DECISION BRIEF

Method for Simulating Returns in the Investment Forecast Model

Scope	This *brief* describes the methods that Rain Dog's Investment Forecast Model (IFM) uses to generate simulated asset-class returns and inflation, the rules that govern which methods apply, and the calibration tolerances the simulated returns are designed to meet.
Author	Steve McConnell, CFP®, ChFC®
Status	Final
Date	2026-07-17

1. Summary

Rain Dog's IFM forecasts household outcomes by running hundreds or thousands of trials, each of which needs a full series of annual asset-class returns across every asset class the household holds. This *brief* describes the methods that generate those returns.

Decision. IFM generates simulated returns from a hybrid of three sources: in runs of 250 trials or more, 70% of trials draw from a percentile-based engine that imposes a weighted historical correlation structure through Cholesky decomposition; 20% draw from actual historical years with bounded jitter; and 10% draw from a proprietary generator that can produce return combinations outside the historical record. Runs below 250 trials use historical returns only and exist for development speed; client recommendations are finalized at 1,000 trials or more. Inflation is simulated from percentile data on the same basis. Simulated returns are calibrated to the goal of CAGR and standard deviation landing within 2.5% of each asset class's historical targets.

This brief covers return generation and its calibration. It does not cover how the returns are used afterward.

2. Scope & Application

The return simulator is the bottom layer of Rain Dog's IFM. Every result the model produces is based on the return series used by the trials, so the quality of the forecast is bounded by the realism of the return series. The simulator's job is to produce annual returns by asset class, across trials, that behave like the returns real portfolios have experienced and could plausibly experience.

This *brief* governs the generation of simulated annual asset-class returns and inflation; the weighted correlation matrix; the percentile return distributions; the positive semi-definite matrix preparation and Cholesky application; historical returns with jitter; the proprietary custom returns generator; the trial-count rules; and the calibration of simulated CAGR and standard deviation to historical targets.

3. Design Criteria & Methodology

3.1 Design Criteria

The return generator is built against four goals:

- **Calibrate to historical targets—**Simulated CAGR and standard deviation are calibrated to land inside the tolerances in §6 against each asset class's historical values.
- **Preserve non-Gaussian return shape—**Each asset class keeps its own empirical distribution, including skew and fat tails, rather than being reduced to a mean and a standard deviation.
- **Preserve cross-asset correlation—**Returns are drawn as a correlated set, not independently by asset class, so that both positive and negative co-movement survive into the simulated years.
- **Allow outcomes outside history—**Subject to the constraints above, allow for combinations of asset-class returns within a single year that have not been observed. The historical record is not treated as the boundary of what's possible.

These are goals for the generator as a whole, and every asset class currently in use meets all four design criteria, subject to the exceptions noted in §3.2.

3.2 Methodology

Weighted correlation matrix. Portfolio optimization and Monte Carlo simulation are both highly sensitive to the period used to compute the correlation matrix. Rather than pick one period, Rain Dog computes three separate correlation matrices—1972–2025, 1990–2025, and 2016–2025—and takes a weighted average of them at 5x, 3x, and 2x respectively. This design puts the most weight on recent returns structure while also including older regimes to some degree.

Percentile return distributions. Each asset class carries a percentile return dataset. The simulator draws a percentile and maps it back to a return for that asset class. Skew, fat tails, and asset-specific shape are carried in the data itself; these all carry into the random returns generated, and nothing has to be assumed about the form of the distribution.

Correlation preservation. Statistically independent percentile draws across asset classes would reproduce each asset class's distribution but destroy the relationships between them. IFM prepares a positive semi-definite correlation matrix, applies Cholesky decomposition to it, and uses the result to impose the target correlation structure on the draws. For this reason, returns must be generated as a set per trial-year rather than asset class by asset class.

Historical returns with jitter. A portion of trials uses returns from randomly drawn actual historical years, varied slightly, under the formula: historical return - 0.25 × standard deviation + random × 0.5 × standard deviation. The variation is bounded to +/- 0.25 standard deviations on purpose. It prevents the model from replaying exact history year for year while keeping the real co-movement and regime behavior of those years intact.

Custom returns generator. A proprietary generator supplies the remaining share of trials and produces combinations of returns that vary further from the historical record of combinations of returns. This is the model's black-swan capacity: not-yet-observed but plausible.

Blend and trial counts. At 250 trials or more, the blend is 70% correlated percentile draws, 20% randomly drawn, jittered historical years, and 10% custom generator. Each trial/year combination draws its returns entirely from one source, because mixing sources inside a trial-year would break the correlation structure. Below 250 trials IFM uses historical returns only. That breakpoint is a development convenience for fast exploratory runs, not a recommendation standard.

Inflation. Inflation is simulated from percentile data as well, for the same reason returns are: it is not normally distributed, and a single fixed assumption would hide the variation that matters to a long-horizon plan. Inflation is treated as uncorrelated with asset class returns.

Calibration. Simulated returns are tuned until they reproduce historical asset-class characteristics. Simulated arithmetic average returns are calibrated to land within 2.5% of the target return—at a 10% target return, that is 10% ±0.25%. Simulated CAGR is calibrated to land within 2.5% of the target return—at a 10.0% target return, that is 10.0% ±0.25%. Simulated standard deviation is also calibrated to land within 2.5% of target. While all asset classes meet the tolerance for standard deviations, some asset classes do not fall within tolerances for CAGR. With current calibrations (2026), Rain Dog accepts 5 asset classes that fall outside the CAGR tolerance. The exceptions are listed in §6. In all cases, asset classes that fall outside the tolerances for CAGR fall within 1.0% of target for their arithmetic-average returns.

Calibration runs and recommendation runs. Calibration runs establish that the generator reproduces target characteristics and use far more trials than any client-facing run—nominally 50,000. Client recommendations use 1,000 trials or more, which is enough for the recommendation once the generator itself is calibrated.

4. Background & Rationale

4.1 Why this approach

The conventional Monte Carlo approach takes a mean and a standard deviation per asset class and samples a normal or lognormal distribution constructed from the mean and standard deviation. This simple approach fails to accurately model returns in fundamental ways. Real asset-class returns are skewed, have fat tails, and have shapes specific to the asset class.

The second failure of the conventional Monte Carlo approach is failure to handle correlation. A forecast that draws each asset class independently produces years that cannot happen, and tends to produce more diverse returns than are historically justified. To evaluate the MVO effectiveness of different portfolios, correlation is a core component and should be modeled accurately.

A third failure mode is over-indexing correlation data on a specific period. Correlations have varied across historical periods, which is why the IFM's correlation matrix is weighted across three distinct periods rather than computed over a single period.

The fourth failure mode is related to the use of overly literal historical data. A model built to simulate only the historical combinations of returns implies that nothing can happen that has not already happened. History shows that black swan events do occur, which suggests that they should be modeled, to the degree possible. The custom generator exists so that the model does not treat the observed record as the outer bound, particularly for combinations of asset-class returns in a single year.

No single method of generating returns addresses all the problems of the conventional approach. The percentile-plus-Cholesky engine represents the distributional shape and the correlation structure and does the bulk of the work at 70%. Jittered historical years carry high fidelity historical regime behavior at 20%. The custom generator supplies rare-event capacity at 10%, enough to matter in the tail without letting invented scenarios drive the central estimate.

Calibration is used to ensure that the combination of methods produce returns and standard deviations in line with what has historically been observed.

4.2 Alternatives considered

- **Normal or lognormal mean and standard-deviation Monte Carlo—**Rejected. This discards empirical return shape, skew, and fat tails.
- **Independent percentile draws—**Rejected. It preserves each asset class's distribution but destroys cross-asset correlation.
- **Historical replay only—**Rejected as the primary engine. It preserves regime behavior but is bound to its sample and cannot produce anything outside it. Retained at 20% with jitter, and used alone below 250 trials for exploratory work.
- **Cholesky-only simulation—**Rejected as a complete method. It is strong on correlation preservation but does not by itself carry historical regime behavior or rare-event capacity.
- **Custom generator only—**Rejected. It is built for rare-event capacity and does not preserve correlation structure as well as Cholesky.
- **Full-period correlation matrix only—**Rejected. It overweights old regimes and understates the correlation structure of the current market. Weighting is used to allow full period correlations to influence the correlation model but not dominate it.
- **Recent-period correlation matrix only—**Rejected. It overfits recent behavior. Weighting is used to give recent period the heaviest weight inside the blended matrix, but not sole influence.

5. References

5.1 Related Documents

- Rain Dog Financial, Investment Forecast Model documentation.
- Rain Dog Financial, mean-variance optimization *Decision Brief*.
- Rain Dog Financial, asset tax-location algorithm *Decision Brief*.
- Rain Dog Financial, dedicated portfolio strategy *Decision Brief*.
- Rain Dog Financial, IFM source code, proprietary returns generator.

5.2 Data Sources

- **YCharts—**historical asset-class and fund-level return data.
- **Rain Dog IFM model and source code—**Percentile return data by asset class, inflation percentile data, the weighted correlation matrix, and returns generators.

6. Key Parameters & Calibration Numbers

Correlation period weights, applied as a weighted average of three separately computed matrices:

Period	Weight
1972–2025	5x
1990–2025	3x
2016–2025	2x

Simulation blend at 250 trials or more, as a share of trials in the run:

Return source	Share of trials
Cholesky / percentile correlated returns	70%
Historical returns with jitter	20%
Proprietary custom return generator	10%

Trial-count rules:

Trial count	Method and use
Under 250	Historical returns only. Exploratory and development work.
250 or more	Hybrid blend above.
1,000 or more	Hybrid blend. Required for client recommendations.
~50,000	Calibration runs.

Jitter formula for the historical-returns share:

Historical return - 0.25 × standard deviation + random × 0.5 × standard deviation

Calibration tolerances:

Measure	Tolerance	Note
Simulated average return	Within 2.5% of target	Relative, not absolute: a 10.0% target admits 10.0% ± 0.25%.
Simulated CAGR	Within 2.5% of target	Relative, not absolute: a 10.0% target admits 10.0% ± 0.25%.
Simulated standard deviation	Within 2.5% of target	Relative, not absolute: a 10.0% target admits 10.0% ± 0.25%.

Asset classes outside the CAGR tolerance for 50,000 trial calibration run:

Asset class	Target Avg Rtn	Sim'd Avg Rtn	Target CAGR	Sim'd CAGR	Target Std Dev	Sim'd Std Dev
International Stocks–Japan	5.1%	5.0%	3.0%	**3.1%***	20.4%	20.3%
International Stocks–Latin America	11.4%	11.3%	6.1%	**6.3%***	33.7%	33.7%
International Stocks–ex-US Small Cap	10.5%	10.5%	7.9%	**8.1%***	22.8%	22.6%
Sector–Utilities	6.0%	6.0%	4.2%	**4.3%***	18.5%	18.5%
REIT–International	5.8%	5.8%	3.0%	**3.2%***	22.5%	22.4%

* ***outside tolerance***

7. Notes & Disclosures

Rain Dog Financial® is a trade name of Rain Dog LLC.

This brief is prepared first for Rain Dog's internal documentation purposes and second for compliance and record keeping purposes. It is also shared with current and prospective clients for sake of offering transparency into Rain Dog's methods and approach. It is informational and does not constitute individualized investment advice. Investing involves risk, including possible loss of principal. Simulated returns are model output, not forecasts of any particular outcome, and depend on the historical data, correlation weights, and calibration tolerances used as inputs. Past return, correlation, and inflation experience is not a guarantee of future results.

Revision history

Version	Date	Change	By
1.0	2026-07-17	Initial draft.	SCM