Commodities Risk Easy Button:
Configure Advanced Risk Analytics from Scratch in 20 Minutes or Less

Albert Hopping, Eric Carriere, Michael Busby, and Christine Jackson, SAS Institute Inc.

ABSTRACT

There is often a misconception that performing risk analysis must be a major endeavor that requires your institution to have advanced quants and technical skills in house. However, with SAS® Commodity Risk Analytics, the needed expertise is built in, but the risk platform is still fully customizable. The easy-to-configure interface runs in your browser and can have you understanding your risk in minutes. Explore your VaR and other risk metrics, and drill down to the position level on the fly. Enable traders to understand the impact of their trades before they make them.

INTRODUCTION

As astute risk professionals know, the capability to produce accurate and timely risk metrics is a very valuable asset to any organization. Foresight into potential landmines in a commodity portfolio allows companies to proactively mitigate these vulnerabilities, rather than scramble to understand and react after the fact when market moves negatively impact the portfolio.

Commodity trading has many challenges due to the physical nature of the business. Obtaining and communicating consolidated risk results across portfolios and lines of business can be daunting for large energy firms and trading desks. Investment in a consolidated commodity risk management system is predominantly driven by the desire to strengthen an organization's strategic decision-making, improve efficiency, and reduce probability of loss. A central commodity risk system that consolidates prices, positions, and counterparty information helps boost profitability by enabling risk managers to analyze and report market and credit risk exposures in an on-demand environment, as well as support commercial intra-day arbitrage opportunities. Such an on-demand risk calculation engine can support accurate decision management and aggregated risk reporting, effectively linking the front and middle offices.

Companies vary in their level of sophistication in risk management and their technological capabilities. Common situations include the following:

- Relatively small but growing companies that have been managing the business through an increasing number of spreadsheets. The business has grown to a point that the processes are no longer reliable, multiple versions of data and reports are not controlled, and analysts spend more time troubleshooting technical errors than performing work that supports the organization's business goals.

- Companies using a logistics system to manage operations and transaction scheduling that require advanced risk metrics their original system was not designed to support. The firm seeks a solution that can both fill substantial gaps in its risk reporting and integrate with existing systems without causing disruption to the business.

- Organizations that have multiple commodity systems across different products and business units. While each system works adequately within its silo, consolidating information to achieve an aggregated view of the portfolio is challenging. The data in

each system might have different structures or naming conventions, requiring a consolidated risk repository that can generate one version of the truth.

- Multiple groups within a company that have varying levels of business, quantitative, and technical knowledge. Some users prefer out-of-box features that can be deployed quickly and run with limited support. Others want to be more hands on with risk modeling, modifying risk models, or creating their own. Either way, the capability to report and distribute results to executive management effectively is as important as the calculations themselves.

Whether your company has an active trading division, operates a collection of physical assets, runs a hedging program to lock in input costs/revenue, or simply has exposure to commodity prices through its normal course of business, SAS® Commodity Risk Analytics is a comprehensive and flexible risk solution to address the full spectrum of risk needs. The **solution's browser**-based interface and its comprehensive data model are simple enough to configure your commodities, enter deals, and create risk metrics in just minutes. Yet, you can still customize it to fit more exacting needs. The challenges inherit in commodities trading and risk management make this no easy task. Luckily, as this paper will demonstrate, the flexibility and intuitiveness of SAS® Commodity Risk Analytics is like hitting the proverbial **"easy**-button.**"**

COMMODITY TRADING CHALLENGES

Trading commodities is not for the faint of heart. The physical nature of the commodities business creates unique challenges when valuing them as assets. Unlike equities, commodity trades include a physical underlying component that requires storage and transmission, such as oil, gas, soybeans, or coffee. This physical component means forward values must be inflated by the costs of storing the commodities. Because commodities are produced and consumed, the market has unique supply and demand dynamics that must be forecast into the future. Time becomes yet another element of the quantification mix because commodities can be perishable and therefore temporal. They also have different physical characteristics and cannot be substituted – a barrel of oil from Texas is not the same as a barrel of oil from Alberta, a gigajoule of natural gas delivered in the winter is not the same as a gigajoule of natural gas delivered in the summer.

These complexities result in extremely volatile assets that must be grouped together by type, location, and time. For example, crude oil is not a unique asset unless referred to by type (light oil), location (Midland, Texas) and delivery period or time (Dec-2020). When hedging or valuing this asset, a barrel of the same type, near the same location, delivered at the same time must be used.

Consensus around what that barrel is worth becomes difficult as proxies for this unique barrel must be found. Fortunately, actively traded futures markets allow consensus to be formed and price discovery to take place. But determining value when reasonable proxies do not exist means the amount of data becomes large and the math gets complicated.

BIG DATA AND COMMUNICATING STATISTICS TO A WIDE AUDIENCE

Large data sets and complicated mathematics are endemic to many complex problems. Communicating these problems to a non-technical audience is difficult. When predictions of the future are involved, as they are for risk, the topic of statistics cannot be avoided. Normal, lognormal, kurtosis, fat-tails – these are just some of terms that enter the lexicon to describe how portfolios behave. The true distribution of future events is unknowable, particularly in the opaque world of commodities. What is a barrel of oil worth in Midland, Texas? How have returns behaved historically and how have they behaved in relation to

other assets in your portfolio? And, most importantly, how will returns behave going forward?

Statistics is the best scientific tool to describe future outcomes, and it is based on the work of giants who have come before. Cardano's game of chance, Pascal and Fermat's correspondence on the problem of points, and Bernoulli's solution to the St. Petersburg paradox[1] create an environment where the questions can be framed in a logical context and predictions of the future can be made. For example, if one selects a random person and knows the characteristics of the population they came from, conclusions can be made about the height of that person if we can accurately estimate the average height of the population and how individuals deviate from this average. Heady stuff. Can we draw the same conclusions for a barrel of oil? It would be nice to think so; however, the devil is in the details.

What constitutes a representative sample of data in describing the potential distribution of outcomes for the return on a barrel of oil? Can we say the return on a barrel of oil behaves in the same way as a stock return on an equity exchange? Research shows equity returns are normally distributed[2], the same distribution used to describe the height of a population of humans. Conveniently, one only needs the mean and standard deviation to describe this distribution, which can be revealed by analyzing the data. Lots and lots of data.

This might seem intuitive, but the world of statistics is full of surprises and fraught with assumptions. Therein lies the difficulty when trying to explain statistical results to the uninitiated. Not only must the results be explained, but why the results are potentially wrong must also be explained. Uncomfortable conversations with executives begin with phrases such as "we have to assume normality" or "markets must be liquidly traded" and, of course, "assets must be independent and identically distributed."

If relatively simple statistics often create a barrier, risk management brings a further set of communication challenges to overcome. This communication is hampered by the fact that risk processes are often opaque: Black-box models are used with questionable assumptions. An inability to validate these methodologies leads to a lack of trust. This issue is compounded by risk metrics being nonadditive, which prevents most systems from supporting drill-down of the results or flexible reporting hierarchies. Further, hedges are often shown to be weaker than many traders or executives expect, for a variety of reasons: misalignments of hedges during delivery, correlation breakdown, basis blowouts, rare events, even force majeure on physical contracts.

Analysts often rely on overly simplistic, potentially inaccurate models because more complicated and representative models are harder to calibrate and cause other downstream complications. Basic statistics provide a logical roadmap to frame a discussion and interpret the results. But it is vital that these weaknesses are communicated, and their shortcomings revealed. The path of least resistance leads to a tendency for black box solutions and opaque systems, where answers are buried within lines of code and reams of data to the detriment of risk management. It is easy to hide the warts and proclaim the models are accurate because the math is hard. However, the logical constructs of mathematics are only as good as the axioms that serve as their foundation.

COMPLICATED MATH AND DEEP UNDERSTANDING NEEDED

The field of quantitative mathematics underlying the creation of risk metrics is not trivial. It is often non-intuitive to those new to the field. This non-intuitive nature leaves many incapable of performing "sniff tests" or "gut checks" on risk results. This is even true for simple methodologies such as Delta-Normal VaR and simple concepts such as non-additivity. Making basic errors is not uncommon among new practitioners when calculating volatilities and correlations, leading to material impacts in risk metrics. Luckily these simple

mistakes can typically be identified with common validations. However, miscalibrations of more complicated price models or systems can be much harder to notice and just as impactful.

When the assumptions and calculations are not tested and validated by a quantitative analytics (quant) team, the natural outcome is error-prone. Unfortunately, many **institutions' risk management teams** do not have quants on staff to perform this validation. The market for quality quant resources that understand the nuances of commodities is tight. Advanced degrees in the field address only the financial services industry with any meaningful depth, leaving an unfilled gap in support for commodity trading activities. It is an unfortunate gap given that commodity volatilities are much higher than most other traded products.

The math is only part of what is needed to add value with risk metrics, perhaps the easiest part. The judgment required to design a risk system is likely rarer. Pricing models must be selected; price curves must be mapped; transition of positions from forward to spot must be handled. Quality adjustments are often needed, missing or bad data must be addressed and newly traded products lack price history. And that identifies only a handful of issues where policy decisions and commitments must be made.

For example, thoughtfully mapping and managing price curves on the input side of a risk project dramatically affects the usefulness of the results, if not the results themselves. It requires disaggregation of prices to a liquid hub with basis curves for location or quality. Without managing curves in this manner, some reporting value might remain unavailable. For future-dated volatility curve creation, judgment is needed to weigh the pros and cons of several approaches. This list is longer than that of the products traded.

INFRASTRUCTURE CHALLENGES

Beyond the abstract math, the infrastructure for a commodities risk management system mounts its own challenge in both the physical and logical sense. One of the largest physical infrastructure issues is the size of the mathematical problem itself. The number of risk factors is likely in the high thousands. Some institutions have millions of positions, especially after risk factors and other component risks are disaggregated into additional positions. When performing calculations for multiple time horizons or scenarios, this volume is literally multiplied. Monte Carlo simulation exacerbates this thousands of times over. Together the desired data and calculations quickly outstrip the capability of a single machine. Coordinating results for a nonadditive measure like VaR across multiple machines is a logistical challenge that further complicates the process and its mathematics.

Even when one machine is able to handle the required memory and computation, other concerns arise. Logistical or other capture systems, with risk largely an afterthought, can be poor sources for risk processes. Beyond the tedium of unit of measure conversions, additional metrics such as quality adjustments, day count issues, and holiday schedules must be managed. Further, historical price data will not align on all dates, will invariably have data quality issues, and might be updated. Each issue might be simple enough on its own; taken together, they create multiple layers of complexity.

As if these issues were not sufficiently daunting, there is also the matter of governance. Source systems and the calculations both need access control. Computing environment access to the input data, the process, and the results needs to be controlled for different groups of users per logical environment. Different reports must be surfaced to various groups. Edit access and execution access need to be more tightly controlled. Yet, ad hoc analysis should be easy for many without risking the integrity of the daily processes and results.

WHAT SAS DOES FOR COMMODITY RISK

To handle the many challenges commodities present, a solution must be able to process data on an industrial scale. It must be able to analyze the data in an efficient manner and present the data in an intuitive way. The solution must address the unique needs of the commodities business, speaking a language the business understands. One must be able to take the massive amounts of data required to describe the characteristics of commodity type, location, and time, and drill-down, aggregate and disaggregate all in real time. Statistical abstractions must be communicated in an intuitive manner, allowing the analyst to easily confirm "gut-check" assumptions. Deep understanding must be achieved through out-of-box results that are open to investigation. Not a black-box that hides the complicated math underlying the analysis. Furthermore, the solution infrastructure must scale and consume valuable computing resources with extreme efficiency, providing for near instantaneous results that are necessary in a world of extremely volatile assets.

SAS® Commodity Risk Analytics addresses all these challenges as well as providing other key benefits such as scalability, flexibility, and transparency. The main challenge with any commodity risk system is how to handle the massive amounts of data required, and then how to process the data to produce meaningful results. Once the processing is complete, the data must be presented in an intuitive way, to be consumed by front-, middle-, and back-office operators, or consumed by other, more advanced analytic processes. To begin to understand how SAS® Commodity Risk Analytics achieves these lofty goals, we must start with the risk engine that drives the solution and feeds the results to be communicated to organizational executives and operators.

BIG DATA AND COMPLICATED MATH - THE RISK ENGINE

SAS® Commodity Risk Analytics is powered by the SAS® High-Performance Risk engine. The engine is capable of consuming position and market data on an industrial scale to generate multiple simulated market/portfolio states, calculating volatilities, and correlations automatically from historical time series data. The risk engine not only handles the math, but it does so in an optimized way for fast processing. The analysis data is kept in memory in a scalable distributed environment that allows for "on the fly" drill down to the trade level to ensure "gut-check" assumptions are verified. You can slice-and-dice analytical results on the fly without having to re-simulate or recalculate your portfolio. This allows the user to verify that hedging strategies are achieving optimal results before markets close and regret sets in.

The SAS® High-Performance Risk engine is transparent, allowing users to open the models and review the assumptions behind them. The engine quickly configures portfolio data and generates risk factors based on logic built into the SAS® Commodity Risk Analytics data model. The risk engine also enables users to define their own assumptions around what constitutes a risk factor as well as the processes governing the behavior of that risk factor. This transparency helps satisfy the deep understanding required to fully appreciate all the risks your portfolio is exposed to, at any level of detail required.

Users can utilize the out-of-box analytical models or configure their own quantitative models and pricing methods. Model results can be compared with the click of a mouse, from day-to-day, week-to-week or month-to-month to ensure models "rolling over" from prompt to spot are all handling risk correctly. Historical simulations can be compared to Monte Carlo results effortlessly to gain insights into statistical assumptions. Having the ability to open the black box and question the assumptions is paramount to helping understand the nature of one's portfolio. Thereby revealing what risks lie in the shadows, waiting to wreak havoc in ways unknowable within the black box paradigm.

Infrastructure issues can be a challenge when dealing with large volumes of data. To address this, the solution is scalable and can use a distributed computing architecture to manage large, complex portfolios. Just as important, the solution can be scaled down for smaller portfolios that just need the basics.

Whether your requirements are complex or simple, global or local, real time or end-of-day, SAS® Commodity Risk Analytics gives users the power to configure their portfolio in an easy, intuitive way using an interface rich in content, built with the commodity business in mind.

RICH CONTENT DESIGNED FOR COMMODITIES

Having the flexibility and power of the risk engine is all well and good, but how do you get those positions and prices configured to harness that power easily? To obtain a quick time to value, a business user must have an intuitive starting point without having to be a quant or programmer.

SAS® Commodity Risk Analytics provides a data model designed specifically for storing the attributes of a commodity portfolio, so that it can be efficiently used for risk reporting and analytics. Pre-configured analysis automatically leverages the risk engine while shielding users from any complex setup.

A graphical user interface provides users an intuitive way to enter data into the heart of the system; a window to view, enter or modify counterparties, transactions, and market curves. The system can manage the full lifecycle of a transaction and be used as a system of record. Alternatively, it can also simply be used as a consolidation layer to aggregate commodity data from multiple sources/systems to create a common structure across commodities and business unit silos. Pre-built web service calls can be leveraged to automate the integration from external sources and existing commodity systems.

COMMUNICATING RESULTS TO A WIDE AUDIENCE

Given an easy way to configure the portfolio and a powerful risk engine, how can results be communicated to various consumers in the organization? Some might get more value from simple summarized results while others are keenly interested in the detail behind the numbers.

Embedded in the solution is SAS® Visual Analytics technology. This data visualization technology gives business users the capability to easily explore relationships between large amounts of data, build reports and dashboards, and surface reports online, through email or via mobile devices. The key detailed risk output is automatically made available so that different graphical and tabular reports can be made as part of production reporting requirements. Ad Hoc reporting can be driven by these production data sets or be produced via custom created output data depending on the needs of the users.

SAS® Commodity Risk Analytics makes it easy no matter how you choose to communicate results to your audience. Whether it be a simple end-of-day process for executive reporting or an ad hoc risk exploration into the hedge assumptions within your portfolio, SAS visualization tools will ensure that the results are intuitive and explainable no matter how wide your audience is.

WORKING EXAMPLE

Let's work through an example to demonstrate how effortlessly this can all be done. Many data entry procedures are automated in operations through back-end integration points. For this example, we will show how the front-end User Interface can be used to quickly set up a portfolio and configure the necessary data to calculate advance risk analytics.

Configuring the Basics - Static and Administrative Data

Before entering portfolio transactions, we must first setup all the attributes that are required to make it easy. Most of the effort is done once, up front with periodic updates when needed. Typical static and administrative data elements include (but are not limited to) defining commodities, units of measures, conversions, counterparties, agreements/contracts, limits, locations, and price indices as shown in Figure 1.
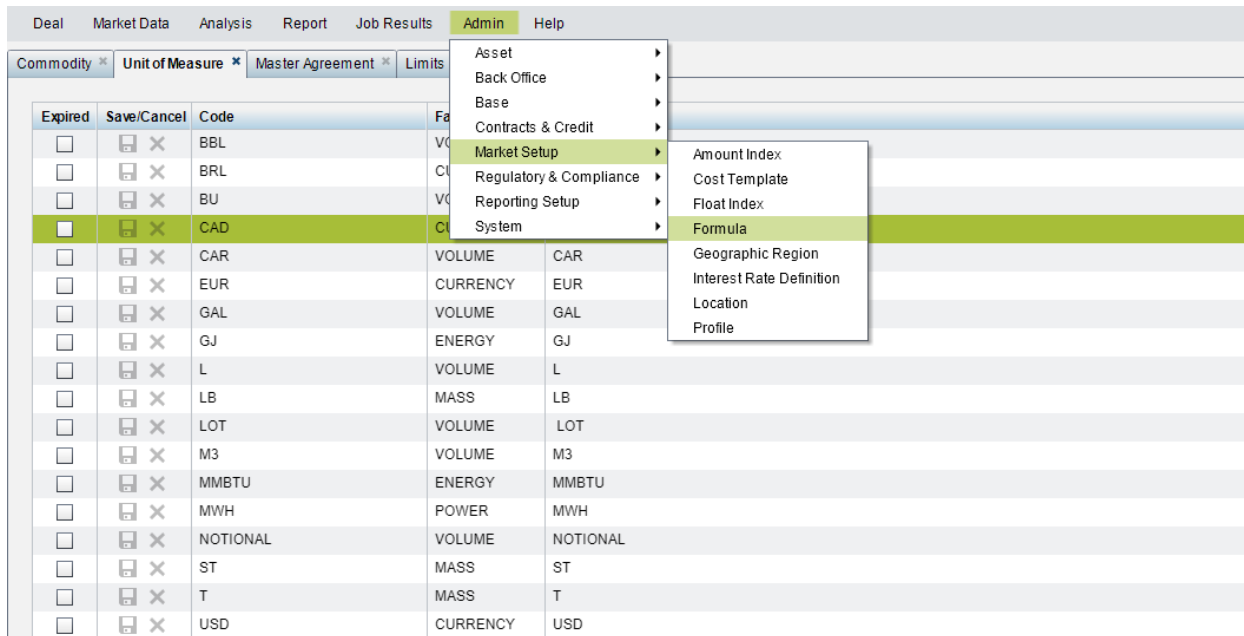

Figure 1 Unit of Measure Example with Market Setup Menu

The setup can be very simple – one or two commodities based on a few, simple price indices – or more complex – multiple commodities and currencies with complex basis and benchmarked price indices. The intuitive layout of the interface facilitates this set up while helping users gain a deep understanding of the data required and how it is used.

Creating the Portfolio

With the administration and static data elements configured, you can now easily enter transaction data in the system. From physical cargos to financial swaps and options, all the required transactions can be represented. Pricing can be simple (fixed price or index for deals) or based on complex formulas entered by the user. Either way the solution will automatically decompose everything into the risk factor components. Each component is broken down by commodity type, location, and delivery period to ensure hedging results match upfront expected results, as shown in Figure 2.

Figure 2 Deal Viewer Displaying Transaction Data Used to Configure Portfolio Attributes

Managing Market Data

The pre-configured valuation methodologies for these transactions require market data. Not only does SAS® Commodity Risk Analytics provide a repository for transactional position data, it also stores and manages current forward market data such as prices, FX rates, interest rates, and volatilities. Figure 3 shows the Price Curve Display.
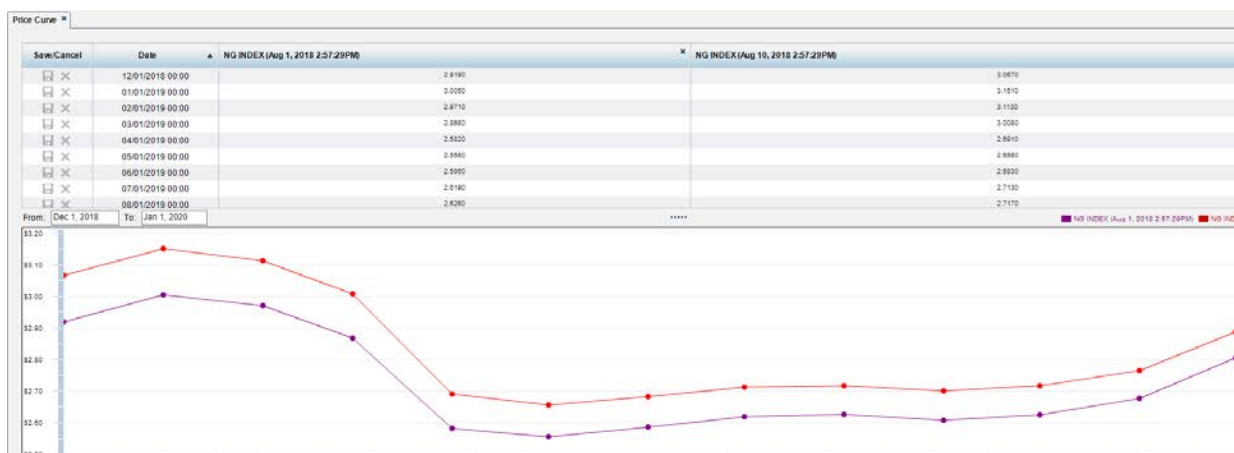

Figure 3 Price Curve Display

To provide risk analytics, historical time series data must also be provided. This enables the solution to calibrate models and configure out-of-box analytics, so your quants don't have to. However, if a quant prefers, parameters can be passed into the solution to better understand the math and the inputs to the model.

Run Out-of-Box or Ad Hoc Real-Time Analysis

Now that transactions and prices are in the data model, the pre-built analysis can leverage the risk engine to generate key risk metrics. It can be limited to running a basic Mark-to-Market or more advanced calculations such as Value-at-Risk (as shown in Figure 4) or Cash Flow-at-Risk. Credit exposures can also be generated to estimate Potential Future Exposure where holding periods can be five years or more, depending on the exposure to the counterparty.
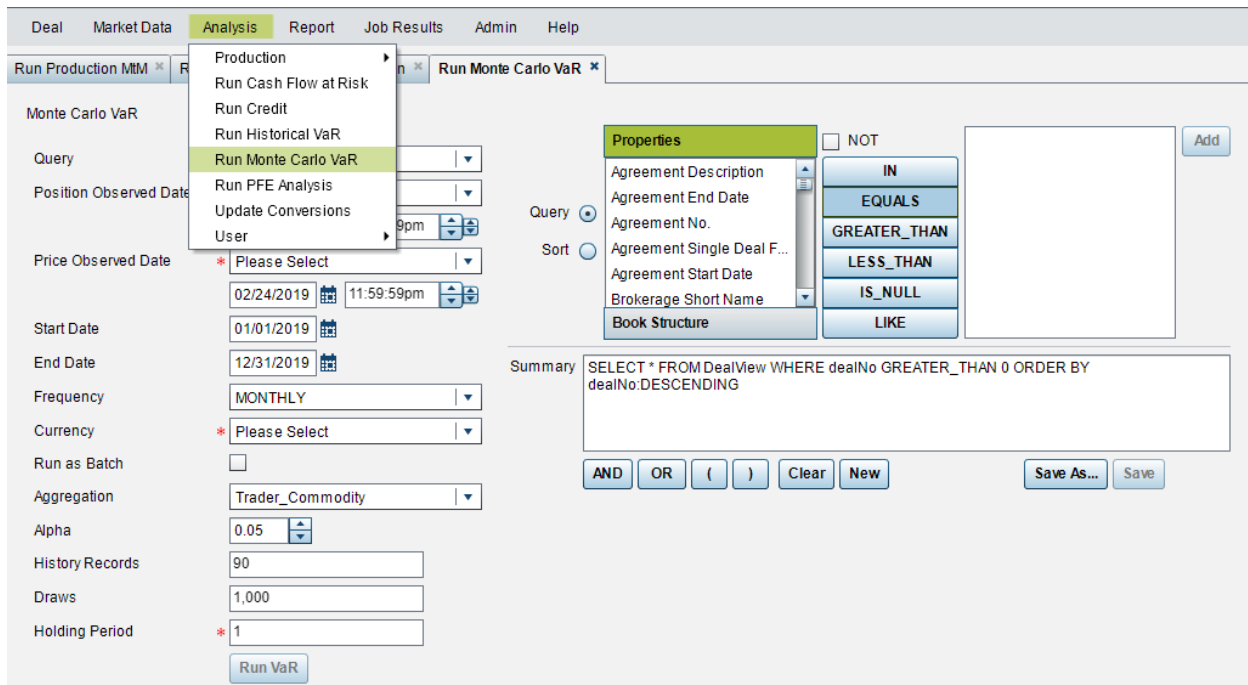
Figure 4 Monte Carlo VaR Analysis Setup

Pre-built analytics are great for quick time-to-value returns. This example demonstrates all the requirements, from administrative data to portfolio transactions to forward and historical price data. After the data is loaded, simply choose which analysis to run. Production reporting can be configured and run daily to satisfy executive requirements. Ad Hoc and real-time analysis can be performed on the fly and in real time.

The output from ad hoc analyses can be viewed in the SAS® High-Performance Risk user interface (shown in Figure 5), which takes a risk cube of portfolio data and processes it in-memory to minimize I/O performance inhibiting. You can re-aggregate the VaR calculations, filter portfolio components, and adjust confidence levels on the fly, without having to re-simulate the portfolio.

Figure 5 SAS® High-Performance Risk user interface

For regular production reporting, the result data sets from out-of-box analysis are surfaced through visualized reports. Both market risk and credit risk metrics can be displayed in the format the consumer finds most intuitive.
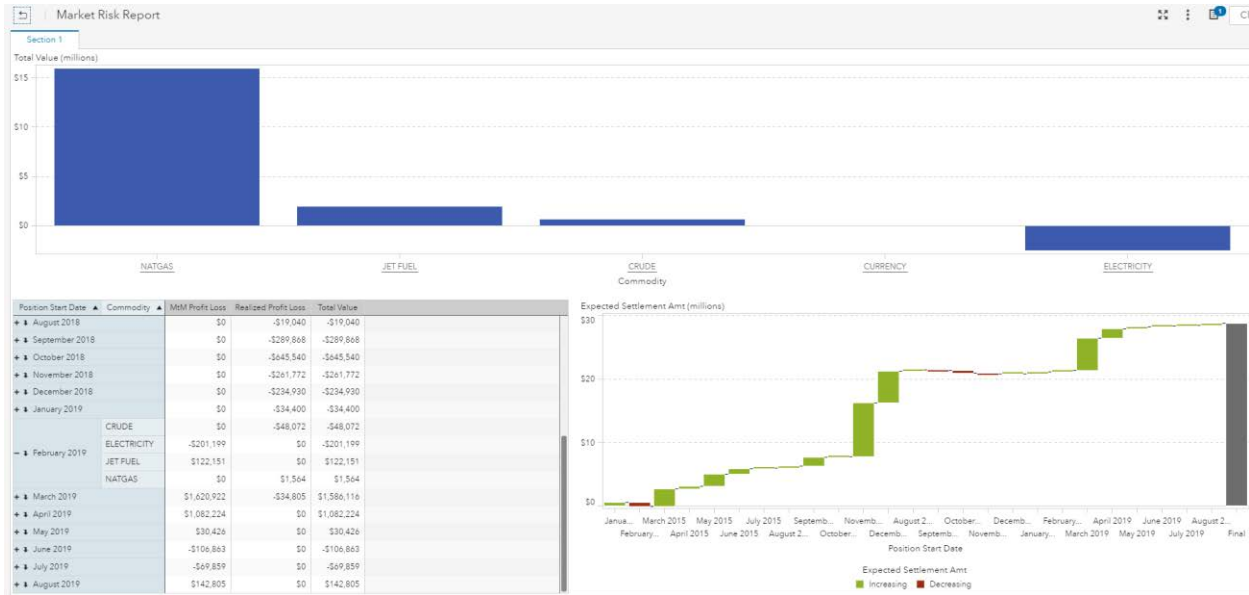

Figure 6 Market Risk Report Example

For example, Figure 6 shows MtM in both a table and graph, while the corresponding cash flows are displayed in a waterfall chart on the bottom right.

This working example demonstrated a set of easy steps to configure a portfolio for processing risk analytics. Static data was configured, transaction data was entered, market data was maintained, and analytics were produced. A range of other key benefits help the solution provide an easy and intuitive way to manage commodity risk.

OTHER KEY BENEFITS – SCALABILITY, FLEXIBILITY & TRANSPARENCY

SAS® Commodity Risk Analytics is built with several key concepts in mind that allow for seamless customizability as the solution is implemented:

- *Robust Data Model* – Clean, accessible data is critical to modeling, risk metrics, and reporting; "garbage in, garbage out." SAS maintains consistent data quality that is validated as it is entered into the system. The data model is structured in a way that is easy to explore and analyze.

- *Comprehensive Auditing* – With growing focus on traceability, transparency, and security, risk solutions are constantly evolving to capture more in terms of auditing and logging. The SAS solution provides insight into activities beyond an audit history of entered and edited data to include visibility into analysis information, system health, and user access/capability reporting.

- *Advanced Integration* – With potentially millions of records flowing in and out of the system on a daily basis, the solution provides a robust series of integration points, web services and upload capabilities to ensure that data is available when it is needed.

EXPAND

Risk teams need to capture, analyze, and report on large amounts of data within short time frames, so automation is a valuable asset, freeing business and technical resources to focus on more value-added activities. Using robust integration technology, the risk solution enables diverse data from multiple locations to flow in and out of the system with little maintenance once initial configuration is completed.

The solution expands the use of its data to service both front- and back-end processes. Position, pricing, and credit exposure data from multiple systems can be compiled and surfaced seamlessly within the solution to create aggregated risk metrics at the business unit or enterprise level. Consolidated or disparate outputs can be automatically integrated to feed other systems, such as passing settlement data to invoicing systems or surfacing trading data and reports in spreadsheets and other business tools.

ENHANCE

As part of its robust data model, SAS® Commodity Risk Analytics supports a modular approach that can be dynamically augmented for additional requirements. One of the most popular requests during implementations is for credit scoring and credit approvals. While one customer might score their counterparties primarily on quantitative financial metrics, another might develop a comprehensive scoring model that varies by volume of business conducted with each counterparty. Yet another might simply subscribe to an external credit rating service. The solution can accommodate each of these approaches and seamlessly transition between them over time as the business evolves.

EXTEND

SAS® Commodity Risk Analytics is driven by a sophisticated risk engine that balances analytical power with usability. The risk engine can rapidly calculate user-defined risk metrics and generate outputs in near real time. For example, additional controls can be extended into out-of-box VaR capabilities to include user-supplied correlation values for positions such as spread options. The solution can implement logic that allows correlations to be derived from internally generated values, if a correlation is not supplied, or from current observations in an existing correlations file. The flexibility to use either pre-built valuation or user-defined values allows risk users of all levels to obtain meaningful results.

SAS® Commodity Risk Analytics can also be extended for applications beyond traditional commodity risk management, and even beyond conventional risk management uses. Although risk is a focal point of most implementations, users quickly realize the value that their investment can contribute to other business areas. For example, midstream companies commonly calculate best use of assets through a blending strategy. Rather than engaging in **what can be a heavily manual process, the solution's capabilities can be directed to** build a model that determines the optimal blend of condensate and different qualities of crude. In additional to affecting profitability, this extension of the solution improves the integrity of previously manual processes and enhances the accuracy of the strategy by incorporating additional metrics such as network parameters, costs/tariffs, and quality constraints.

ENLIGHTEN

Whether drilling down to risk positions, exploring source data or performing additional calculations, risk professionals want reliable information at their fingertips. The ability to run ad hoc analysis that yields valuable results is paramount. The SAS solution has responded to the growing demand for additional insights by providing data management, ETL, visualization, and other features that enable a self-service approach backed by concrete data and valuations. It accommodates users of all levels by including both coding

capabilities and a click-and-drag interface for non-technical users. For example, a small risk group might generate **ad hoc analysis on each trader's open position**s while simultaneously fielding requests from counterparties for up-to-date credit data.

CONCLUSION

The challenge can be simple to explain, but it is not trivial to address. The challenge is best met by a single solution that can pragmatically address basic requirements and provide quick time to value while also providing the advanced analytical capabilities needed to manage the complexity inherent in the commodity industry. Big data requirements and complicated math are handled by the SAS® High-Performance Risk engine, which is in-memory and scalable. The intuitive user interface allows for data to be entered and organized in a logical manner. The results of the analytics can be visualized via SAS® Visual Analytics technology, allowing a focus on producing value from the results rather than troubleshooting data errors generated by a black box solution.

Understanding and communicating commodity risk analytics can be easy with the right solution and support. SAS® Commodity Risk Analytics delivers a scalable, transparent, flexible, and governed environment that can achieve the main objectives of risk teams around the globe. Designed for both technically savvy quants and non-technical users, its dynamic interface and flexible user experience fuel collaboration and innovation. Critically, this flexibility is grounded in reliable, robust data and functionality that delivers value out of the box while remaining highly customizable.

Simple and out-of-box versus advanced and flexible. Is it possible to get the best of both worlds? Some might say it is not only possible, it is easy.

REFERENCES

Bernstein, Peter L. (1998). "Against the Gods – The Remarkable Story of Risk." John Wiley & Sons, Inc pp 45-51, 61-63, 102-107.

Samuelson, P. (1965). "Proof That Properly Anticipated Prices Fluctuate Randomly." Industrial Management Review 6: 41.

Taylor, Katherine. (2017). "Managing Real-Time Data Streams to High-Performance **Analytics Engine." SAS Global Forum 2017, Cary, NC: SAS Institute Inc. Available at** http://support.sas.com/resources/papers/proceedings17/SAS0563-2017.pdf.

Kulkarni, A., Hopping, A., and Xiang, L. (2015**). "**Real-Time Risk Aggregation with SAS® High-Performance Risk and SAS® Event Stream Processing**." SAS Global Forum 201**5, Cary, NC: SAS Institute Inc. Available at https://support.sas.com/resources/papers/proceedings15/SAS1958-2015.pdf.

Chen, W., Skoglund, J., and Iyer, S. (2014). "Effective Risk Aggregation and Reporting **Using SAS." SAS Global Forum 2014, Cary, NC: SAS Institute Inc. Available at** http://support.sas.com/resources/papers/proceedings14/SAS193-2014.pdf.

ACKNOWLEDGMENTS

RECOMMENDED READING

- *SAS® High-Performance Risk Procedures Guide*

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the authors at:

Albert Hopping
SAS Institute Inc.
Albert.Hopping@sas.com

Eric Carriere
SAS Institute Inc.
Eric.Carriere@sas.com

Michael Busby
SAS Institute Inc.
Michael.Busby@sas.com

Christine Jackson
SAS Institute Inc.
Christine.Jackson@sas.com

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are trademarks of their respective companies.