A Methodical Approach to Retiring Strategies

Rasheed Sabar
Managing Director
Ellington Management Group

September 2015

1 Introduction

From a bird's eye view, the capital allocation process consists of the following:

1. Form beliefs about the risk-adjusted returns of managers
2. Allocate capital to managers on the basis of these (and other) beliefs
3. Observe the managers' go-forward return outcomes
4. Modify beliefs based on observed outcomes, and restart capital allocation loop

The aim of this paper is to think critically about Step 4: how exactly should observed return outcomes update an investor's beliefs? And how should these updated beliefs translate into a **practical decision to increase or decrease capital?**

We offer investors a methodical, quantitative framework for addressing these questions. Oftentimes investors rely on heuristics when deciding whether to cut capital (e.g., 'cut capital if the manager loses more than X%'). These heuristics are useful but leave important information on the table, and can result in actions that are either too early or too late.

The actual decision to increase or decrease allocated capital (to a manager, strategy, or asset) will ultimately be holistic rather than mechanical. It will take into account the environment, relative performance, the portfolio's other holdings and mandate, etc.; nevertheless, we believe our simple quantitative framework can provide a helpful guidepost. We employ a similar framework in our own investment process to help determine when a strategy should be upsized or downsized.

This exposition deliberately avoids jargon and assumes little prior knowledge. We begin by building reader intuition for the core intellectual tool used to update beliefs given new evidence: Bayesian inference.

2 Bayesian Reasoning: A Revealing Example

In everyday life, we make probability judgments constantly, even if we don't conceive our decision-making in those terms. What is the likelihood I will enjoy reading the book, given that my friend recommended it? What is the likelihood I will get the promotion, given that my boss enjoyed my presentation? What is the likelihood we are compatible, given that my partner dislikes the Yankees?

Based on my innate preferences and accumulated life experience, I *already* have a view of how likely I am to enjoy the book (e.g. it is fiction and I don't normally enjoy fiction). My friend's recommendation is *additional information.* Is it enough to 'tip the scales' and make it worthwhile for me to read the book? What is the right way to quantitatively trade off my prior views with this new information?

Errors in judgment often arise from making this trade-off incorrectly. Consider the scenario below—which many parents encounter—relating to diagnosis of food allergies in babies. Repeated studies have shown that about 85% of doctors get the answer to this problem very wrong[1].

> 1% of babies have peanut allergies. 80% of babies with peanut allergies get positive skin-test results. 10% of babies without peanut allergies also get positive skin-test results. Now suppose a baby gets a positive skin-test result. What are the chances this baby actually has peanut allergies?

Most doctors reason that the baby's chance of peanut allergies is between 70% and 80%. In reality, it is about 7%! The 70%-80% number seems 'intuitive' if one doesn't distinguish between these two statements:

IF the baby has allergies, THEN the test shows positive 80% of the time. (1a)

IF the test shows positive, THEN the baby has allergies 80% of the time. (1b)

The problem states (1a) is true, and it can be tempting to conclude that therefore (1b) is approximately correct. To see clearly that this reasoning is flawed, consider the pair of statements below:

IF it is snowing, THEN it is winter 95% of the time. (2a)

IF it is winter, THEN it is snowing 95% of the time. (2b)

On the East Coast, (2a) is accurate, but (2b) is far from accurate: only a small minority of days in winter have snowfall[2]. Analogously, (1b) is inaccurate because only a small minority of positive skin-test results come from babies with allergies. This is because the test yields false positives 10% of the time on the large set of babies (99% of the population) who do not have allergies.

How do we correctly calculate the percentage in 'flipped' statements like (1b) from statements like (1a)? The key is Bayes' Theorem; it is a powerful tool for backing out 'the chance of X given Y' from 'the chance of Y given X' and a couple of other pieces of information. Specifically, letting Prob(X) denote the probability of X, and Prob($X|Y$) denote the probability of X given Y, the theorem shows that

$$\text{Prob}(X|Y) = \text{Prob}(Y|X) \cdot \frac{\text{Prob}(X)}{\text{Prob}(Y)} \tag{B1}$$

In the allergy problem context, X is having allergies and Y is a positive skin-test result. The problem statement tells us the quantities on the right hand side of (B1). We are told Prob(X), the percentage of babies with allergies, is 1%. We are also told Prob($Y|X$), the percentage of babies who get positive test results given they have allergies, is 80%. Lastly, Prob(Y), the percentage of babies with positive tests, is $(80\% \cdot 1\% + 10\% \cdot 99\%) = 10.7\%$, i.e., 80% positives from the 1% of the population with allergies and 10% positives from the 99% of the population without allergies. Putting these into (B1) gives a value of 7.4% for Prob($Y|X$), the percentage of babies who have allergies given a positive test result.

You can verify the arithmetic in (B1) works by imagining a specific number of babies, say 10,000. Then 100 babies (1% of 10,000) have allergies and the remaining 9,900 have no allergies. Out of the 100 with allergies, 80 get positive skin-tests. And out of the 9,900 with no allergies, 990 get positive skin-tests (10% false positives). So a total of 80+990 = 1,070 babies screen positive. Of these positives, only 80 actually have allergies. So, the percentage of babies with allergies among those who test positive is $80/1070 = 7.4\%$.

Bayes' Theorem thus provides a way to quantitatively combine *prior views*—e.g., 1% of babies have peanut allergies—with *new information*—e.g., this baby got a positive test result—to form *updated beliefs*—e.g., the chance this baby has allergies increased from 1% to 7.4%.

[1] See e.g. Casscells, W., Schoenberger, A., and Grayboys, T. (1978): "Interpretation by physicians of clinical laboratory results." New England Journal of Medicine 299:999-1001.

[2] Less than 10% of winter days in Manhattan get snow.

3 States of the World

In the allergy problem, there are only 2 possible *states of the world*: the baby either has allergies or she doesn't. Our belief prior to observing the positive test result is a probability of 1% on the former state and 99% on the latter state.

Let's now move to a portfolio management setting, where there will inevitably be more than 2 states. Suppose we are allocating capital to a strategy[3] and believe its go-forward Sharpe Ratio[4] is between -2 and 2. We conduct investment diligence and this leads us to form the beliefs shown in Figure 1. Specifically, we believe there is a 60% chance the strategy is a Sharpe of 1, a 15% chance it is a Sharpe of 0, etc. Note that the weights on the 5 Sharpe Ratio states[5] add to 100%.

Figure 1: Prior Beliefs For Strategy Sharpe Ratio

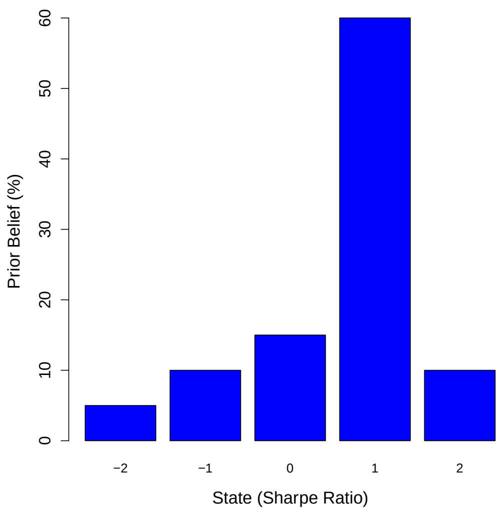

Figure 2: Strategy Returns Post-Allocation

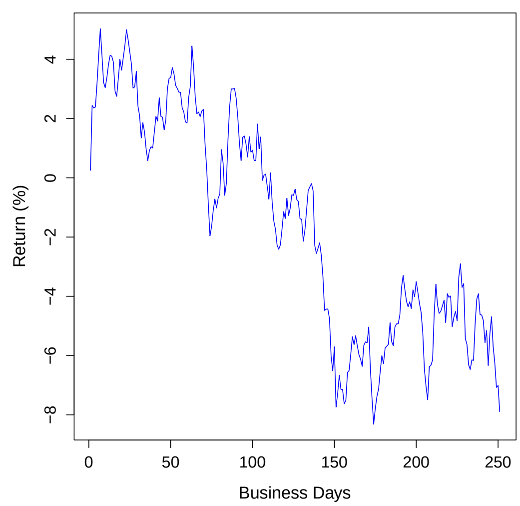

Having formed these beliefs, we allocate capital to the strategy and monitor its performance. One year later, the strategy—run at 10% annualized volatility—produces the returns in Figure 2. The strategy loses 8% and realizes a Sharpe Ratio of -0.8 over the sample window.

Is this window long enough to make definitive conclusions? How should the original beliefs in Figure 1 be revised given the post-allocation performance in Figure 2? Intuitively, because the strategy loses money, there should now be less weight given to Sharpes of 1 and 2, and more weight given to the lower Sharpe Ratio states. The weights in Figure 1 can be updated by utilizing the **exact same procedure used for the allergy problem**—Bayes' Theorem (B1). The resulting updated beliefs are shown in Figure 3.[6]

The shifts in weights from Figure 1 to Figure 3 are mostly intuitive. There is almost no chance now that the strategy is a Sharpe of 2. There is a 30% chance it is a Sharpe of 0 and over 25% chance the strategy is a Sharpe of -1. This makes sense given it realized a Sharpe of -0.8 over 1 year. The weighted average of the Sharpe Ratios is -.05; we therefore believe the strategy is a money loser going forward.

At first glance it may seem surprising that a Sharpe of 1 in Figure 3 is assigned greater than 30% weight, given the strategy's negative performance. This occurs primarily because we originally assigned a large 60% weight to the Sharpe Ratio 1 state (see Figure 1). If, instead of the weights in Figure 1, we had assigned an equal 20% chance to each of the 5 states, then our updated beliefs would be those shown in Figure 4.

[3] Everything we say applies equally well to allocation to a *manager* or to an *asset*.

[4] This is the ratio of annualized expected return to annualized volatility. We ignore risk free rate.

[5] We could just as easily have divided the interval into more than 5 buckets (resulting in more states) or widened the interval.

[6] Please see Appendix for details of applying equation (B1) to get Figure 3.

Figure 3: Updated Beliefs

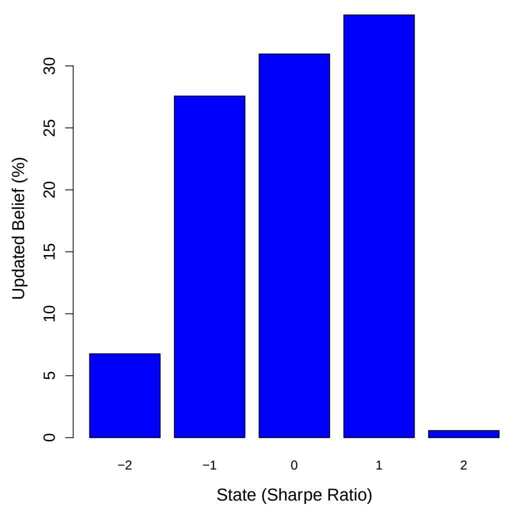

Figure 4: Updated Beliefs Given Equal-Weighted Prior

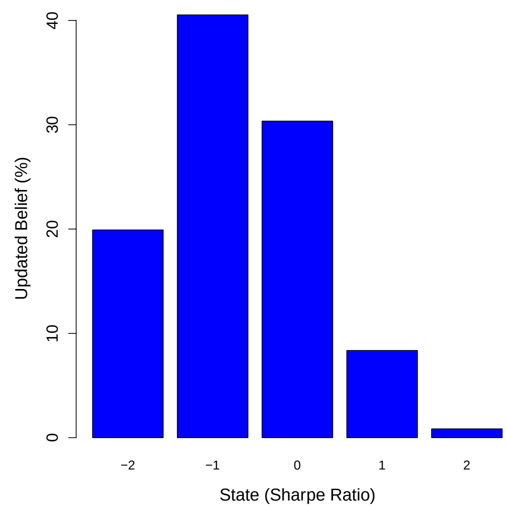

As Figure 4 shows, starting from equal-weighted prior beliefs across states, there is less than 10% chance the strategy is a Sharpe of 1. The largest weight is on a Sharpe of -1, which is close to the realized Sharpe of -0.8. The takeaway is not that the prior beliefs in Figure 1 are bad, or that placing 20% weight on the 5 states is better; instead, the moral is that *updated beliefs* are sensitive to the specification of *prior beliefs*. Ultimately, updated beliefs combine prior beliefs with realized performance. The longer the realized performance window, the more that updated beliefs will align with realized results; and the shorter the window, the more they will align with prior views.

In practice, then, we need a robust method for specifying prior views, one that minimizes psychological biases and arbitrariness. We also need to be able to quickly reduce risk in a strategy, without having to wait for a long realized performance window. We address both of these points in the next section.

4 Operationalizing the Framework in Practice

In the preceding section, we analyzed a simplified setup with 5 states of the world. Each state corresponded to a *static* Sharpe Ratio. In reality, Sharpe Ratios *change* over time. An initially high Sharpe strategy can decay to a lower Sharpe as more traders exploit the same signal[7]. For example, daily mean reversion strategies in equities were high Sharpe until 2010, when more capital deployed and lower barriers to entry (due to better technology) caused a structural decay in Sharpe.

To reflect potential alpha decay, we can make our states dynamic rather than static. Figure 5 shows states with Sharpe Ratios that change over time. In the Figure, state S_1 is a constant Sharpe of 1.5 across time, static like the 5 states in the previous section. In contrast, the other two states are dynamic. State S_2 is a Sharpe of 1.5 for a period of time (1 year), after which it decays to 0 instantaneously and stays there. State S_3 is initially a Sharpe 1.5 and decays to 0 exponentially (with a 1 year half-life).

With such dynamic states, we can detect decay in realized performance *much more quickly*[8]. The framework gives us wide latitude in specifying states of the world, so long as the prior weights on the states sum to 100%. This then raises the practical question: if we include dynamic states like S_2 and S_3 (and potentially dozens more) for a given strategy, how should we select prior weights for these states? The number of required decisions is now much larger, and so is the room for arbitrariness.

This problem of selecting states and weights for a strategy can be transformed into the much simpler problem of *categorizing* the strategy. Specifically, we can identify base strategy types and choose natural

[7]That's why a robust research process that regularly produces new signals is key to avoiding alpha decay.
[8]The calculations in the Appendix can be readily modified to accomodate dynamic Sharpe Ratio states.

Figure 5: States With Time-Varying Sharpe

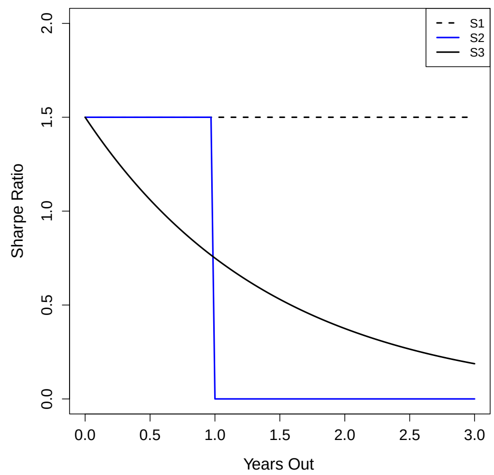

dynamic states for these base types. Then, whenever a new strategy comes along, we can map it to one or more base type and have it inherit the corresponding states and weights.

Figure 6 shows the particular strategy types we use in our implementation. Note the differences between, for instance, strategy types 1 and 2. Risk premia strategies—e.g., carry and value—tend to have lower Sharpe ratios and are usually slower to decay. In contrast, information-edge strategies—e.g., ones relying on proprietary data sources—tend to have higher Sharpe ratios and are susceptible to instantaneous decay at some random point in the future (when the proprietary data source becomes popular, for example). Each type, by its nature, has its own range of applicable states. We keep proprietary the exact configuration of states that we use; it has been calibrated by observing the behavior of hundreds of strategies over time.

Figure 6: Alpha Types

1. Structural Anomaly / Risk Premium	
Intuition	Makes money by providing liquidity or taking on risk correlated with market
Examples	Carry strategies; Put selling; Value strategies
Starting Sharpe	Ranges from 0.2 to 0.6
Decay Speed	Generally does not decay
2. Information or Friction-Based Alpha	
Intuition	Makes money via proprietary data, technology barriers, or trading in new market
Examples	Mobile phone geolocation data
Starting Sharpe	Can be above 1, with wide distribution based on exclusivity of access
Decay Speed	Decays abruptly to 0 at some future point (e.g. when data becomes public)
3. Modeling-Based Alpha	
Intuition	Makes money through superior modeling, analytics or choice of instruments
Examples	Cross asset-class lead-lag strategies
Starting Sharpe	Can be above 1, with wide distribution
Decay Speed	Decays exponentially to 0 with some half-life as other participants catch on
4. Data Mining	
Intuition	Appears to make money historically or backtested, but alpha is actually illusory
Examples	A technical strategy which backtests well due to data mining / over-fitting
Starting Sharpe	Centered at 0
Decay Speed	Decays immediately to 0

To see how this comes together in practice, suppose we are allocating capital to a new strategy. The strategy tactically trades small-cap vs large-cap stocks, long-short balanced. It is long small-caps most of the time and uses a mix of company-specific and macroeconomic data as drivers. Through investment diligence, we determine that the strategy is 50% Risk Premium (since it is mostly long small caps), 10% information edge (in using company-specific information), 25% modeling edge (in using macroeconomic information), and 15% data mining (potentially getting lucky historically given a small number of inflection points).

Determining these four percentages is all that is required to operationalize the framework. From there, the collection of states and the prior weights applicable to the strategy are implied automatically, as we have already configured dynamic states for each of the 4 canonical strategy types. We can now observe the strategy's realized performance and update our beliefs using Bayes' Theorem (B1), reducing risk as our updated beliefs dictate.

Note that mapping a strategy to base types not only makes it easier to select states and prior beliefs for that particular strategy; it also makes it easier to be consistent across multiple strategies, increasing objectivity. Ultimately, it makes operating the framework robust and interpreting its results intuitive.

5 Conclusion: Bayes Rule and Drawdown Rule

This paper focused on Step 4 of the capital allocation loop sketched in the Introduction: how to modify beliefs based on observed outcomes. We have seen that Bayes' Theorem is a powerful tool for updating beliefs, applicable as much to allergy probabilities given positive skin-tests as to Sharpe Ratio probabilities given realized performance.

There are two obvious obstacles to implementing the framework in practice. Specifically: (1) it can require a long realized performance window to identify deterioration, and (2) updated beliefs are sensitive to prior beliefs. We have shown how to address both of these via dynamic states and canonical strategy types. The result is a robust, consistent framework capable of quickly detecting alpha decay.

This framework complements heuristics such as drawdown rules. These rules are easy to interpret and do not depend on prior beliefs. However, they do not help with strategies that decay to flat without major losses. Moreover, they leave information on the table, unable to distinguish poor outcomes which are bad luck from those which strongly challenge prior expectations.

Ultimately, the capital allocator should leverage both a drawdown rule and Bayes' rule. The decision to retire a strategy can have significant impact on a portfolio's bottom line. Giving up on a strategy too soon will result in paying unnecessary fixed costs (e.g. time spent developing a new strategy or hiring a new manager) while acting too late risks further poor performance. Any tool that helps the allocator improve this tradeoff can drive significant portfolio alpha.

Appendix

We show how to apply equation (B1) to calculate updated beliefs in Figure 3 from prior beliefs in Figure 1 and observed strategy returns in Figure 2.

Let S_1, S_2, S_3, S_4, S_5 denote the 5 Sharpe Ratio states (-2,-1,0,1,2), and let $R = \{r_1, r_2, \ldots, r_{251}\}$ denote the 251 observed daily returns in Figure 2. We'll focus initially on updating our belief in S_1, the state of the world in which the strategy is a Sharpe Ratio -2.

We started from the view (depicted in Figure 1) that S_1 has a 5% chance of being true. Now that we've observed R, we need to revise this 5% number. In other words, we want to calculate $\text{Prob}(S_1|R)$. From equation (B1) we have

$$\text{Prob}(S_1|R) = \text{Prob}(R|S1) \cdot \frac{\text{Prob}(S1)}{\text{Prob}(R)} \tag{A1}$$

Observe that $\text{Prob}(R)$ is a fixed number that does not vary across the 5 states. Hence, we can write (A1) as

$$\text{Prob}(S_1|R) \propto \text{Prob}(R|S1) \cdot \text{Prob}(S1), \tag{A2}$$

where '$\propto$' means 'proportional to'. The probabilities of the 5 states add to 100%, i.e.,

$$\text{Prob}(S_1|R) + \text{Prob}(S_2|R) + \text{Prob}(S_3|R) + \text{Prob}(S_4|R) + \text{Prob}(S_5|R) = 100\%.$$

Thus, we can calculate the right hand side of (A2) for each of the 5 states and then rescale so the sum is 100%. Below we evaluate the quantities on the right hand side of (A2). We assume that daily returns R are normally distributed and independent across time. This assumption and the corresponding calculations below can be readily modified.

From Figure 1, we have $\text{Prob}(S1) = 5\%$. Next, the term $\text{Prob}(R|S_1)$ means the probability of observing R assuming that S_1 is true. In a world where S_1 is true, the strategy has an expected annual return of -20% (because it has a Sharpe of -2 and an annualized volatility of 10%). This translates into a daily expected return of -20%/251 = -.08%. Further, the annualized volatility of 10% translates into a daily volatility of $10\%/\sqrt{251} = 0.6\%$. Thus, in a world where S_1 is true, the daily returns R are normally distributed with a mean μ of -.08% and a standard deviation σ of 0.6%. From the normal density function, the likelihood of observing R is then

$$\text{Prob}(R|S_1) = \left(\frac{1}{\sqrt{2\pi\sigma^2}}\right)^{251} \cdot \exp\left(-\frac{1}{2\sigma^2}\sum_{j=1}^{251}(r_j - \mu)^2\right)$$

We repeat this calculation and plug the results into (A2) for each of the 5 states to arrive at the updated beliefs depicted in Figure (3).

Important Notes

The information in this document does not constitute an offer to sell or a solicitation of an offer to purchase any securities of any entities described herein. Any such offer will be made solely to qualified investors by means of private placement memoranda and related subscription materials. Investors should rely on such memoranda, and not this report, in connection with a purchase of any such securities. The information in this document should not be taken as a recommendation to purchase or sell any particular security.

The strategies described herein are speculative and involve substantial risks, including risk of loss.

Information in this document is as of the date indicated. The author assumes no duty to update the analyses contained herein.