
Making simulation and stress-testing more intuitive and insightful

An overview of existing methods and introduction of a novel approach

Dane Rook, Dan Golosovker

1. Introduction

Uncertainty is unavoidable in investing. To properly handle it, investors need tools to help them explore how sources of uncertainty — such as inflation, interest rates and market volatility (to name just a few) — impact the potential outcomes of investment decisions. Simulation and stress-testing (S&ST) methods are among the most powerful of such tools, and using them is a best practice.

Unfortunately, the dominant approaches to S&ST (including Monte Carlo analysis) are often opaque, burdensome to implement, constrained in the insights they deliver and challenging to communicate to clients. We're developing a novel technique, the History-Based Multi-Lens (HBML) approach that alleviates these shortcomings. We use HBML extensively in our own research, and we're currently building applications to allow investors to do so too. The aim of this piece is to give a high-level overview of existing S&ST methods (as they apply to investing), including their flaws. We also introduce our HBML approach, and describe how it improves on the most popular S&ST methods.

Simulation and stress-testing encompass a broad range of analytical techniques that all share a common goal: to assess how asset performance varies in response to changes in the assumptions that drive that performance. Such assumptions may include the future state of the economy, the scale and duration of a market crash, growth rates, dividends, the pacing of capital calls from



private equity funds or practically any other value-influencing variable whose future level isn't perfectly known. Often, the aim of S&ST isn't just to establish what's possible, but also what's likely — for instance, investors often turn to S&ST to explore probabilities of specific outcomes. Further, outcomes alone aren't the only thing to care about: most investors are also concerned about the paths an asset might follow in reaching a given outcome, and the riskiness of those paths.

To the extent that a particular S&ST method is able to fulfill these objectives, it can potentially be useful for various investment activities, such as:

- **Portfolio construction:** identify portfolios with ideal risk-return profiles
- **Risk management:** reveal both the types and degrees of risk that portfolios are likely to face under diverse market conditions
- **Back-testing:** investors study how portfolios might've performed in previous market environments
- **Scenario exploration:** investors envision the portfolio impacts of specific events (e.g., spiking unemployment, falling oil prices)
- **Wealth planning:** investors uncover the likelihood of meeting their risk-adjusted return objectives, and map these to client needs (such as funding education, retirement, philanthropy or other ambitions)
- **Client engagement:** investors better communicate the riskiness of particular investment choices to their clients, and more concretely explain the foreseeable impacts of future (or current) market events

Yet, despite the potential breadth of applications for S&ST, its usefulness remains hampered in practice by deficiencies of existing S&ST methods. To better understand those deficiencies, it's helpful to review the methods themselves.

2. Existing methods

S&ST methods can be divided into two categories: those that include a role for randomness (and are stochastic), and those that are deterministic. Most deterministic methods are simply chains of 'hard-coded' assumptions that must be specified explicitly, for example: the S&P-500 will fall by 2% for every 1% increase in inflation. To generate significant insights, a large number of these explicitly



specified assumptions are usually required, which poses two problems: 1) formulating so many assumptions can become burdensome; and 2) the likelihood of analytical error increases with the number of assumptions. These problems mean that most deterministic methods are both assumption-intensive and 'brittle,' in the sense that they rely tremendously on the quality of the assumptions that underpin them: if the assumptions are weak, the analysis crumbles.

Another shortcoming of existing deterministic methods is that they typically span only a small range of possibilities for outcomes, which means they have a hard time assigning meaningful probabilities. Said differently, most deterministic methods lead to under-exploration of possible outcomes, and therefore struggle to: 1) comprehensively depict risk; and 2) indicate the true likelihood of specific risks materializing. Stochastic methods partly bypass these flaws by assigning probability distributions to key variables, rather than single values. Random draws from these distributions can then be taken to generate prospective outcomes for analysis; with enough draws, it's possible to precisely characterize not just the range of possible outcomes, but also assign probabilities to those outcomes in a principled way, which can furnish answers to pressing questions, such as the odds of a portfolio underperforming its benchmark, or the range of returns that might be expected for a stipulated level of confidence.

However, stochastic approaches do little to resolve the problem of assumptions' intensity, and in many cases worsen it. For instance, while a deterministic method would require a single assumed value for the return on an asset, stochastic methods would require an assumption about the type of probability distribution for that return (normal, lognormal, hypergeometric and so on), as well as the parameters for that distribution (e.g., mean and variance, in the case of a normal distribution). These additional assumptions are the price of the extra analytical depth stochastic methods can provide; but whether that price is worth paying isn't always obvious: the brittleness of stochastic methods — in terms of sensitivity of their results to how accurate the underlying assumptions are — can be difficult to assess. That said, stochastic S&ST approaches are favored by many investors.

Among stochastic methods, Monte Carlo (MC) techniques are perhaps the most popularly used.¹ The basic mechanics of MC entail: generating large numbers of samples (usually thousands or millions) from explicitly specified probability distributions; performing various mathematical operations on those samples (such as applying correlations across assets, compounding returns and more); and then calculating statistical values (e.g., averages, variances) based on aggregate

¹ It's worth noting that 'Monte Carlo' is actually a very broad class of analytical tools, not a single approach. There are many variants, which differ substantially in their complexity.



results. When conducted at the portfolio level, MC for S&ST purposes requires specifying covariance structures in asset returns (that is, how returns between different portfolio assets are correlated). Generating assumptions for covariances can be difficult for many analysts, particularly because these can vary over time and under distinct market conditions.

Apart from the added labor of covariance assumptions, MC methods pose issues of transparency. Alongside the many explicit assumptions that are needed to use such methods, there are a host of embedded assumptions that must be properly understood in order to assess whether the method is valid for the desired analysis (these might be considered 'hidden' assumptions, in that a user implicitly adopts them by using the method). A major embedded assumption that lurks in basic MC approaches to S&ST is that there's no correlation in asset returns across steps over time, (such that random draws from probability distributions are independent).² In reality, this assumption is routinely violated, which threatens the realism of MC results and the quality of any insights that might be derived from them. In fact, this predicament holds true more generally: the realism of MC rests squarely on the quality and intricacy of its assumptions (both implicit and embedded), and more assumptions are required for greater realism; but added assumptions also increase the possibility (and severity) of error.

A further strain inherent in MC methods is their time intensity. These methods gain their power from volume, and routinely involve tens or hundreds of thousands (if not millions) of random draws and transformations to yield passable results. Unless one uses optimized algorithms or special hardware, a single MC analysis can take minutes or even hours to compute, which isn't ideal in circumstances when one is trying to quickly explore investment hypotheses. Moreover, it's often unclear how many random draws are necessary to get 'good' results. From what we've observed, many analysts embrace a more-is-better ethos when it comes to MC; however, having too many random draws can create improperly small confidence intervals (thus tempting analysts to be overly confident in results).³ Despite the number of draws being a consequential choice in MC analysis (hence, it can be seen as another influential assumption that must be explicitly made), guidance on how to choose that number is scant (or else too mathematically complicated to be accessible for most users).

² There are more sophisticated ('non-vanilla') variants of MC S&ST that allow for serial correlations in returns, but we've noticed these are less commonly used by most investment professionals, and those variants require intricate additional assumptions.

³ This problem seems to be under-acknowledged, both in research literature and among investment practitioners.



Another limitation of MC analysis comes by way of conditioning. While users of MC are free to specify whatever probability distributions (and parameters for them) they wish (bearing in mind that these choices will impact the quality of the analysis), in practice users often frequently choose distributions that resemble historical ones. This approach is hazardous (for MC, at least) when one uses historical distributions that are based on very short histories, which might be done, for instance, when one is trying to explore observed asset behavior under very specific economic conditions. These conditioned distributions may contain too little data to make MC results statistically valid (in such situations, users must either forgo S&ST analysis via MC, or willingly accept dubious results). Yet, despite all of these limitations, MC is often treated as the gold standard of S&ST, and many investors have clients who specifically request that MC analysis be performed.

Our reaction to this state of affairs has been twofold. First, whenever investors are asked by their clients to conduct Monte Carlo analysis, the investor should strive to understand what the client is truly seeking. In our experience, “doing Monte Carlo” is, within the investment community, a byword for deriving likelihoods and confidence levels for potential outcomes: when clients ask for MC analysis, more often than not it’s probabilities that they’re really after. But MC methods (and stochastic methods generally) aren’t the only technique that can generate probabilities and confidence intervals in a principled way — a realization that was key to our second reaction to the difficulties that MC poses for S&ST: inventing the History-Based Multi-Lens (HBML) method. It’s to HBML that our attention now turns.

3. The History-Based Multi-Lens (HBML) method

As indicated above, the steepest challenge that existing S&ST methods pose for investors lies in the intensity and quality of assumptions required. Formulating assumptions about capital markets (both in terms of the behavior of markets overall and the interrelationships among assets in them) is an intensive process that can consume considerable resources (e.g., professionals’ time, costs of research materials and fees paid to source assumptions from third parties). The motivation behind HBML was to lessen the ‘assumptions’ burden placed on investors. HBML fulfills that objective by sourcing assumptions directly from history.

HBML applies a rolling partition to history, thus splitting it into a series of episodes. For instance, one might choose to apply a 24-month window to a time-series of asset returns from 1950–2022. The first episode would run from January 1950 to December 1951, and the second would run from



February 1950 to January 1952, and the third would run from March 1950 to February 1952, and so on. Different window lengths can be used, with the caveat that ‘sample size’ (i.e., the number of episodes generated per time-series) decreases as window length increases.⁴ Applying rolling partitions directly to empirical time-series is what makes HBML history-based.⁵ The main advantage this confers is preservation of many subtle relationships that would otherwise need to be explicitly assumed, or else omitted. These relationships may include phenomena such as momentum effects in returns (whereby returns in one period are affected by returns in earlier periods), sensitivity of returns to changing macroeconomic conditions, and the correlation structure in returns across different assets. This final relationship is vital for accurate assessment of risk in portfolios (which, after all, is the chief motivation for doing S&ST), and is something that must be explicitly estimated in other S&ST approaches.

Notably, HBML is a deterministic S&ST approach.⁶ Yet it overcomes many of the deficiencies that plague other deterministic methods. For example, HBML can generate valid probabilities and confidence intervals for specific outcomes: all that’s needed is to calculate relevant frequencies and statistics across episodes. Figure 1 presents an example of simulated return paths for a ‘60/40’ portfolio under HBML, using a 12-month rolling window derived from historical returns for January 1970 through April 2022 (giving 616 12-month episodes). It shows simulated trajectories of top-quartile (i.e., 75th percentile), bottom-quartile and median returns — with 95% confidence intervals estimated for the median. (These intervals can be estimated through bootstrap resampling, to give a sense of how reliable the estimated median is, from a statistical standpoint.)

Per HBML, the probability of this 60/40 portfolio exceeding a 5.0% annualized return over a 12-month period is 60.2% (5.1%, at a 99% confidence level). This sort of analysis is both straightforward and quick under HBML — it takes just seconds to compute, whereas for even basic MC methods it can consume several minutes (even when using fairly efficient algorithms). Moreover, unlike MC, there’s no need to question the realism of simulated runs under HBML, because runs are directly extracted from events that actually occurred. Further, there’s no need for the user to judge how many runs are required (all the user needs to decide is how far back in history to look).

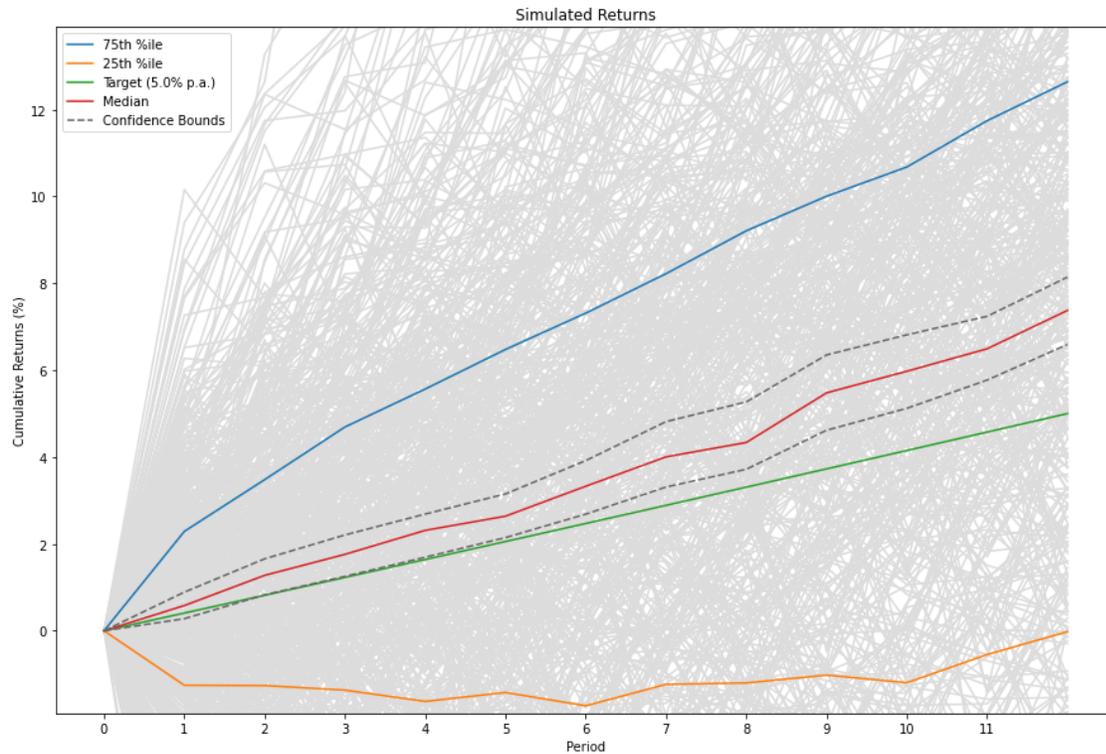
⁴ Similarly, one might choose to work with quarterly or daily data; the granularity of HBML is limited only by data frequency.

⁵ It’s common for stochastic methods to extract some of their assumptions from empirical data, but they mostly do so from statistical aggregations, rather than directly from time-series. For example, in MC analysis, a typical practice is to use the mean and variance of historical returns for an asset as assumed parameters for the probability distribution of its returns.

⁶ There are ways to enrich HBML with stochastic features, but the version we describe here is entirely deterministic.

Figure 1

Returns of a 60/40 portfolio over 12-month horizon, simulated HBML



Source: Addepar Research

Notes: Median annualized return across episodes is 7.37%, and mean is 5.96%.

Episodes account for the 'history-based' part of HBML. The 'multi-lens' component stems from being able to combine episodes with various perspectives on history, or *lenses*. One lens that's meaningful for most investors is economic regimes, such as whether inflation is rising or falling. HBML can readily answer questions on how portfolios perform in different economic regimes:

Figure 2 continues with the previous example of a 60/40 portfolio, and shows the probabilities of it exceeding a 5.0% annualized return over various horizons.

**Figure 2**

Predicted success probabilities for 60/40 portfolio under different inflation regimes

Regime	6 months	12 months	24 months	60 months	120 months
Rising	47.3%	51.2%	43.4%	36.6%	46.8%
Falling	65.6%	70.1%	68.5%	68.4%	75.7%

Source: Addepar Research

Notes: Results are based on rolling-windows of excess monthly returns on a portfolio that's composed of 60% stocks and 40% bonds. Regimes of rising inflation are those for which inflation was increasing during at least half of the months in the specified horizon. Regimes of falling inflation are those for which inflation was increasing during less than half of the months in the specified horizon.

Regime-based analysis is straightforward under HBML, which (as we noted earlier) isn't often the case with MC-based S&ST. And it's not only success probabilities under inflationary regimes that HBML is able to readily simulate: Figure 3 shows annualized returns estimated via HBML for regimes of high and low market volatility (again, for a 60/40 portfolio). This sort of analysis can be performed for any conceivable regime, as long as one is able to identify the regime that existed in a given period (e.g., whether inflation was rising or falling in December 2022) — which in many cases is uncomplicated.

Figure 3

Annualized returns estimated by HBML under different volatility regimes

Regime	Horizon				
	6 months	12 months	24 months	60 months	120 months
High volatility	2.3%	3.7%	2.4%	3.9%	3.9%
Low volatility	8.8%	8.5%	7.8%	8.1%	7.3%

Source: Addepar Research

Notes: Results are based on rolling-windows of excess monthly returns on a portfolio that's composed of 60% stocks and 40% bonds. Predictions for the high volatility regime come from rolling-window episodes that had top-quartile volatility,



historically. Predictions for the low volatility regime come from rolling-window episodes that had bottom-quartile volatility, historically.

It's worth noting how light HBML is regarding assumptions: practically no explicit assumptions are required (except for how far back in history to look), and the lone embedded assumption is that the future is likely to resemble the past. (And this is really only an 'assumption' if one uses HBML for predictive purposes, and not just exploring historical outcomes.) Because HBML has so few assumptions, it's a robust S&ST method: the opposite of brittle. Simulation quality under HBML stems from data quality, rather than the quality of assumptions plus data.

Another way in which HBML improves on existing S&ST methods is its ability to accommodate multiple risk lenses. Many MC methods are relatively constrained in the varieties of risk they're able to realistically capture. For example, because basic MC methods do a poor job of capturing serial correlation in returns, they're not suited to stress tests for drawdown risk. These methods can thus be described as 'monocular' regarding risk: they may do a decent job of reflecting one or two forms of risk, but generally can't realistically capture the complete spectrum of risks that investors might care about. HBML isn't constrained in this way: it enables a 'macroscopic' view of risk.

For instance, the drawdown status of a portfolio across episodes can provide a probabilistic view of drawdown risk: during a one-year period, our illustrative 60/40 portfolio has a 16.9% probability of being 10% or more below its starting value, at some point during the year. Also, one can easily include confidence assessments: in this example, the 99% confidence interval runs from 13.1% to 20.8% (i.e., $\pm 3.9\%$). Quite often, these are exactly the sorts of figures that investor's clients really want when they request "Monte Carlo analysis"; HBML can yield these figures, in a way that's far more transparent, robust and readily explained to clients. It's for these reasons that we believe HBML is a way to make simulation and stress-testing more intuitive and insightful, for both investors and their clients.



Figure 4

Comparison of Monte Carlo and HBML methods

Property	Monte Carlo	HBML
Treatment of randomness	Stochastic	Deterministic
Requires explicitly specifying returns distributions for individual assets?	Yes	No
Requires specifying covariance structures in asset returns?	Yes	No
Embedded assumptions	Many	Few
Requires determination of suitable number of runs?	Yes	No
Realism of generated paths	Not guaranteed	Guaranteed
Risk perspective	Generally monocular	Inherently macroscopic
Scale of run-time	Usually minutes or hours	Seconds or milliseconds
Ease of examining asset behavior under specific conditions (e.g., particular economic regimes)	Moderate to high difficulty	Low difficulty
Explainability	Challenging for non-technical audiences	Straightforward, even for audiences with minimal understanding of statistics



4. Conclusion

All investors face uncertainty. Simulation and stress-testing (S&ST) are best-practice methods for properly handling it. However, not all S&ST approaches are created equal. In this piece, we've attempted to highlight some of the most consequential shortcomings of popular S&ST techniques — namely, the intensity of time and assumptions they require.

We've introduced a novel alternative, the History-Based Multi-Lens (HBML) method, which both alleviates those deficiencies and performs on par with (and in many cases, surpasses) existing techniques, including Monte Carlo analysis (our performance comparisons are technical in nature, and we therefore reserve them for a forthcoming research paper). Yet, regardless of which S&ST approach an investor uses, it's vital that they understand what their chosen approach can actually achieve, how it does so and how well that aligns with what clients ultimately need.



Authors

Dan Golosovker, Deputy Head of Addepar Research (dan.golosovker@addepar.com)

Dane Rook, Sr. Research Director (dane.rook@addepar.com)



Notice and disclaimers

All information provided by Addepar, Inc. or its subsidiaries (collectively, “Addepar”), including without limitation, all text, data, graphs and charts (collectively, the “Information”) is the property of Addepar and is provided for informational purposes only. The Information may not be modified, reverse-engineered, reproduced or re-disseminated in whole or in part without prior written permission from Addepar. All rights in the Information are reserved by Addepar.

The Information may not be used to create derivative works or to verify or correct other data or information without prior written permission from Addepar. For example (but without limitation), the Information may not be used to create indexes, databases, risk models, analytics, software or in connection with the issuing, offering, sponsoring, managing or marketing of any securities, portfolios, financial products or other investment vehicles utilizing or based on, linked to, tracking or otherwise derived from the Information or any other Addepar data, information, products or services.

The user of the Information assumes the entire risk of any use it may make or permit to be made of the Information.

Addepar makes no express or implied warranties or representations with respect to the information (or the results to be obtained, but rather the use thereof), and to the maximum extent permitted by applicable law, Addepar expressly disclaims all implied warranties (including, without limitation, any implied warranties of originality, accuracy, timeliness, non-infringement, completeness, merchantability and fitness for a particular purpose) with respect to any of the information.

Information containing any historical information, data or analysis should not be taken as an indication or guarantee of any future performance, analysis, forecast or prediction. Past performance does not guarantee future results.

The Information should not be relied on and is not a substitute for the skill, judgment and experience of the user, its management, employees, advisors and/or clients when making investment and other business decisions. All Information is impersonal and not tailored to the needs of any person, entity or group of persons.

None of the Information constitutes an offer to sell (or a solicitation of an offer to buy), any security, financial product or other investment vehicle or any trading strategy.

Addepar does not recommend, endorse, approve or otherwise express any opinion regarding any issuer, securities, financial products or instruments or trading strategies and Addepar’s research products or services are not intended to constitute investment advice or a recommendation to make (or refrain from making) any kind of investment decision and may not be relied on as such.

Addepar, Investment Sentiment Index and other Addepar brands and product names are the trademarks, service marks or registered marks of Addepar or its subsidiaries in the United States and other jurisdictions.

© 2023 Addepar. All rights reserved.