude is best expressed by the adage that “history doesn’t repeat itself, but it does rhyme”. Looking backwards does also have one significant advantage, in that there are fewer areas to dispute – the data is the data. That differs from looking forward, when, particularly for a country like South Africa, we can extensively debate Capital Market Assumptions.

To understand what the optimal multi-asset portfolios over the last 20 years would have looked like, we must first establish some ground rules.

Firstly, we must agree the universe of assets available to the allocator over the last two decades. We think a practical approach is to limit our universe to what we consider to be the core “beta building-blocks” for a South African investor constrained by Regulation 28. To make our analysis relevant to as many readers as possible, we have excluded niche asset classes, or those that require domain-specific expertise. We have only included asset classes which meet the following requirements:

1. Available to investors via cost-effective passive products (ETFs, index futures, or tracker unit trusts).
2. Sufficient market depth to be able to provide daily liquidity, especially during times of stress, which unfortunately removes South African credit and inflation-linked bonds.
3. Sufficiently differentiated not to be a sub-asset class or a style-premia⁴ of the headline beta.
4. Long and robust total return data – as close as possible to 20 years of continuous total-return data that is easily available via Bloomberg or a similar service.

With that in mind, the following “beta building-blocks” qualify:

- South African Equity
- South African Bonds
- South African Listed Property⁵
- South African Cash
- Global Equity
- Global Bonds
- Global High Yield
- US Dollar Cash

Gold bugs will complain that the yellow metal is missing from our universe of asset classes. I can see their point – physical gold meets all the conditions we have specified. We have excluded it because it has no nationality. It is accepted as both domestic and international and therefore adds no value to the offshore allocation discussion⁶.

⁴ Coded into our DNA is the view that investors have an opportunity to earn “alpha” by incorporating risk-premia strategies, sector and thematic portfolios and other strategies structurally and opportunistically. But to avoid muddying the waters of our analysis, and comparing apples with avocados, we exclude all but the core beta here.

⁵ Listed property is contentious for our inclusion rather than exclusion. As we will discuss shortly, the historical return and volatility for both is comparable, so we have sympathy for the view that SA Listed Property is a sub-asset class of SA equity, rather than its own asset class. But the practical reality is that ASISA composites give Listed Property its own allocation and therefore it is considered a separate asset class for many investors. Its correlation to headline benchmarks is also lower than for other super sectors.

⁶ Beyond being stateless, gold does not always seem to know its raison d’être – is it a currency, a precious metal, an industrial commodity or simply a store of value? Also, it does not generate a cash flow, making modelling prospective returns even more difficult. Although there is a reasonable body of evidence supporting gold as an asset class, we tend to apply a higher hurdle rate for the inclusion of gold into strategic allocations.

The Technical / Wonky Bit

For all “technical” choices, we favour the easiest implementation that we feel is robust. Our main effort is to provide insights into the relative merits of offshore allocations, not impress readers with displays of quantitative skill. Most of our results could be replicated by anyone with access to Bloomberg data, moderate coding skill (in something like Python, MATLAB, R) or access to an off-the-shelf optimiser.

Our dataset for our core universe spans more than 18 years of weekly total returns from 14/11/2003 to 01/07/2022, corresponding to the inception of the FTSE/JSE SA Listed Property Index (JSAPTR). We considered “splicing” proxy data to SA Listed Property to lengthen the series, but when we ran our analysis excluding property from the universe the results were barely affected, and we conclude the cost to the cleanliness of the data was not worth the extra few months of data.

Most results are presented as the output of a mean-variance optimiser (MVO). We do “stabilise” the sample variance-covariance matrix (VCV) using Ledoit-Wolf shrinkage (Ledoit & Wolf, 2004), which is a relatively common transformation, and for our purposes makes a much better estimator of covariance, for minimal implementation effort.

Generally, MVO is not our preferred optimisation methodology for real world problems. Issues are well documented (Michaud, 2004), but the main problem is that, in the typical implementation, one runs into so-called “corner solutions”. One optimisation tool we use regularly is “resampling”, the full details of which are beyond the scope of this note. Resampling starts with the assumption that our expected returns are wrong, instead using a multivariate copula to estimate statistically similar returns which we use to average thousands (typically 10 000) efficient frontiers of different return profiles. The result is a smoother efficient frontier with greater diversification.

Before running our historical optimisations, we will take a step back and look at how our building blocks have actually performed over our historical timeline. The chart below shows historical volatility and realised returns for each of our building blocks. South African building blocks appear as black diamonds, while each offshore building block is represented by two diamonds: red for local currency and green for currency-hedged.

Historical Returns: Nov 2003 to June 2022

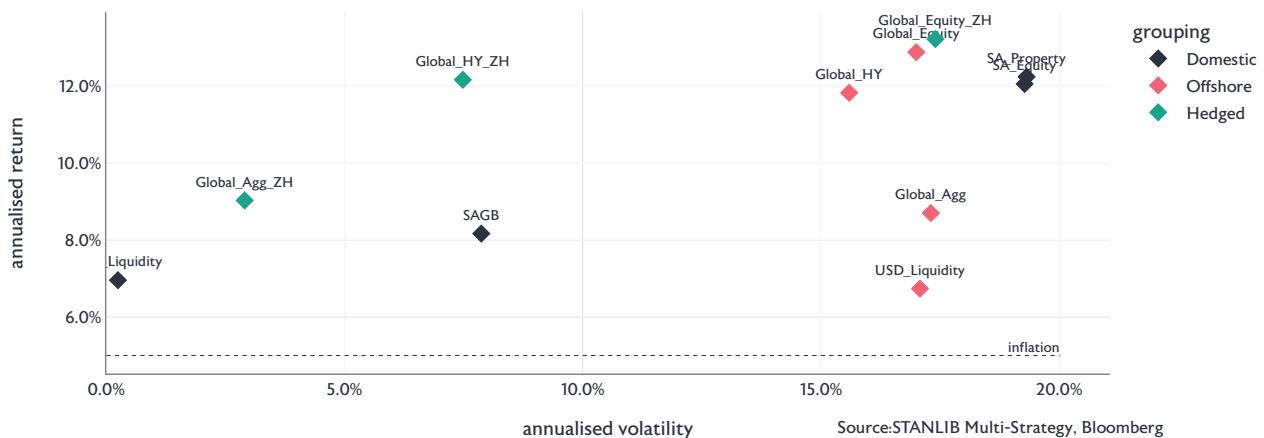


Figure 1: Historical returns and volatilities.

This is an interesting chart, which offers the following insights:

1. SA has experienced positive real cash rates (ZAR Liquidity point on the far left), meaning that a rand investor could earn inflation-beating returns with little or no risk just by holding cash and money-market instruments.

2. Unhedged USD cash (USD_Liquidity) and Global Bonds (Global_Agg) sit in the bottom right of the chart, meaning that they delivered relatively poor returns for the volatility a South African investor would have had to endure, mostly due to movements in the rand. The rand's volatility is a massive issue: dollar cash in the bank is often considered a "risk free" asset, but for a rand investor the volatility of the rand dominates and it is as volatile as global equities.
3. Hedging the currency materially improves the risk-adjusted returns on offshore bonds for two reasons. Firstly, hedging dollars back to rands is a "positive carry" trade, meaning that one is paid to take the hedge, and secondly, hedging strips the excess volatility from the asset class, improving the diversification benefits. If this feels illogical to you, you certainly are not alone, and we have written a whole section on this later.
4. Global equities, either hedged or unhedged, achieved a higher realised return for slightly less volatility than SA Equity.
5. Over our historic window, the risk-adjusted returns on SA Equity are very similar to those of SA Listed Property.

With that historical perspective in hand, we now calculate the strategic asset allocations (portfolio weights) that would have maximised the annualised return for a given level of volatility. Finance types will recognise this as the 'efficient frontier'. To build these portfolios, we recruit a mean-variance optimiser (MVO) using minimal constraints⁷ including a 15% to 45% range for offshore allocation. These portfolios can be understood as the asset allocations of investors with perfect foresight, with the proviso that we allow them to allocate up to 45% offshore, the new limit.

We find that the optimal portfolio for an investor, using historical data, depends on two key factors:

1. Are they willing and able to hedge currency risk?
2. What is their targeted real return (or how much 'risk'⁸ are they willing to take)?

The result of the mean variance optimisation is unequivocal: investors who can strategically hedge the currency should maximise their offshore allowance at 45% and then hedge their currency exposure to a meaningful degree. We found that the optimal weight of unhedged offshore allocation rose from 10% to 30% in line with the investor's targeted return and appetite for risk.

Based on the results of a single historical optimisation, our rule of thumb would be to hedge half the offshore allocation, hedging a bit more in low-risk portfolios and hedging a bit less in more aggressive portfolios.



Figure 2: Mean Variance Optimisation using historical data, including FX hedging, grouped by region. On the right-hand axis we plot the efficient frontier, showing the modelled return for each volatility level. On the left-hand axis we show the aggregated weights for the corresponding volatility level.

⁷ In all cases, we set a floor at 2% in South African cash, which is for efficient portfolio management purposes.

⁸ Technically we are talking about volatility rather than "risk", although in this paper we use risk in the common investment parlance. Philosophically, we strongly argue that for most investors volatility and risk are not the same thing. Risk is your inability to meet future liabilities (of lifestyle goals) due to insufficient capital, whereas volatility is how uncomfortable you feel when looking at fluctuations in your investment account balance. Having cash in a current account which is being eroded by inflation is a large risk to your ability to buy "stuff" in 20 years' time, but you will hardly notice the risk on a week-by-week basis when you log in to internet banking.

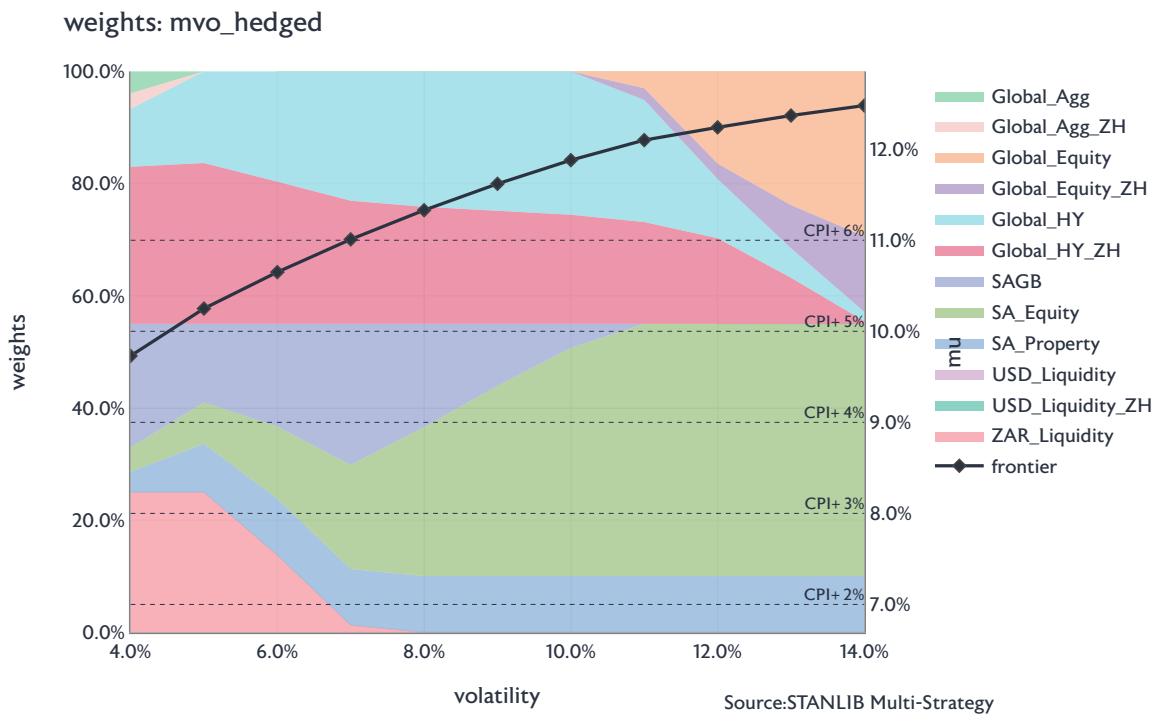


Figure 3: MVO using historical data, including FX hedging; grouped by asset class.

If an investor cannot or will not hedge some offshore currency risk, optimal offshore allocation rises in line with their required return and risk appetite. For those targeting a return of CPI+ 4%, an offshore allocation of 20% would have been sufficient to meet their return goal. To achieve a return of CPI +6%, the investor would have needed to allocate the maximum 45% of the portfolio offshore.

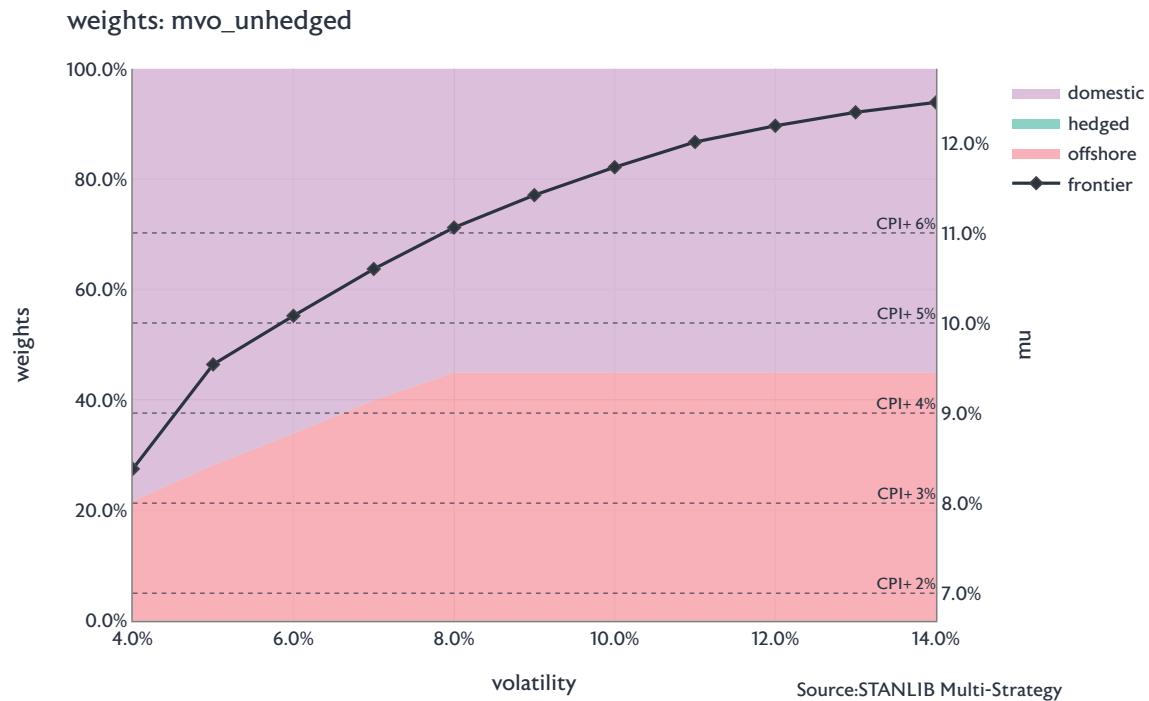


Figure 4: MVO using historical data, excluding FX hedging; grouped by region.

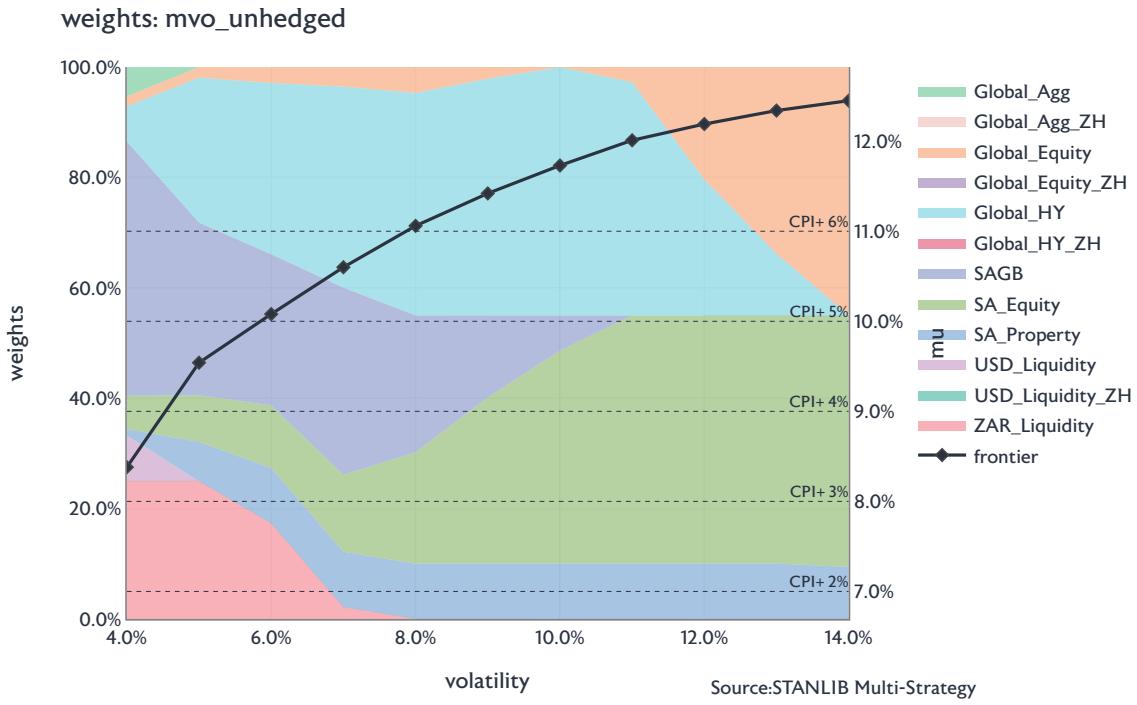


Figure 5: MVO using historical data, excluding FX hedging, grouped by asset class.

Diligent readers may have noticed from the charts above that the star asset class from a South African investor's perspective would have been Global High Yield bonds, whether hedged or unhedged. While we like the asset class, we cannot credibly look our clients in the eye and recommend they allocate 20-40% of their portfolio to High Yield bonds. As a sanity check we re-optimise, introducing a 10%⁹ cap on High Yield.

For the currency-hedged investor, this additional constraint makes no difference to the headline observation – you would still have been best served by maximising your offshore allowance. Where we observe a change is that the slope of the ratio of unhedged vs. hedged offshore assets has steepened, meaning more offshore currency hedging for lower risk portfolios, and less currency hedging for higher risk portfolios. Our 50% rule of thumb still holds for portfolios targeting CPI+ 5-6%.



Figure 6: MVO with high yield cap, including FX hedging, grouped by region.

The High Yield cap changes nothing for an investor who cannot hedge the currency. The optimal offshore allocation starts at about 20% for the more conservative investor, scaling linearly to 45% for investors targeting a higher real return.

⁹ This "feels" appropriate for a total return investor, given that high-yield credit beta tends to sit between equity and short duration bonds. When building Income or Absolute Return solutions, we consider the allocation differently, as the requirements change from pure return maximisation to considering a required yield and drawdown protection.

For one final test, we analyse what our historically optimal portfolio would have looked like over the last decade. We do not believe that the next ten years will necessarily look like the last ten, but it is an interesting exercise to see how our historical optimisation would have fared in such an extraordinary period.

Working backwards this time, for the unhedged investor the result, once again, does not change. This time the numbers suggest that the low-risk investor should have held around 25-28% offshore, scaling linearly for risk again to 45% for the investor who was targeting a result of CPI+ 6%.

When an investor can hedge, the initial result of maximising offshore allocations remains, but the unhedged proportion rises more quickly as the required return increases. This is a function of the outstanding returns of unhedged global equities compared with pretty much anything else. In this case the optimal hedged proportion of the offshore allocation declined from two-thirds for the ‘low-risk’ investor to zero for investors chasing a real return of CPI+ 8%.

Peeking into the Cloudy Ball: Incorporating our Prospective Returns

In the aftermath of the Global Financial Crisis (GFC) the global economy received an unprecedented level of monetary support, driving bond yields to multi-century lows and forcing capital outwards on the risk curve to find meaningful returns. I will leave the macroeconomic analysis to our Chief Economist, Kevin Lings, but from a market perspective we do not expect a resumption of the ‘Goldilocks’ conditions that risk assets have enjoyed since the GFC (or really since the 1980s).

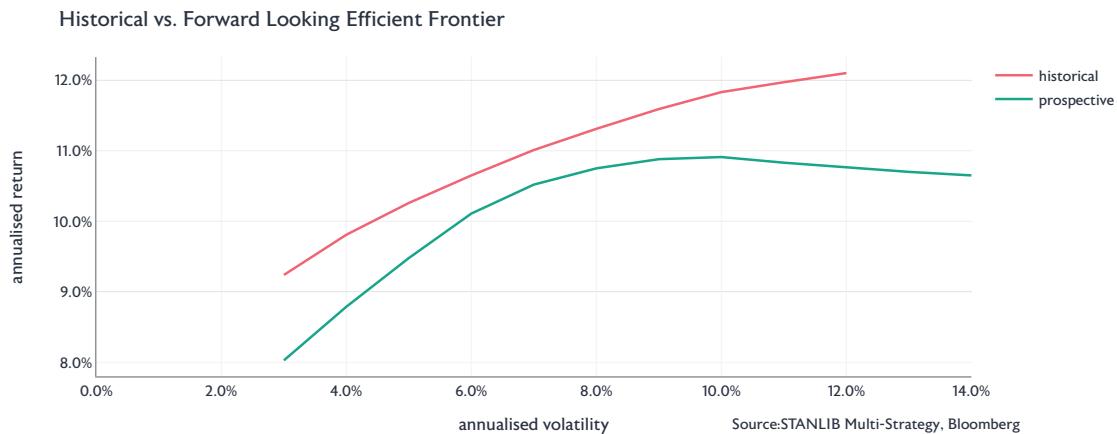


Figure 7: Comparing the historical efficient frontier to a frontier using STANLIB Multi-Strategy prospective returns.

Every month STANLIB’s Multi-Strategy team refreshes its view of prospective¹⁰ returns on various asset classes; which we use to help guide our strategic allocation. Spoiler alert: today we see forward returns far below those achieved over the last couple of decades!

The chart above compares our historical and prospective efficient frontiers, we were inspired to include it when told by a respected client [we paraphrase]

*“Everybody talks about lower expected returns, but it’s **much starker** when you see the frontiers”*

The following chart shows the returns and volatility that we expect from our ‘beta building blocks’ over the next ten years.

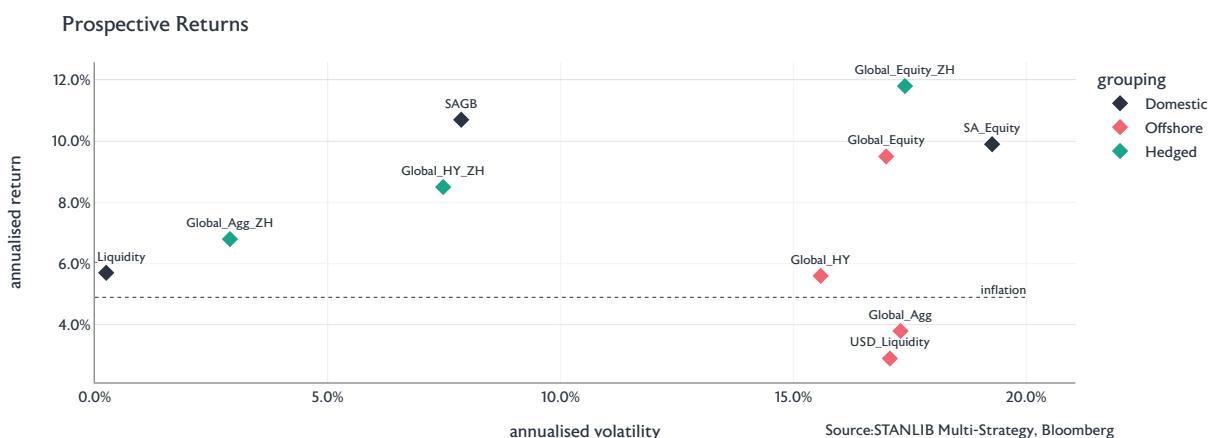


Figure 8: STANLIB Multi-Strategy Prospective Return and Volatility.

- Our first assumption is that South African inflation will average around 5% over the next 5-10 years, which combines market implied estimates, survey forecasts and central bank guidance.
- South African money markets have historically offered real returns of nearly 2%. We do not expect they will be so rewarding for investors in the future, but we still expect them to offer a real return of about 70 basis points per annum.

¹⁰ Our monthly prospective returns are built on a quant basis using risk premia building blocks – which helps us with consistency across the premia. For live portfolios we may incorporate update economic or market forecasts based on the long-term views of the team.

¹¹ This is not a new observation, it is one the team has written about on several occasions over a number of years.

- By far the most attractive asset class is South African government bonds. Assuming the government does not default on its obligations, we see a very strong risk-return profile¹¹ over the next 5-10 years.
- We think unhedged global bonds are relatively unattractive, offering negative real returns, although investors who can and will hedge the currency will be adequately rewarded for the risk they are taking. The backing up of US yields and spreads so far in 2022 has made global bonds a far more compelling strategic asset class than they were last year.
- Our expected forward return on SA Equity is approximately 10%, fractionally higher in absolute terms than unhedged global equity, although with higher volatility. Hedged global equity in our models looks like it will outperform both.

Those who were so inclined to read the box we entitled “*The Wonky Bit...*” will already know that, in our view, Mean Variance Optimisation can be a blunt tool. While the optimised portfolio is mathematically correct, it is subject to the bold assumption that we got our input forecasts right. If there is one thing I am 100% sure about, it is that our point forecasts are wrong.

Constraining the MVO to a range of 15%-45% offshore allocations and no less than a 2% position in rand cash, our efficient frontier for a currency-hedged investor shows a curious pattern, like an hourglass on its side.

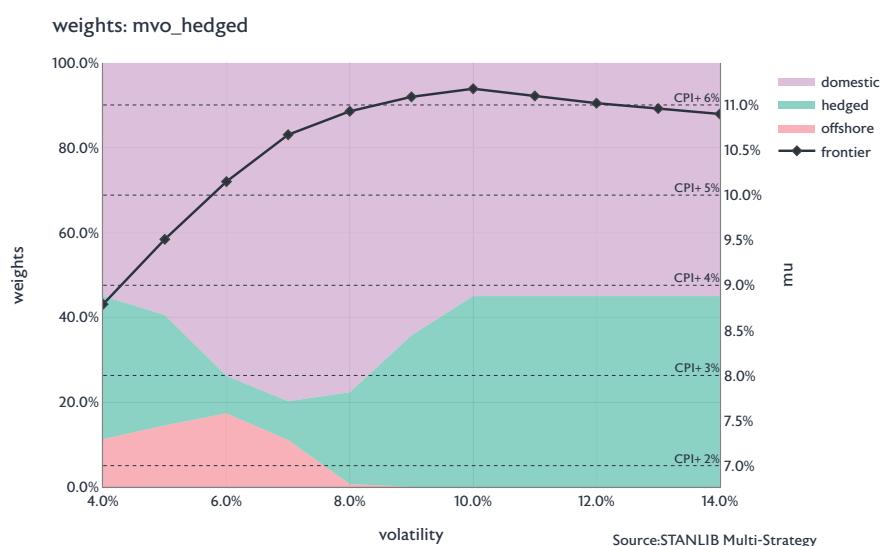


Figure 9: MVO using prospective returns, including FX hedging, grouped by region. We observe a quirky hour-glass shape which is caused by the strong prospective returns on South African Bonds on our measures.

“If we use prediction as the measure of a model, traditional finance makes precisely wrong predictions.”

Richard Thaler

That hourglass shape is explained by the conspicuous 11% prospective return on South African government bonds. If one could guarantee an 11% annual return with the low volatility of sovereign debt, why would one allocate to offshore bonds at all, let alone to riskier assets like SA Equity? I certainly would not, and I suspect you would not either. Unfortunately, nothing is certain in this world, and we need to remain aware of the risk that those attractive starting sovereign yields are not realised.

Below we show the results of a resampled optimisation, which applies a level of “uncertainty” to our return forecasts. This method diversifies away from SA Bonds, acknowledging the risk in the asset class. Our optimal range of offshore allocations narrows to between 32% and 40%, depending on target return, with approximately 25% to 50% of the offshore allocation left unhedged. By comparison with our historical heuristic, the forward-looking portfolios follow a similar pattern, but they do not fully utilise the offshore allowance and have a flatter ratio of hedged-to-unhedged in the offshore component.

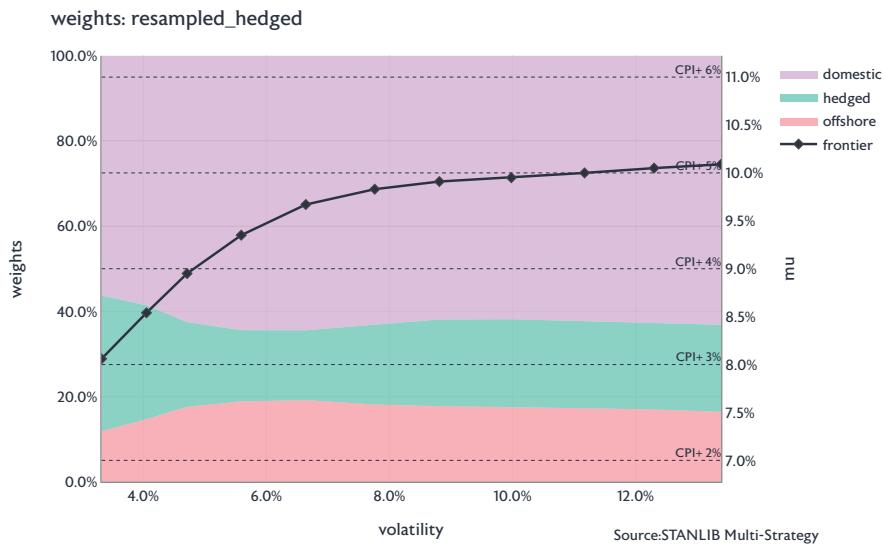


Figure 10: Resampled optimisation using prospective returns, including FX hedging. This methodology forces diversification and smooths the weights by assuming “uncertainty” around our prospective returns.

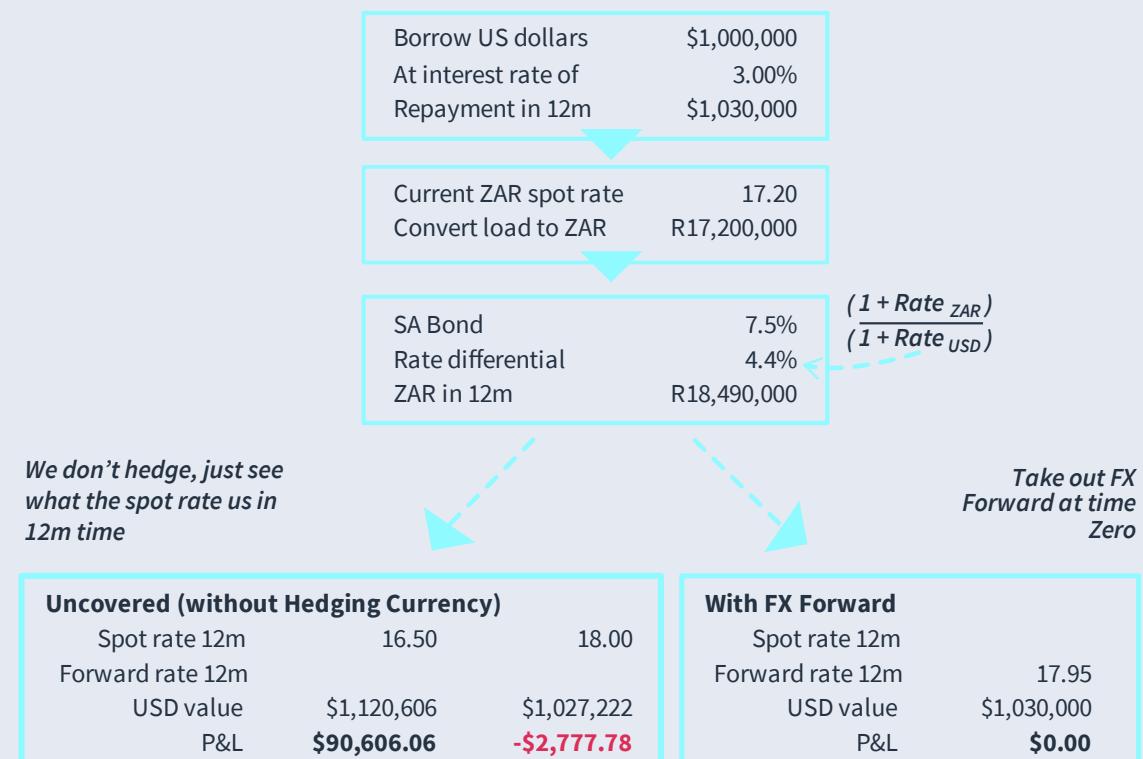
Finally, for the investor who is unable to hedge offshore currency risk, the results are a less extreme version of what we saw in the historical analysis, showing an offshore allocation of between 25% and 35%, increasing in line with target return.

Hedging: Tailwinds on Choppy Seas

It feels like a truism to say that the rand always falls, but what seems to surprise most South Africans we talk to is that historically the rand has depreciated by less than the positive carry¹² that is available from hedging dollars into rands. Hedging offshore currency risk is a big topic, and comprehensive detail is beyond the scope of this paper, but we can familiarise ourselves with the broad strokes.

Firstly, to remind ourselves what is going on with regards to the positive carry, we need to dust off our CFA textbooks and remind ourselves of the basic mechanics of “covered interest rate parity”. The principle is that you should not be able to make an arbitrage (risk-free) profit from borrowing in one currency, holding cash in a second and having hedged the currency risk from the start. Flipping the principle on its head, we mean that the “cost” of hedging a currency can be approximated by the difference in the short-term interest rates between two currencies. It is paid as a cost by the low interest rate currency and received as a “carry” by the holder of the higher interest currency. A worked example is in the box.

We borrow \$1 million from a US bank at an interest rate of 3%, repaying capital and interest of \$1.03 million in one year. At an exchange rate of R17.2/\$, we have R17.2 million, which we invest in a 12-month South African bond earning 7.5%, giving us R18.5 million next year. We could risk it, as in the left box (Uncovered), hoping the spot rate goes in our favour, but the right box is what matters. Here we take a forward currency contract, meaning we lock in now the FX rate we will use when converting our dollars back to rands in a year’s time.



Covered interest parity means we should not be allowed to make a “risk-free” profit from such a trade – if such an opportunity existed, an arbitrageur would step in and remove it. Thus, the fair value of the forward rand is the level which takes our trade P&L to zero! In our example, where we started with the rand at 17.2/\$, the fair value for the forward is 17.95/\$, which is an adjustment from our starting rate of 4.4%, or the differential¹³ between the 7.5% on the SA bond and the 3% cost of borrowing US dollars.

¹² A South African holding dollar assets is paid what is essentially interest to sell dollars back to rands. It is often underappreciated that the interest earned has typically been higher than the amount by which the rand has depreciated.

¹³ You may wonder why the differential is 4.4% rather than 4.5% (=7.5% - 3%). The answer is that we need to use the geometric differential rather than the arithmetic, where the geometric is given by $((1 + \text{rate}_{\text{ZAR}}) / (1 + \text{rate}_{\text{USD}}))$. Proving the difference is beyond the scope of this paper, but there are lots of good sources on the internet, or you could easily rebuild our example in Excel and prove it to yourself.

Earning carry helps, but if we imagine our returns as a boat, while the carry is a helpful tailwind, we are still sailing on the volatile sea of the rand. To highlight this, in the following charts we proxy¹⁴ the carry using the difference in one-month JIBAR vs LIBOR (or SOFR), and we compare that with the monthly spot return on the rand. Over the long haul, we find the monthly excess return is positive, and the distribution is negatively skewed (remember that, counter-intuitively, negative skew is a good thing for investment returns with more weight on the upside than the downside).

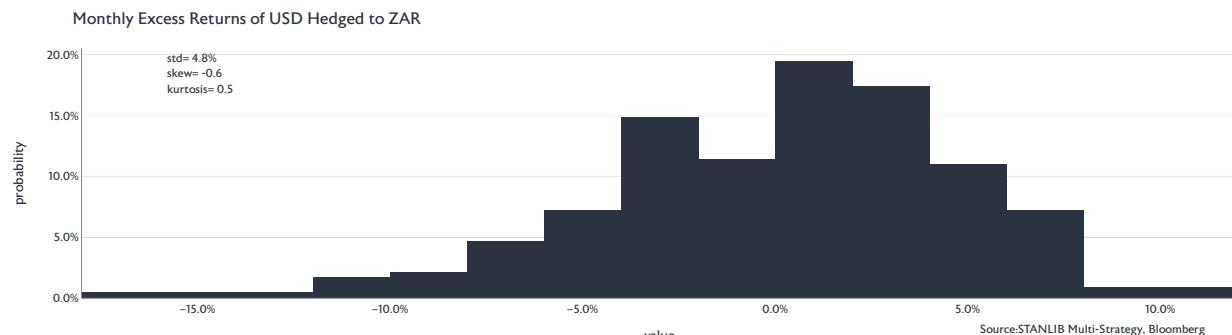


Figure 11: Histogram of monthly excess returns of hedged vs. spot rand.

The following heatmap shows whether a hedged currency strategy made, or did not make, money in any month. About 60% of the time the strategy had a positive excess return¹⁵. Scanning the heatmap reveals no obvious pattern in the excess returns, supporting our intuitive conviction that the benefits of hedging are volatile but are structural, rather than characteristic of a particular period of time.

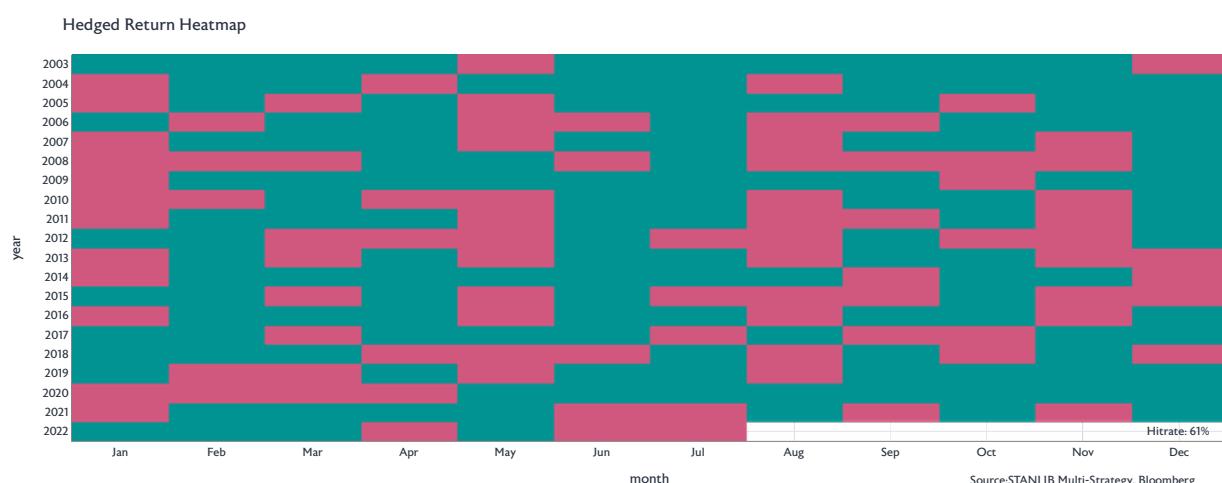


Figure 12: Heatmap by month of excess returns of hedged vs. spot rand.

Looking at our results, we are regularly asked why currency hedging matters less for equity than bonds.

The simple reason that currency hedging makes such a difference to the risk-adjusted returns of global bonds is that the volatility from the rand (circa. 17%) dominates the volatility of the underlying dollar bonds (circa. 3%).

¹⁴ A truer picture would be painted by inferring the carry from the “FX forward points” which are closer to the rates achievable using an FX forward. We know in the real world we are not really borrowing at SOFR or earning JIBAR, and that there is an arbitrage channel based on competing uses of capital. We prefer the proxy for this style of analysis because month-end levels of JIBAR and SOFR are less affected by market vagaries. Looking at the forward points, we find that where there is a bias it tends towards higher carry, not less, making any positive cases for hedging stronger than we have shown in these numbers.

¹⁵ Hit rate for just the rand, without hedging, over the same period was about 52%.

The bond's final return in rand terms is therefore almost entirely driven by movements in the rand rather than the dollar-based return of the bond itself.

Meanwhile, the volatility of global equities is about 15%, comparable to that of the rand. Thus, each price movement should be, on average, of a similar magnitude, so from a return perspective they could be mutually reinforcing or vice-versa. If we existed in an uncorrelated, random-walk world, the expected excess return would be close to zero, but we do not. The rand is often regarded as a proxy for global risk appetite and trades as such. The rand's 'risk on' reputation is justified when we regress¹⁶ the monthly spot return of dollar-rand vs. the return on global equities (chart below).

That combination of the higher intrinsic volatility and the underlying correlation between the rand and global equities explains why the currency hedge matters less for global equities than for global bonds.

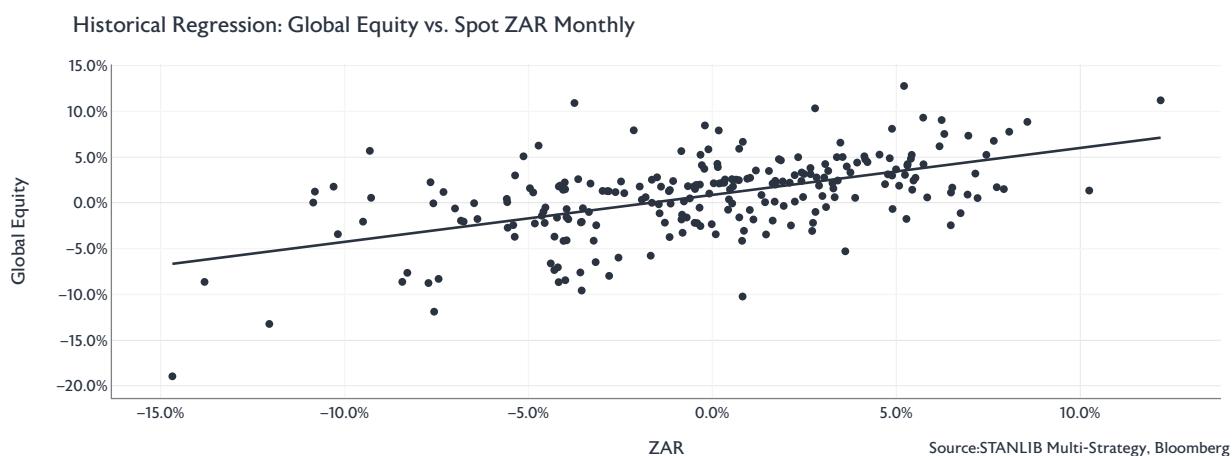


Figure 13: Monthly regression of the returns on spot rands vs. Global Equity in US dollars.

If you want to ignore another of my metaphors then feel free to skip the next couple of paragraphs. Those who want to know how I can link volatility to a rugby club night out, read on.

Imagine a rugby club annual dinner. At the end of a long evening, pairs of players compete in a blindfolded three-legged race¹⁷. If the 125kg prop gets tied to the 70kg scrum half, the practical reality is whether our scrumhalf steps a bit left, or a bit right on any given step. The pair's destination will be wherever our blindfolded prop takes them. In this case, our prop is obviously the rand, and the scrumhalf is the offshore bond.

In the next race, our prop gets tied to a second row¹⁸, a gentleman of similar weight. Without any communication they randomly push and pull one another to the left and right, weaving down the course but ending up travelling in a roughly straight line. Finally, they grunt every few steps, which (being forwards) is their version of communication, and can usefully represent correlation in our example. They get to the end of the pitch still with some volatility, but in a marginally straighter line than before.

Wrapping up our whistle-stop tour of offshore forex hedging, we are often asked an important question which is hard to answer. Given SA's fiscal backdrop, can we reasonably expect the hedge premia to continue? The short answer is 'yes' but it deserves some qualification.

Nominal rate compression (i.e. lower interest rates) means the headline carry from a hedge is lower than it has been through most, although not all, of history. However, arguably, the metric we should consider is the real interest rate differential. Using a lazy approximation¹⁹ of real rates, the current interest rate differential is relatively wide from a historic point of view. This should normalise, unless inflation in the US remains persistently higher than in SA.

¹⁶ Using monthly data from Bloomberg for MSCI World NTR in US dollars and rands, we find an R² of 0.3.

¹⁷ We are stretching the bounds of credulity with the blindfold three-legged race, but replace blindfold with tequila and race with stumble and we have a less family-friendly, but more realistic, rugby-based example.

¹⁸ A small reward for anyone who has stuck with me thus far – a poem by Mick Collis, the ex-Australian forward, entitled "My Son the Back" https://www.youtube.com/watch?v=zbYAyalgjsE&ab_channel=mickcolliss66

¹⁹ There is heated debate about the appropriate deflator for interest rates. The lazy option, which we show here, simply subtracts the contemporaneous year-on-year inflation. For more comprehensive analysis, and in our models, we used a smoothed blend of Core and Headline CPI, survey data and "market implied" metrics where we feel they are robust.

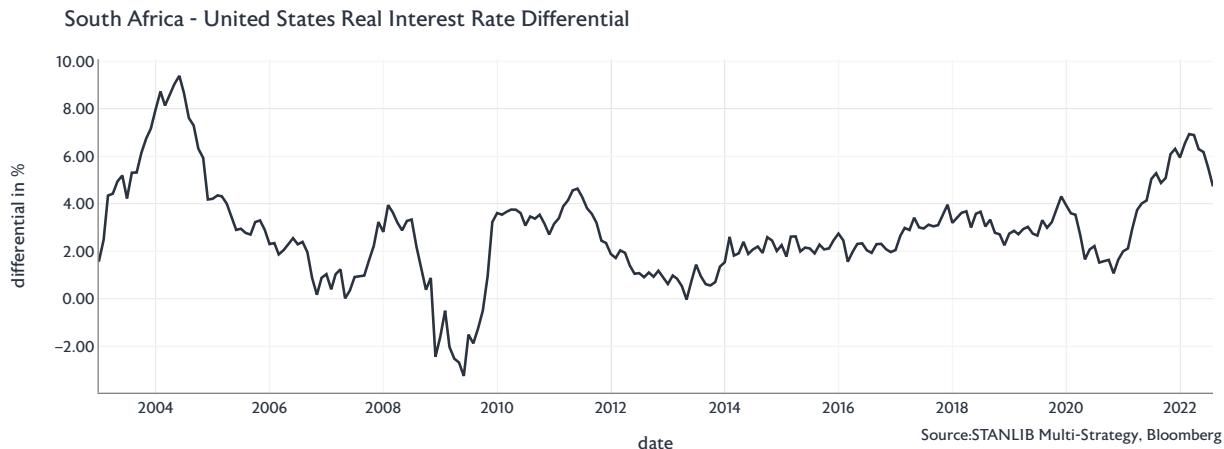


Figure 14: Short term interest rate differential between SA (using JIBAR) and the US (using LIBOR / SOFR).

The more contentious discussion is expectations around real²⁰ exchange rates, which historically, across currencies, tend to mean revert²¹. On a real exchange rate basis, the rand looks slightly cheap, which would imply that on the current strategic investment horizon it is not unreasonable to expect a very mild appreciation in the real rand.

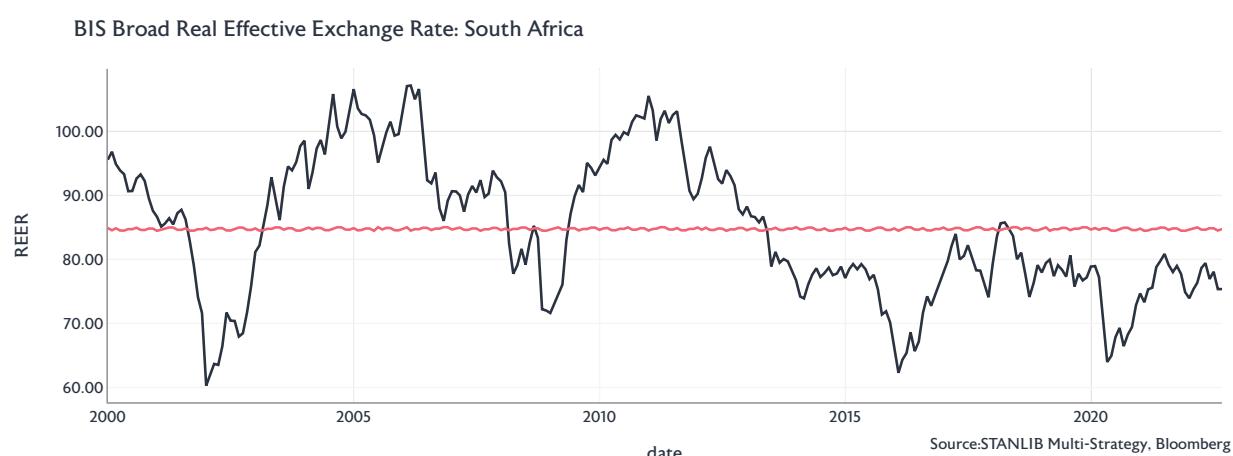


Figure 15: Real, inflation-adjusted, rand exchange rate with the long-term trend.

For the hedged offshore component to add value to a portfolio, we do not necessarily require mean reversion of the real exchange rate. Practically, we could see a mild weakening of the inflation-adjusted currency and still benefit at the portfolio level. However, if you are forecasting a complete collapse, or a multi-year decline in the real exchange rate, then the tailwind from hedge carry is going to be a fan pointing into a gale.

²⁰ Here we use the BIS Broad Real Effective Exchange Rate series.

²¹ See (Ca' Zorzi & Rubaszek, 2018)

Conclusion

We have deliberately presented our work in the domain of strategic allocations because we do not want tactical considerations over the next few months to distract us from our view on a multi-year timeline. We have also written this paper as food for thought for investors wrestling with the implications of the new Regulation 28 limits, but our analysis equally applies to any South African multi-asset investor. We can, and do, develop our own strategic solutions using the same framework, but with differing asset universes, constraints, and investment objectives. Whether looking back through time or forward, our conclusion is that an investor's optimal offshore allocation is a function of two variables:

1. The investor's ability and willingness to hedge currency risk
2. Their required return

For a generic investor who is able and willing to hedge currency risk, our analysis suggests that they should utilise most (if not all) of their 45% offshore allowance and then hedge some of the currency risk. More conservative investors should hedge more of the currency risk, while more aggressive investors (i.e. those with higher targeted returns) should hedge less. Historical data suggests that mitigating currency volatility improves risk-adjusted returns for South African investors across all offshore asset classes but particularly in less volatile instruments like global bonds.

Our rule of thumb is that the optimal level of currency hedge reduces as risk appetite increases. As targeted returns (risk appetite) increases, optimal portfolios tend to be increasingly weighted towards equities. As discussed above, the portfolio construction benefit of the currency hedge is much less for equities than bonds (remember the three-legged race?).

For investors who are unable or unwilling to hedge currency risk, offshore allocation should simply correlate with risk appetite, rising from 20% for the most conservative investors to the maximum of 45% for the most aggressive. The difference is that South African cash has provided positive real returns with zero volatility (by definition), meaning a conservative South African investor could grow the real value of their savings with very little risk. We expect positive real returns on cash to persist, but at lower levels than we have seen in the past.

Simply put, the greater an investor's targeted rate of return, the greater their offshore allocation should be, and if they can hedge the exchange rate risk then they should do so.

Final Considerations and Q & A

The work we have presented is anchored on strategic asset allocation at an index level, which would be the base of the pyramid in a hierarchy of investment needs, but it is not the “be all and end all”. Over the last few months, we have fielded questions and feedback from interested peers and clients.

Below we provide short answers to some of those questions, although our thoughts are largely intuitive rather than the product of supplementary analysis. We reserve the right to change our minds as we dig deeper.

What about non-Regulation 28 portfolios?

We've framed our paper around the new Regulation 28 limits, which reflects the genesis of the question, but our observations hold for any South African portfolio when considering a strategic horizon. The Multi-Strategy team applies the same process and thinking to our non-Regulation 28 portfolios.

What about the higher domestic alpha opportunity?

For equity, the argument broadly runs that there are fewer analysts covering domestic names, therefore there is an informational advantage to be gained by local analysts. That is a case we fundamentally agree with. However, to capitalise on this for clients we need to seek sustainability of that alpha, which is harder to find.

The same argument must also hold for any under-covered market, which means the same opportunity exists offshore in active emerging market funds and, more recently, in small- and mid-cap strategies, where changes in bank capitalisation rules and regulation (like MiFID II in Europe) have had the same effect.

As a counter, we did some preliminary research into the drivers of equity market returns. The chart below uses a Principal Component Analysis (PCA) of the last five years of weekly excess returns (stock returns minus benchmark returns) of direct stocks²² in SA and developed markets. It shows that in SA nearly 60% of the variance of excess returns could be explained by a single factor, whereas for developed markets the primary factor explained 40% of the variance. We infer from this that, although there may be less information efficiency in the South African market, which offers an opportunity to a good active manager, more than 80% of the variance is driven by only two factors.

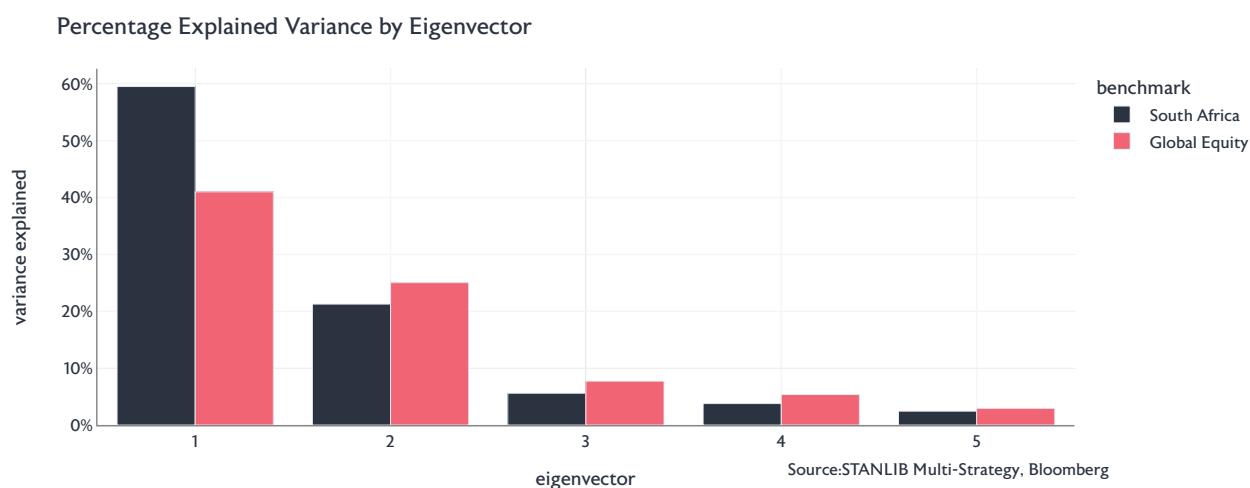


Figure 16: Percentage of the variance of excess returns which is explained by the first 5 “principal components” for SA and Global Equity. This demonstrates the percentage of alpha which can structurally be explained by each factor.

²² Source names and weights were taken from MSCI SA and MSCI World, which are also used as benchmarks for calculating excess returns. For global equity we limit the number of names to the top 150 by free float in US dollars, to roughly match the number of names in the JSE All-Share. Results do not vary much with larger or smaller subsets of global equity.

How stable are your proposed allocations? Is there anything today, given forward-looking expected returns, that might change in, say, six months' time?

Our proposed allocation incorporates everything we know about the world today and is designed to serve an investor who wishes to deploy capital and then sail round the world for a couple of years. The world is an uncertain place, however, and the last couple of years have taught us a hard lesson about nature's ability to blindside us. If we wake up to another new reality tomorrow, then our recommended asset allocations will change accordingly.

As a rule of thumb, we think investors should review their strategic asset allocations at least annually, although we do not necessarily think it is appropriate to update long-term weights each year. 2022 has been an extreme year. For example, we have seen credit spreads and bond yields move from historic tights (meaning low yields) to more normal levels, improving our perception of the high yield asset class from very unattractive to something close to fair value for a strategic portfolio. This is a big enough change to justify an adjustment to our strategic allocation.

How much conviction do you have in the SA bond expected return?

There is a comprehensive²³ body of work showing that the starting yield is a very good approximation of the expected return of a government bond. This is partly explained by a weird quirk of bond maths, which states that using the starting yield as an approximation holds even if yields blow out, because short-term capital losses end up being recouped by earning higher interest coupons in subsequent years. The short version is that a yield on the 10-year government bond of more than 11% makes us comfortable with our SA bond expected return.

That said, most bond research is based on developed market sovereign issuers who do not face default risk and have ostensibly no risk of being locked out of parts of their curve. Sadly, SA does not really belong in that company. Although we are not uncomfortable with SA bonds strategically, we prefer to diversify some of that risk offshore.

What impact do you think the change to Regulation 28 will have on domestic assets?

This is obviously just speculation, but our analysis suggests that South African managers will eventually increase their offshore allocations. The process may play out over years rather than weeks, as managers need time to hire additional global expertise and/or build offshore partnerships.

Following the intuition from our optimisation results, we anticipate that more managers will fund their increased offshore allocation positions from domestic equity rather than domestic bonds – good news for National Treasury, but less so for the JSE. Presumably, over the longer term, some dual-listed international names could delist from the JSE as South African capital becomes less captive.

²³ See (Bogle & Nolan, 2015) for an overview and (Leibowitz, Kogelman, & Bova, 2014) on bond math's.

References

- Bogle, J. C., & Nolan, M. W. (2015).
Occam's Razor Redux: Establishing Reasonable Expectations for Financial Market Returns. *The Journal of Portfolio Management*.
- Ca' Zorzi, M., & Rubaszek, M. (2018).
Exchange rate forecasting on a napkin. *ECB Working Paper Series*.
- Ledoit, M., & Wolf, O. (2004).
Honey, I Shrunk the Sample Covariance Matrix. *The Journal of Portfolio Management*.
- Leibowitz, M. L., Kogelman, S., & Bova, A. (2014).
Forward Curve Shifts and Return Convergence. *The Journal of Portfolio Management*.
- Michaud, R. O. (2004, January).
Why Mean-Variance Optimization Isn't Useful. *New Frontier Newsletter*.
Retrieved from <https://www.newfrontieradvisors.com/media/1184/why-mv-optimization-isnt-useful.pdf>