The Development of a Systematic Forecasting Model
How Evolution Became Revolution

The ability to forecast returns is essential to any strategy, and it is the inspiration for our firm's name. "Racon" is a maritime radar transponder that identifies obstacles, hazards and landmarks for ships at sea allowing navigation through darkness or severe weather without the benefit of actual sight. When we set out to design the Revolution model, the mission was to build a robust forecasting engine that could navigate these inevitable financial market storms. In doing so, we targeted several key objectives:

1. Replace discretionary speculation with systematic methods driven by data and rigorous statistics.
2. Standardize measures of risk and premia to capture opportunities across markets in a way that is unbiased and agnostic to asset class labels.
3. Develop a flexible and dynamic engine that avoids the traps of static correlation assumptions because markets change in unpredictable ways.

These objectives represent the byproduct of decades of experience and understanding about the evolution of market efficiencies and data.

Why Systematic?

The late 1980s and early 1990s were fertile ground for discretionary global macro trading. The strategy offered flexibility to trade across markets and geographies unlimited by the silos employed by traditional equity or fixed income managers. Provided one invested the time and effort and possessed the intellectual horsepower, opportunities were plentiful. In 1987, Paul Tudor Jones tripled his money being short on Black Monday. And In 1992, George Soros & Stanley Druckenmiller became legends by breaking the Bank of England. These trades put a spotlight on the strategy, and soon after, pension plans, endowments and sovereign wealth funds were quick to adjust asset allocations in the quest for this uncorrelated alpha.

Following the 2008 credit crisis however, market dynamics changed in several ways, making it increasingly difficult to monetize macro-oriented sources of risk premia. Forces including overcrowded trades, increased central bank manipulation and shorter-term performance demands made timing and sizing trades more difficult, introduced more speculation and created real challenges for traditional approaches to macro investing.

Despite headwinds, significant opportunities still exist for macro managers who can break through the noise that exists in today's markets. However, capitalizing on them increasingly requires the ability to use modern data science with plentiful data and advanced statistical methods to generate forecasts in a systematic way.

Why Cross-Asset?

A core philosophy of our firm is that asset classes react to economic information at different speeds, and asset correlations change frequently to influence market behavior (premia and risk). Thus, the application of market data and economic information should not be confined by traditional asset-class silos. In fact, global macro strategies have the unique ability to exploit them for superior risk-adjusted returns.

Over the years, we have used a series of stand-alone models to generate buy/sell signals within FX, equities, rates, credit and commodities. While each model operated independently, frequently signals from one model informed moves or premia in another asset class. On the other hand, these disparate models could, from time to time, result in conflicting information regarding risk environments which lead to frustration and unprofitable activity. In developing the Revolution forecasting engine, we aimed to mitigate the impact of conflicting signals by combining a large number of signals across asset classes to standardize premia (ex. compare the return, vol, dispersion etc. of a single name equity to commodity and vice-versa) and identify (forecast) hidden premia wherever it may exist.

Why Dynamic Correlations Assumptions?

The chart below highlights the correlation of a commodity (Aluminum) across four assets: Gilt bonds (GBP 10y), Canadian bonds (CAD 10y), domestic growth equities (US Growth), and emerging market equities (EM). In the 3-month window these assets have very low correlation, yet in the 1-month window they are all highly correlated. Clearly, some extraneous factor is driving this change. We designed our forecasting tool to be flexible and dynamic recognizing that correlation structures and market regimes change unexpectedly.

		3M	**1M**	**Change**
GBP 10y	Aluminum	5	81	76
CAD 10y	Aluminum	3	70	67
Aluminum	WF US Growth	18	85	67
Aluminum	MSCI EM	24	88	63

Data Source: Goldman Sachs

Key Forecasting Decisions

Having outlined a clear set of objectives supported in practice and logic, the challenge of designing, testing and implementing a technically robust forecasting engine took years of time, energy and resources. The process also revealed some key foundational design considerations, mainly:

- Structure
- Signal
- Scope

Below we discuss each, noting how Revolution compares and contrasts with other more common forecasting methods.

Structure: Causal vs Relational

The first consideration is whether a forecast is built using a model that (1) specifies the cause and determination of forecasted returns (causal) or (2) relies exclusively on statistical relationships in the data without regard to why they exist (relational).

	Causal	**Relational**
Higher reliance on...	Assumption	Data
Comparative advantage in...	Understanding market relationships	Predicting market outcomes
Widely used in...	Macro, events	Statistical arbitrage, High-frequency

Causal

A causal model imposes structure derived from theory, so that the full power of the data is used in adding precision and quantification to the assumed relationship. The output is a predicted outcome, but there is also a model of the mechanism by which the signals form predictions. The results hinge on whether the assumed causal relationship is valid.

Causal forecasts for returns are numerous and wide-ranging. They include many of the classic macro investment strategies, given that one is often dealing with limited data and trading on events with short or non-representative time series. They would also include many forecasting approaches which rely on economic theories about how interest rate movements translate to investment, discount rates, or credit. These channels are assumed based on economic theory, and the data is used to predict returns according to these links.

For example, a popular model for currency trades relies on the idea that the correlation between the USD and crude oil is strongly negative. This idea is founded on two mechanisms: the fact that crude oil is priced in USD and the fact that the U.S. had been importing large amounts of oil for decades. However, over the past decade, the U.S. became an exporter of oil. The strong negative correlation between crude oil and the USD shifted, which invalidated some currency pair models. While the economic assumptions behind this idea were sound for many decades, taking the empirical relationship on its face is a mistake. The danger of counting on stable economic relationships can also be seen in the famous "Fed model" relating earnings-price ratios of equities to 10-year treasury bonds. Among the various reasons this model does not have consistent predictive power is the fact that the correlation between equities and bonds moves substantially over time due to changing risk premia, flights to quality, etc. Relying on the Fed Model performs well only when the underlying correlations remain valid; however, these relationships can be very unstable.

Relational

Contrast this with a relational model, where the causation of price is left unspecified. Rather than using data to fine-tune assumed relationships, the data must be used to uncover the relationship and predict the result. Accordingly, these relational models are much more flexible with regard to how they connect signals and predictions. There is not an explicit theory on the formation of returns, but rather a statistical likelihood of what will occur in the future. This flexibility comes at a cost—the need for more data, and the increased difficulty in addressing questions of "why" the signal translates to a particular prediction.

Relational forecasts are most commonly used in statistical arbitrage, relative value, and technical analysis. These approaches all involve rich data environments and little economic theory. The forecasts are generated by relationships constructed from the data rather than founded on economic theory.

Consider the difference in trading the effects of events versus the probability of the event itself. For example, an inverted yield curve has been used as a signal to forecast U.S. recessions. This signal has performed well in the post-WWII period in the U.S, and there are economic theories as to why the relationship exists, but they are reliant on strong assumptions. However, having experienced only a handful of observations, predicting a recession remains an entirely empirical matter. Thus, if one wants to identify U.S. recessions, it is hard to avoid these causal, structural models, whether based on the inverted yield curve or related methods. Conversely, if one is interested in trading the broader effects of recession without the need to identify a cause, relational models are helpful. There is a rich set of data through time and across assets to that can provide statistically significant forecasting power for equity discount rates, changing yields, liquidity premia, carry and more. In the end, the relational approach focuses on the effects directly rather than on identifying a recession to forecast a derivative outcome.

Our model, Revolution, is unusual in that it is a macro strategy, yet it is relational, not causal. It processes medium-term signals into medium-term forecasts, but not per a tightly integrated model of the macroeconomy. Instead, it constructs forecasts by using massive amounts of data to flexibly identify and utilize relevant signals. Forecasts are consistent across asset classes, not a bundle of asset-class-specific forecasts. Latent factor models are a key ingredient to achieving consistency and harnessing the data necessary to address macro forecasts with a relational model. The model assumes factor premia influence data widely, but through unobserved, unseen, and evolving channels. This allows the model to consider consistent, cross-asset linkages, and to uncover them with many readings related to them—our signals. Importantly, this framework allows for highly correlated signals, which is the only way to get the rich set of data to give the relational model enough statistical power. Ultimately, Revolution's forecasts make greater use of the cross-sectional data, rather than trying to predict why some premia, factor, or market goes up or down.

We believe that the challenges inherent in using a relational model for a macro strategy are worth the effort. As discussed above, there have been substantial changes to markets and their interaction with macro elements—particularly since the 2008 crisis and perhaps going forward in response to COVID-19. On the other hand, advances in statistical and computational methods continue to improve, allowing for approaches that make use of much larger amounts of cross-sectional data and can consider more flexible relationships in the data. This massive increase in data provides enough statistical power to compensate for the lack of specified channels, and the flexibility inherent in the relational approach allows Revolution's forecasts to identify cross-asset links that continue to evolve with changing markets.

Selection: Reducing Bias and Variance

Causal versus relational models are not just a matter of flexibility; they also raise issues of forecast bias and variance. Any forecast seeks to be unbiased, so as to be correct on average, even though the forecast will be too high or too low at times. The forecast should also minimize variance, such that it is not just right on average, but also avoids large misses. There is tension in these two goals: making a relational model unbiased requires a model flexible enough to infer the expected market values. This flexibility intended to make the model unbiased risks that the model is too variable in its conclusions. Accordingly, reducing forecast bias and forecast variance is an important consideration in making relational models effective. Below we discuss ways in which this can be done.

Lower Bias	**Variance**
Relational signal processing	Ensemble forecasts
Large set of signals	Latent state model

Large Signal Sets Lower Bias

To get an unbiased forecast, it is important to model the signal processing correctly. Causal models do this by using theory-based specifications, while relational models do this through data-based construction. For example, in a causal fundamental approach to forecasting equities, the forecast for the security depends on that same security's fundamental data, and possibly some other macro data. These models rely on various mechanisms by which a security's attractive fundamental data is overlooked or undervalued. Contrast that with relational approaches including the so-called "quantamental" approach of using the fundamental data from many other related equities in order to better forecast the equity of interest. This allows the relational model to better ascertain the value of the target by considering it as one noisy expression of a broad market result. The drawback to using large, relational approaches is that the data is noisy, so it can be difficult to ascertain which signals matter. The large number of potential signals is both an opportunity and a challenge—particularly given high correlations. Causal models often avoid this complication by using a small number of signals. If the model is overly simplistic, it may be precise yet biased due to abandoning relationships that matter for the outcome but were eliminated for the sake of statistical clarity. Such a model may not be flexible enough to deal with the changing dynamics and cross asset relationships that exist in today's markets. For example, classic fundamental models based on dividend yields ran into trouble as markets changed, large-cap firms retained more cash, and discount rates decreased.

Machine Learning Techniques Mitigate Variance

Our forecasting model, Revolution, is built with these competing issues in mind. In order to use the relational approach, we needed to make the model flexible enough to estimate the relevant relationships and signals, while still providing a precision level that makes the forecasts powerful. To achieve these dual goals, it produces a wide set of signals and thus a very high dimensionality, while maintaining precision by incorporating machine-learning tools, including latent state models and ensemble forecasts. Latent state models allow the forecasting engine to process thousands of signals to assess marginal information without trying to model each signal separately. Ensemble methods allow the model to run many thousands of flexible forecasts (thus maintaining the unbiased results) and average across them to reduce the variability in the output.

Together, these techniques help Revolution implement an extremely flexible model to detect the relationships between signals and predictions, in line with our desire above to use a relational approach. While this is more commonly done with denser, higher frequency domains, machine learning tools give the model ways to implement this for our macro, medium-term domain.

Scope in Application

The last of the three key design elements of the Revolution forecasting model is its wide scope.

Wide Scope Improves Signal to Noise

The scope of a forecast refers to how widely the forecast is applied. A model may predict a single event, security, sector, asset class, or factor. Models with a narrow domain result in more specificity

	Narrow Scope / Application	**Wide Scope / Application**
Advantage	Make use of idiosyncratic information	Scale consistently and systematically
Frequency	More variability to the forecast	Narrower range of forecast values
Aggregation	Combine many micro insights	Broadly utilize core macro insights

of prediction. This would include models for a single market, trade flow, or security. These narrow forecasts often estimate a wide range of predictions, allowing the system to trade substantially in that one area. Examples would include CTA models for a specific futures contract, long-short equity models built for a single sector, or models built for a particular trade-flow, such as on-the-run versus off-the-run treasuries.

On the other hand, models with a wide scope in application are forecasting a broad phenomenon, not the specifics of one security. Examples include relatively general models of carry, momentum, and value, as applied to entire asset classes or beyond. These forecasts have increased stability of the forecasts due to using broader data. However, it is typical to see wide forecasting scopes generate narrower ranges of predicted values. This is because the system is forecasting a more general component, one which applies more broadly across markets, and such general components are less likely to have fast and large ranges of predicted values. Still, the wider scope gains substantially by improving the signal-to-noise ratio through diversification of the forecasts.

This breadth of application is a cornerstone of how Revolution operates. It forecasts the latent factors mentioned above which express and impact the markets widely—across more than one style, security, or asset class. While broad market factors move more slowly and in narrower bands, Revolution is able to use those forecasts over an extremely diverse application space. That scale and scope of predictions is valuable downstream in the portfolio construction. Though individual signals are noisy, combining low correlation sources boosts the total signal-to-noise ratio.

Ultimately, we believe forecasting using a relational model and advanced statistical methods applied through a broad set of signals is both unique in its approach, and perhaps more importantly, successful in its objectives: (1) replacing speculation with relationships validated by data; (2) standardizing risk premia agnostic to asset-class labels; and (3) avoiding static correlation traps. The design evolved through years of trading experience, extensive research and a curiosity that recognizes the fragility of traditional causal models and other classic linear approaches to forecasting. And while no model is perfect, we believe acknowledging the existence of hidden states in the market and giving data the statistical flexibility to stand on its own presents a distinctive approach to modern macro investing. Demonstrating this through consistent performance remains the benchmark. Our hope is that this paper provides a narrative framework for investors to understand why, as financial markets continue to shift and traditional forms of alpha generation are challenged, managers who can evolve and apply modern data science to investing (e.g. forecasting) are keenly positioned to provide uncorrelated returns over time.

Disclosures:

permitted by law, by permission of the Investment Manager's senior management following the completion of an investor suitability analysis. Any investment in the funds managed by the Investment Manager involves a high degree of risk and may not be suitable for all investors. Before making an investment decision, potential investors are advised to carefully read the applicable CPOM, the limited partnership agreement (if applicable), and the related subscription documents, and to consult with their tax, legal and financial advisors. While all of the information prepared in this document is believed to be accurate, the Investment Manager makes no express warranty as to its completeness or accuracy, nor can it accept responsibility for errors appearing in the document. The manager relied on certain third-party sources of information when completing this report. The details, parameters and descriptions found in this document may be modified in the future by the Investment Manager in any manner. The information in this report is provided for informational purposes only and should be not be relied upon or construed as creating formal investment rules or requirements. In the case of any inconsistency between the descriptions or terms in this document and the CPOM, the CPOM shall control. Opinions expressed by the author(s) of this article do not necessarily reflect those of the Investment Manager or its affiliates. Examples in this report are not intended to indicate performance that may be expected to be achieved. Past performance in not indicative or a guarantee of future results. No assurance can be made that profits will be achieved or that substantial or complete losses will not be incurred. The Investment Manager, its affiliates and employees do no provide tax, accounting or legal advice. Please consult your tax, legal or accounting advisor regarding such matters.

Racon Capital Partners LLC
234 W. Florida St. — Suite 700 — Milwaukee, WI 53204
Phone 414.877.1270 — www.raconcapital.com