Strategic Implications of Commodity Risk

Optimizing Business Strategy in Light of the Commodity Risk Faced by Electric and Natural Gas Companies

The NorthBridge Group

About The NorthBridge Group

The NorthBridge Group is a leading economic and strategic consulting firm serving the electric and natural gas industries, including both regulated utilities and other companies active in the competitive wholesale and retail markets. We apply market insights, rigorous quantitative skills, and regulatory expertise to complex business problems, always seeking to preserve and build our clients' shareholder value. A unique set of attributes allows us to create value for clients:

- Experienced, senior staff with deep industry and analytic expertise
- Focus on practical, high-impact solutions to real world problems
- Ability to integrate across regulated and competitive functions
- Collaborative working style

All our work is tailored to the individual needs of each client, but the main areas of our practice focus on helping clients develop strategies to deal with the problems and opportunities raised in eight broad areas:

- Risk Management and Finance
- Forecasting and Wholesale Strategy
- Asset Valuation and M&A
- Market Structure and Transmission
- Environmental Risk and Opportunity
- Pricing and Cost Recovery
- Regulatory Strategy
- Expert Testimony

Founded in 1992, we have approximately 25 consultants who serve a national client base from our office located just outside Boston in Concord, Massachusetts.

Contact Information

Contact:	David Coleman Principal The NorthBridge Group 30 Monument Square Concord, MA 01742
Phone:	(781) 266-2632
Email:	dcc@nbgroup.com
Website:	www.nbgroup.com

Table of Contents

Why It Is Critical to Anticipate Market Uncertainty

The U.S. energy sector is characterized by enormous amounts of market uncertainty, including uncertainty related to commodity prices, interest rates, capital costs, and environmental compliance costs. Both producers and consumers of energy operate within a world of continually changing markets, where the bounds of what is considered plausible are stretched daily. Events in recent years have demonstrated that 'Tail Events', or extreme outcomes, happen with surprising regularity. For example:

- In 2008, forward prices for natural gas fell by over 50% in the course of six months,[1] creating catastrophic conditions for some power producers. At the same time, credit costs skyrocketed and created a financial double whammy for entities that needed to post collateral for out-of-the-money forward positions.
- Copper prices rose over 90% during the year leading up to June 2006, remained relatively 'stable' for two years, and then dropped over 30% between September 2008 and March 2009.
- In August 1990, Iraq invaded Kuwait and crude oil prices jumped almost 40% in a single week.
- During the three months preceding November 2008, crude oil prices rose from under $70/barrel to just short of $100/barrel – an increase of over 40%.
- NO_X allowances which once traded at over $3,000/ton dropped to as little as $25/ton in only a few years due, in part, to changes in environmental regulations.
- During the financial crisis of 2008, many companies were simply unable to fund working capital through the financial markets, or found that the cost of financing working capital increased hundreds of basis points. Credit spreads on BBB corporate debt rose from 185 basis points in 2007 to 404 basis points two years later – an unprecedented increase.

What makes these events notable is that they fell so far outside the realm of historical experience that risk analyses and business strategies constructed a priori generally failed to include them as plausible outcomes. Strategic planning often focuses on predicting what <u>will</u> happen and too often fails to consider what <u>could</u> happen. No one would have expected these events would take place when they did, but neither should anyone have been entirely surprised that such dramatic changes in commodity prices and market structure could happen, and in ways that might prove disastrous to an otherwise well-reasoned business plan. In order to ensure long-term growth, a company that is financially exposed to the energy markets, interest rates, inflation, etc. must adopt sophisticated approaches to anticipate uncertainty and risk.

[1] Annual Contract for delivery during 2009 between June 27, 2008 and December 19, 2008.

Both power producers and consumers face financial risks due to exposure to market uncertainty.[2] In the absence of firm contract pricing, a power producer faces revenue uncertainty due to the volatility in wholesale electric prices. A power producer may also be exposed to substantial fuel cost uncertainty, as well as risks associated with operational performance, market structure (e.g. changes in capacity price calculations), and existing and possible future environmental regulations and legislation. Furthermore, a power producer's portfolio risk (e.g. aggregate net revenue) is often impacted by the complex correlations (or lack thereof) between fuel and electricity prices at different locations and is often incorrectly assessed. While prudent hedging strategies serve to reduce risk, the lack of liquidity and market depth associated with energy commodity contracts makes it impossible to hedge these risks fully.

Consumers, such as Local Distribution Companies (LDCs), and ultimately their retail customers, are also subject to market uncertainty. If an LDC adopts an electricity procurement approach in which it manages a portfolio of the component parts of the overall supply obligation, it is then subject to volatility in the prices of block energy, capacity, ancillary services, renewable energy credits, balancing spot energy, etc. In addition, the LDC must also contend with uncertainty pertaining to customer loads, customer switching/aggregation, energy efficiency, etc. LDCs that seek to have third parties assume the bulk of these risks (e.g. through full requirements supply solicitations) face uncertainty about the prices that they must pay their third-party suppliers in the future. LDCs must have a thorough understanding of market uncertainty in order to determine the magnitude of risk to which their customers are exposed and to demonstrate to regulators that their recommended procurement approach represents the best balance of cost and risk for customers.

A wide variety of key questions stem from concern about market uncertainty and about commodity prices and relationships:

<u>Revenues and Earnings</u>

- ✧ How uncertain are the company's net revenues (*i.e.*, gross revenue less fuel and allowance costs)?
- ✧ How wide should earnings guidance bands be?
- ✧ How volatile is the company's cash flow and who are the natural investors for this type of investment?

<u>Debt, Dividends, and Working Capital</u>

- ✧ What debt levels are sustainable given some acceptable level of default risk?
- ✧ Will increasing the dividend payout impede our ability to fund growth projects?
- ✧ What is the likelihood that our need to access capital markets will coincide with a general tightening of credit conditions?

---

[2] Consumers may refer to end-use customers with direct exposure, or to the entities such as Local Distribution Companies which participate in wholesale markets on behalf of their ratepayers.

Hedging and Risk Reduction

- How effective are 'dirty' hedges (e.g. selling gas forwards to hedge electric price uncertainty) at reducing my revenue risk?
- How large is the diversification benefit of owning generating assets in different regions?
- How should supply be procured for the customers of an LDC?
- Can acquiring assets in different business segments create a natural hedge?
- Are options (e.g. calls and puts) an attractive risk reduction tool? Can we create a 'synthetic' option cheaper than one purchased over-the-counter?

Capital Investment Opportunities

- Is there 'option value' in accelerating or deferring one-time capital investments?
- What is the likelihood that an in-the-money proposed new capital investment will be out-of-the money by the time construction is completed?
- What is the appropriate cost-of-capital for a project with uncertain cash flows?
- What 'off-ramps' can be built into a project so that plans may be modified as new information becomes available?

This paper introduces The NorthBridge Group's unique analytical expertise in developing strategy recommendations for executives facing real-world challenges stemming from market uncertainty. Several of the largest companies in the electric and natural gas sectors have hired NorthBridge to help them formulate high-level strategic plans when market uncertainty is an inescapable aspect of their business. The expertise and cutting-edge tools described in this paper have been integral to those efforts.

Significant Problems Associated With Common Approaches to Assessing Risk

In many cases, the methodology used by a company to assess uncertainty and risk is inadequate – either it fails to identify outcomes which could derail an otherwise sound business plan, or it fails to provide sufficient information for managers to structure an optimal strategy. Often the company does not even recognize these inadequacies and the resulting sub-optimal business decisions until it suffers an unforeseen catastrophic outcome or a lost profit opportunity. Many of these methodologies can be categorized as sensitivity analyses. However, there are also significant problems with the more complex probabilistic and/or simulation-based methodologies as they are commonly employed. In this section, we will briefly discuss the problems associated with both types of commonly used approaches.

Problems Associated With Sensitivity Analyses

Many risk analyses consist of constructing simple high/low cases of outcomes for the most important value drivers such as natural gas prices or electricity prices. These types of sensitivity analyses present the future in simple, discrete states-of-the-world in order to illustrate that future outcomes may deviate from today's expectations:

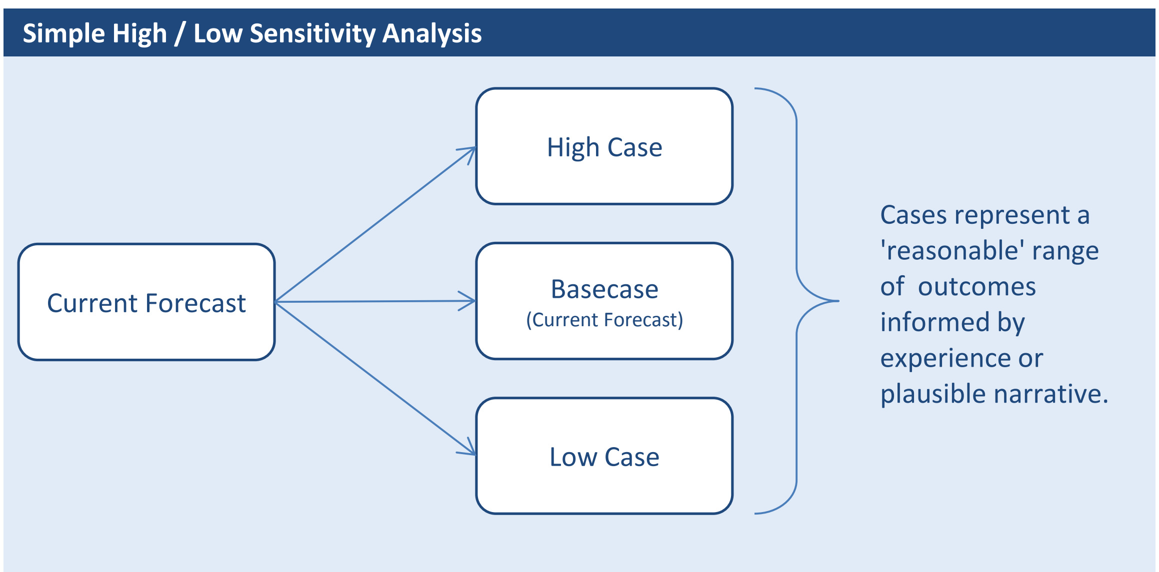

Unfortunately, this type of approach, even when it includes several cases, is often inadequate. It may provide some information about the sensitivity of an important metric (e.g. earnings, mark-to-market exposure, etc.) to a driver (e.g. natural gas prices), but it often provides very little other useful information and may even be misleading. For example,

- ✧ What is the likelihood of each outcome? Does the lack of further detail suggest they are equally probable?

- Do the 'high' and 'low' cases represent upper and lower bounds, respectively? If not, how might the company be affected by even more extreme outcomes?
- Is the impact to the company path-dependent? (*i.e.*, does 'how we get there' matter just as much as 'where we end up'?)
- How do correlations (e.g. between prices and customer demand) impact the relative likelihood of different scenarios?
- What variables are not varied in the different cases? Might they represent a 'hidden risk'?

Even in the simplest case, say when we are only concerned with the risk associated with the price of a single commodity, an analyst attempting to illustrate risk using such high/low sensitivity cases generally relies heavily on subjective and ultimately arbitrary judgments about their relatively likelihoods. In more complex situations, as most situations are, there are innumerable permutations of potential paths and outcomes of different variables; this often makes sensitivity analyses simply intractable.

Sensitivity analyses are often simplified to the point where they provide little useful information to guide a business strategy. For example, imagine constructing a sensitivity analysis that is designed to determine whether someone should purchase a health insurance plan. We identify the possible outcomes: In the 'high' case, we assume that no ailments occur and therefore no medical treatment is required. In the 'low' case, we assume that a serious illness, resulting in many thousands of dollars of medical treatment. Both scenarios are valid in the sense that they are possible, but they are far from adequate in answering the relevant question.

First of all, as the analysis is presented, we have no basis for assessing the likelihood associated with each case, which is critical to the decision. Further, the cases ignore many other scenarios in which other types of medical treatment are needed, with differing amounts of cost. They also ignore possible outcomes that may be less probable but are even more extreme and/or are overlooked. Even when several additional cases are added, this type of case-based analysis invariably omits key considerations. This error-by-omission may mislead decision-makers to inferior decisions while providing a false sense of security.

Risk analyses like the one above are often conducted in the energy industry and are just as flawed. A company with exposure to natural gas prices may produce a risk analysis in the same way, by creating cases of high and low outcomes for key variables and identifying the resulting earnings, cash reserves, etc. in each case. Just as in the health insurance example, this type of approach overlooks or ignores information that is vital to the decision.

Some companies attempt to attribute probabilities to various possible scenarios. When properly performed, probabilistic analyses may provide valuable insights. However, probability assignments are often entirely subjective, or are based on rudimentary analyses of market data which fail to capture the complex dynamics of the markets. Subjective judgment from experienced experts is extremely valuable in understanding risk,[3] but overreliance on such judgment may ignore hidden information revealed by a sophisticated analysis of market data. Further, experts may be biased by their own experiences and dismiss outcomes that fall outside of their own historical observations. Numerous studies indicate that people, even experts, tend to underestimate uncertainty.

The recent history of the energy industry is full of meaningful events that were either so extreme or otherwise improbable that they likely never appeared in anyone's sensitivity analysis. Several of these events have already been discussed: the 53% drop in forward prices for natural gas in 2008, the 40% increase in crude oil prices over one week when Iraq invaded Kuwait, the increase in crude oil prices of over 40% over three months in 2008, drops in NO_X allowances from over $3,000/ton to $25/ton, and the financial crisis of 2008. Sensitivity analyses are correct in illustrating that outcomes are uncertain, but they generally fail to acknowledge just how far outcomes can deviate from expectations. Extreme events do happen and are often a key determinant of a business strategy's success or failure.

Outcomes in the real world are better represented by a wide continuum of values, where the highest probability outcomes are those closest to current expectations and the lowest probability outcomes are those furthest from current expectations (*i.e.*, on the 'tails'):

Full Range of Future Outcomes

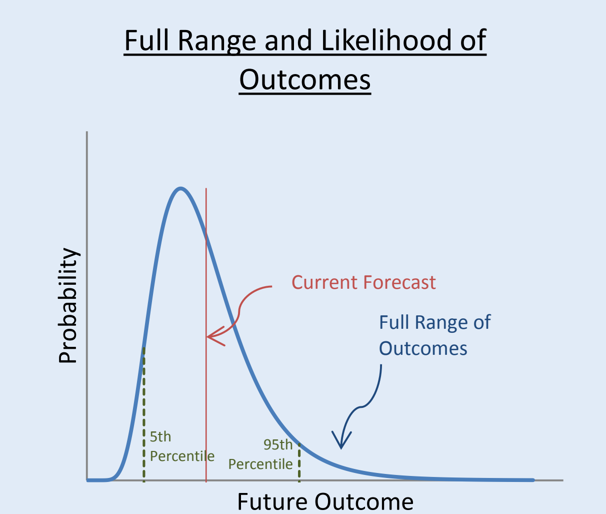

Likelihood of Joint Outcomes

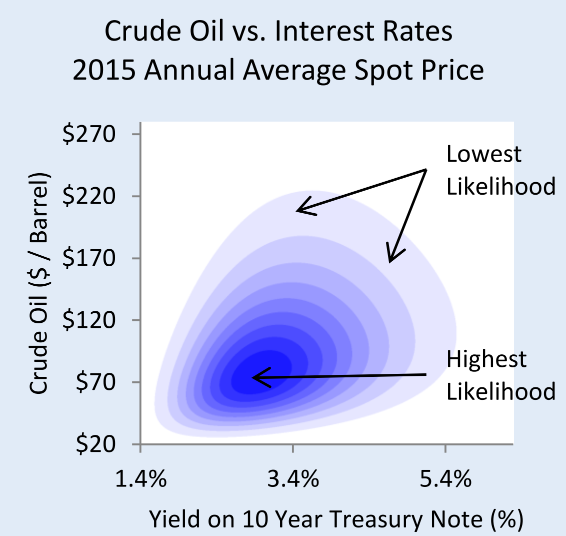

[3] This is particularly true in new or evolving markets where historical data may not exist or may not be a reliable predictor of the uncertainty in future outcomes.

In sum, sensitivity analyses are often inadequate in characterizing risk because they do not provide information regarding the likelihoods of various outcomes, fail to describe complex relationships between variables,[4] and often omit consideration of outcomes that are possible, but extreme.

Problems Associated With Probabilistic and/or Simulation-Based Methodologies

Several energy industry companies have adopted risk modeling approaches that have gone well beyond simple sensitivity analysis. Unfortunately, even the most complex risk models, which incorporate a probabilistic assessment of possible outcomes, often do not adequately glean insights from market information and consequently fail to assign appropriate probabilities to various outcomes. As a result, outcomes that could have the greatest impacts on a company's bottom line are often mischaracterized or even omitted.

Furthermore, some companies attempt to simulate many different drivers of supply and demand, and by extension price. This often results in a focus on the minutia rather than on the important drivers of risk and models which are overly complex and difficult to debug and/or verify. Often times, the complexity of these models is falsely reassuring – problems with the models may not be transparent and company employees may believe the models are sound simply because they explicitly model many factors. Further, complex 'fundamental' models are often bulky and inflexible, making them difficult to adapt when managers need to assess new risks or opportunities.

[4] There are certainly occasions in which it makes sense to leave assessments of likelihood to discussion rather than incorporating them in a quantitative analysis. However, in practice this is often done without first have investigated whether a quantitative assessment of likelihoods is possible.

NorthBridge's Unique Expertise and Cutting-Edge Tools

NorthBridge is uniquely poised to assist decision makers in the electric and natural gas sectors with challenging high-level strategy questions. NorthBridge draws upon its consultants' decades of experience working with high-level industry executives and its cutting-edge quantitative tools to address the strategic challenges industry players face today. Our engagements are characterized by rigorous analysis, innovative solutions to problems, and the acceptance and implementation of strategic insights.

Industry Expertise

The energy and utility industries today are at the nexus of competition and regulation. Traditional regulated utilities face pressures in their ratemaking treatment due to visible wholesale market prices and alternatives to owned generation such as demand side resources and purchased power.

Unregulated competitive market participants face both market pressures and challenges to their business from government regulation and changing market structures. Preserving and building shareholder value in this new environment is difficult, and we are adept at helping clients navigate the interaction of competition and regulation at the state and federal levels.

All our work is tailored to the individual needs of each client, but the main areas of our practice focus on helping clients develop strategies to deal with the problems and opportunities raised in eight broad areas:

- ✧ Risk Management and Finance
- ✧ Forecasting and Wholesale Strategy
- ✧ Asset Valuation and M&A
- ✧ Market Structure and Transmission
- ✧ Environmental Risk and Opportunity
- ✧ Pricing and Cost Recovery
- ✧ Regulatory Strategy
- ✧ Expert Testimony

Rigorous Analytical Support

NorthBridge bases its strategic advice on both its wealth of industry expertise and its rigorous analytical analyses. The firm's approach to quantifying market price uncertainty involves applying a flexible, theoretically sound, and very sophisticated method of identifying and gleaning key information that is embedded in actual market data and that is overlooked by other modeling approaches. As such, NorthBridge is better able to reflect the spectrum of possibilities of future market conditions. It has adapted concepts and tools originally developed for quantitative finance for use in the electric and natural gas sectors, allowing it to analyze risks and opportunities that might ordinarily be considered too complex to assess.

One of the reasons why NorthBridge is able to produce useful insight into commodity price uncertainty is its seasoned ability to characterize commodity price behavior as financial time series with complex, but observable, dynamics rather than relying on fundamental models that are necessarily simplifications of the real world. When simulating how market prices might unfold in the future, NorthBridge is careful to account for real-world characteristics such as:

- Mean reversion (*i.e.*, the tendency of prices to trend back to long-term averages)
- Stochastic and conditional volatility (*i.e.*, prices undergo both periods of relative stability and uncertainty)
- Correlations between different commodities
- Linkages between spot and forward price movements
- Likelihoods of extreme outcomes (*i.e.*, 'Tail Risk')

The dynamics of commodity price movements are complex and replicating these dynamics in a simulation requires a sophisticated model of how and why prices change. NorthBridge has had great success simulating wholesale prices using a model that blends two approaches commonly used in quantitative finance: The first model, known as the Geometric Ornstein-Uhlenbeck Model, captures mean reversion. The second model, known as the Heston Model, captures stochastic volatility. Neither of these models individually can replicate all the important characteristics of commodity price movements, so NorthBridge has developed a hybrid model and the necessary proprietary computational tools to simulate realistic price paths and outcomes.

The sophisticated models developed by NorthBridge have the capability to replicate the types of price dynamics observed in the real world, but must first be calibrated or 'fitted' to specific commodities before they can be used to generate insight. The calibration process is in many ways even more sophisticated than the simulation models themselves. NorthBridge has developed an application that performs the calibration using 'Maximum Likelihood Moment Matching', and utilizes both historical spot and forward prices as inputs and can utilize forward looking volatility forecasts, such as option implied volatilities, if available. The calibration tools view historical spot and forward prices as a 'sample' drawn from a complex system. Through extensive simulation and non-linear optimization, NorthBridge's tools are able to identify the most likely structure of the underlying system, given what has been observed in market prices.

The powerful quantitative tools deployed by NorthBridge are used in conjunction with considerable expertise and experience within the industry to provide one-of-a-kind strategic insight. When combined with the firm's vast experience in the energy markets and the regulatory arenas, NorthBridge's sophisticated market modeling capabilities play a key role in the firm's ability to provide useful assessments and insights regarding market price uncertainty and its impacts on a company's overall financial risk.

A Sampling of the Types of Valuable Information/Insight that NorthBridge Provides

In developing its recommendations for its clients who seek strategic guidance in the face of market uncertainty, NorthBridge provides both qualitative assessments of risks based on years of market experience as well as quantitative assessments developed with the cutting-edge tools previously described. With respect to the latter, NorthBridge is able to assess risks effectively and communicate insights in clear and comprehensive ways – clients often comment that this is one of the differentiating aspects of an engagement with NorthBridge. The following is just a sampling of the types of information and insight that a NorthBridge analysis provides.

Sample Engagement: Wholesale Hedging Strategies

A key business risk for a merchant energy company is revenue uncertainty, and one of the principal responsibilities of a chief executive is to communicate the magnitude and character of that uncertainty to stakeholders. Equity shareholders, debt holders, rating agencies, and financial counterparties value this information. Companies which appear to misunderstand their own risk exposure, or which do a poor job of communicating the nature of their business risk to stakeholders, are exposed to adverse business outcomes and jeopardize their access to capital markets. As a result, understanding revenue uncertainty is critical for corporate strategy development and for communications with the company's investors and potential investors. Shareholders may look at several different metrics to assess a company's level of risk, such as earnings-per-share (EPS) guidance bands, EBITDA confidence intervals, Value-at-Risk (VaR), share price volatility, etc. The NorthBridge Group has expertise assisting clients in quantifying risk as measured by many several different metrics, and then going one step further by identifying how different actions might affect that risk. In order to illustrate the types of information and insight that a NorthBridge analysis provides, the following sections of this document discuss applications of NorthBridge's capabilities in the context of EPS risk.

Unhedged Revenue Risk Exposure

The first step in developing an accurate understanding of earnings risk is to simulate how commodity prices could evolve over time, and then show how that uncertainty translates into revenue risk. In the simplest case, where the firm's cash flows are directly exposed to spot price risk, calculating and illustrating EPS bands is straightforward. From today's perspective, near-term earnings are relatively more certain, while earnings in more distant periods are relatively less certain. For example, a firm with long exposure to wholesale electric market prices could face substantial EPS risk:

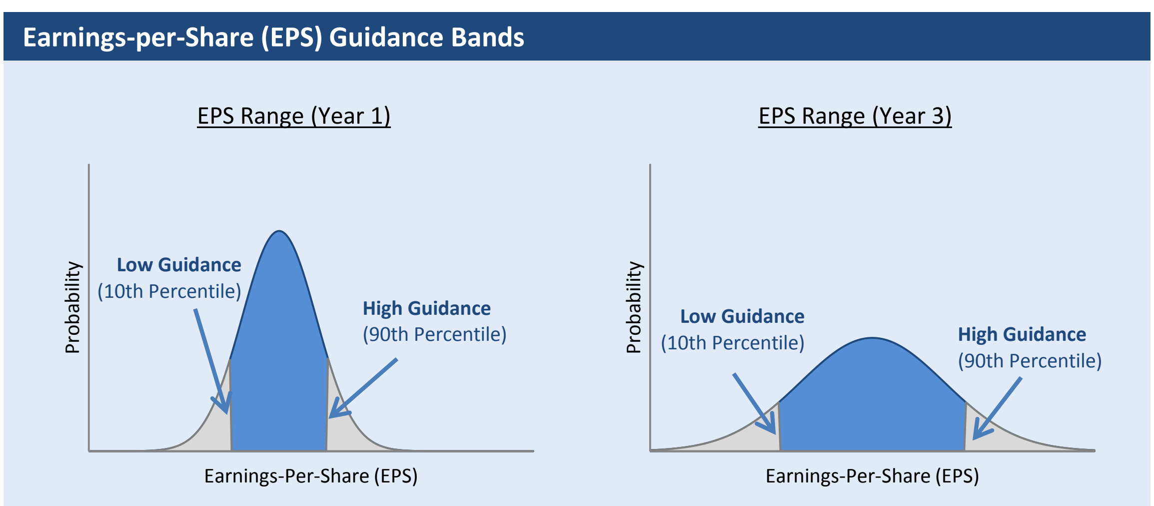

Earnings-per-Share (EPS) Guidance Bands without Hedging (Years 1-30)

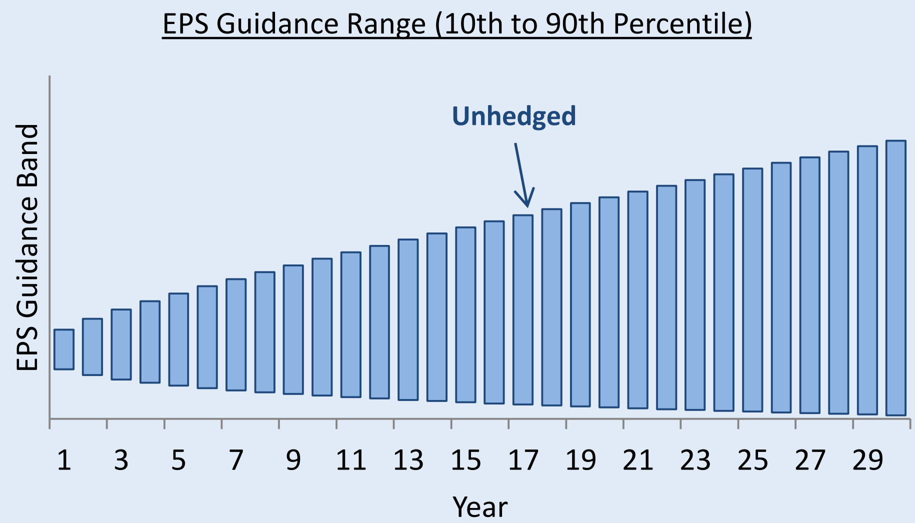

Insight Developed

The charts above illustrate that EPS risk stemming from commodity price uncertainty is substantial, even in the near term, and continues to grow over time. Diligent managers often seek to reduce that uncertainty on behalf of investors using one or more of several tools available to them. Fortunately, NorthBridge has the tools and expertise to help managers understand the efficacy of each approach.

Revenue Hedging with Energy Forwards

Forward markets for wholesale electricity have been operating since the late 1990s. Liquidity for long-term products (5+ years) remains thin at best, but firms may have the opportunity to hedge market price exposure using shorter-term contracts. Consider the firm described above; it has long exposure to wholesale electric markets and may want to reduce its price risk by selling forwards.

Each year the firm's traders sell forward contracts to hedge 1/3rd of the firm's long exposure. This is called a three-year rolling hedge and results in a laddered portfolio of forwards. In practice, this means that output will be sold at a price equal to the average forward price observed over the prior three years. This strategy is attractive because it fixes the sale price of some output up to three years in advance and results in a 'fixed' price for all of the output for the prompt year. This approach has the effect of 'accelerating' the pricing of sales into earlier periods and may reduce uncertainty regarding the prices obtained.

By simulating both spot price outcomes and the evolution of forward contract prices, NorthBridge can illustrate the efficacy of hedging with a three-year rolling strategy. The EPS risk for the prompt year is reduced (practically to zero):

Earnings-per-Share (EPS) Guidance Bands with Hedging (Year 1)

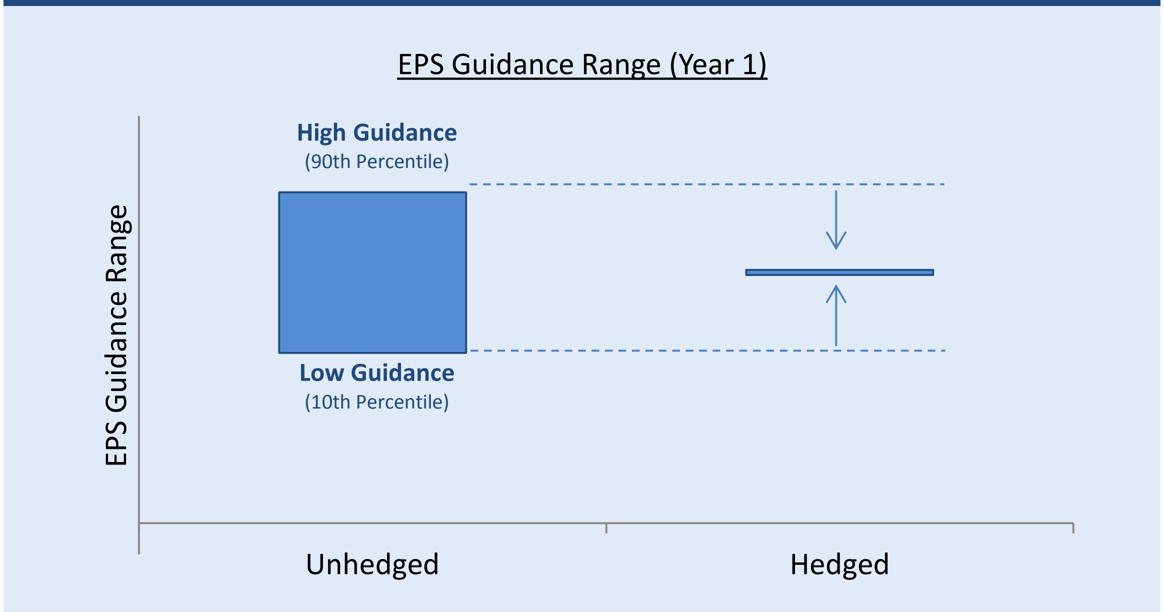

This hedging approach may initially appear to mitigate substantial risk and materially change the nature of the uncertainty to which investors are exposed. However, EPS risk for later years looks very different:

Earnings-per-Share (EPS) Guidance Bands with Hedging (Years 1-30)

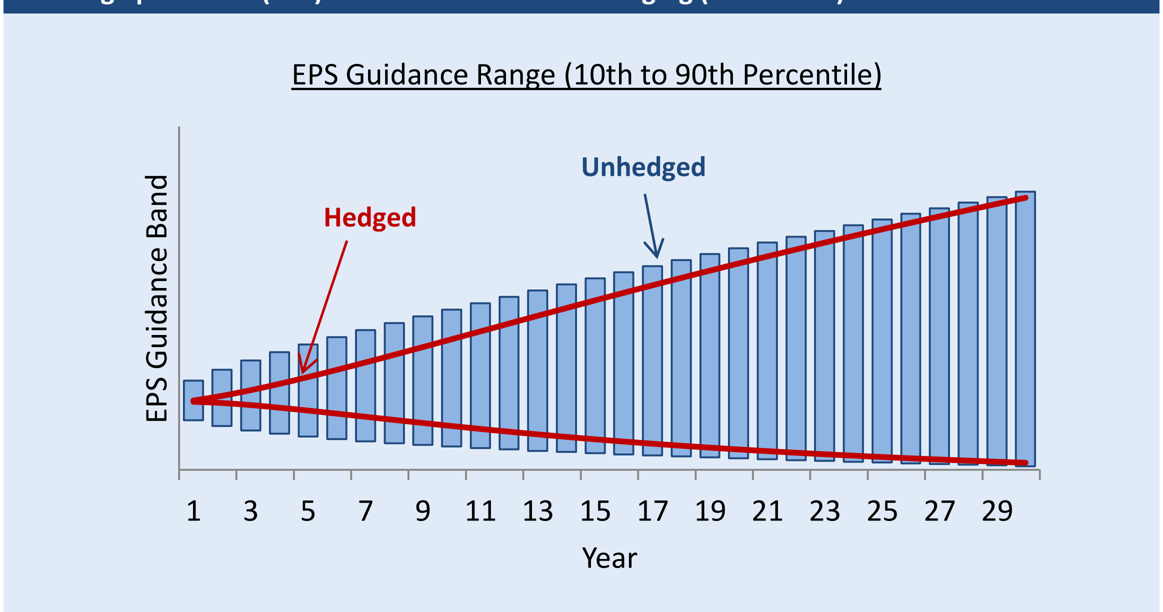

Insight Developed

While all output for the prompt year may be hedged, output for the second year is only 2/3 hedged, meaning that the price for the remaining 1/3 is still uncertain. Similarly, EPS for the 3rd year is only 1/3 hedged; the remaining 2/3 will be priced at forward prices as they evolve over the next two years. As we look out further, the difference in risk exposure between hedged and unhedged outcomes diminishes quickly. For example, in year 30, the rolling hedge allows us to sell at the forward price observed during years 27-29. The program allows us to avoid spot market price uncertainty, but investors are still subject to forward price uncertainty over the next 29 years! The inescapable conclusion is that a company that plans to be fully hedged at some point in the future is still exposed to market uncertainty prior to that time.

Revenue Hedging with Fixed-Price Full Requirements Default Service Supply Products

Solicitations by LDCs for fixed-price full requirements default service supply products may provide power producers with another opportunity to mitigate earnings risk. In these solicitations, producers (and other market participants) bid a fixed price to satisfy the full requirements supply needs of customers who do not choose an alternative retail supplier. The products in these solicitations may initially appear more attractive to producers than exchange-traded futures for several reasons: longer term, lower credit requirements, natural buyer, etc. However, they differ from block forward sales in the sense that the sales volume is not fixed; the customer, an electric LDC in this case, pays a fixed price for only the volume of power needed. When a producer hedges its long exposure by selling a full requirements product to an LDC, it in turn exposes itself to new risks in the form of uncertainty in overall customer demand, customer switching/aggregation, and the correlation between unexpected changes in customer usage and price.

The decrease in hedge efficacy due to customer switching / aggregation is easily demonstrated:

Efficacy of a Fixed-Price Full Requirements Supply Product as a Revenue Hedge

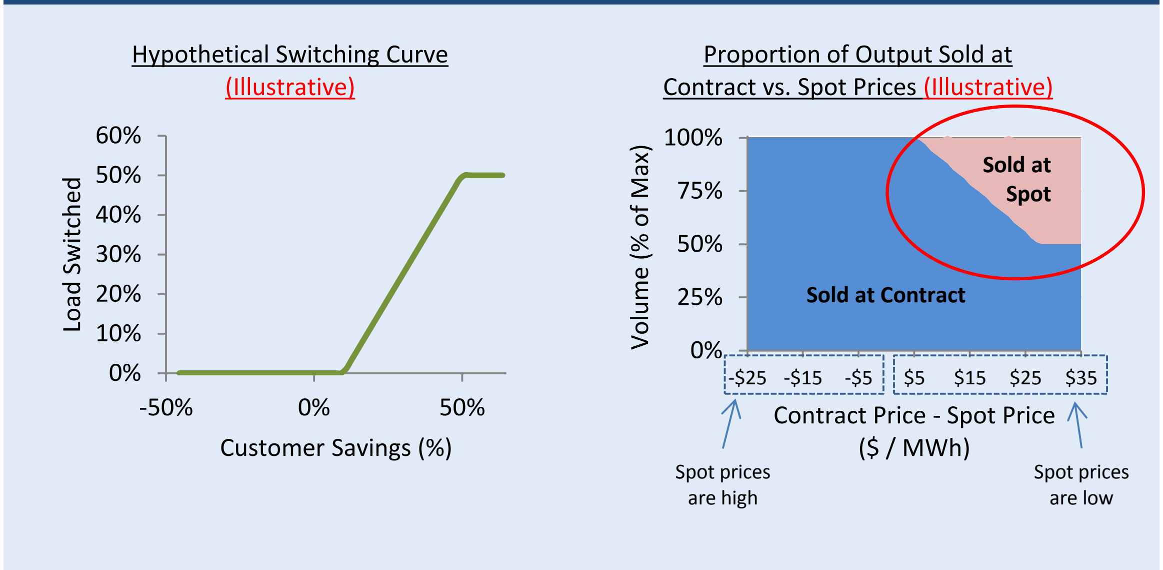

Efficacy of a Fixed-Price Full Requirements Supply Product as a Revenue Hedge

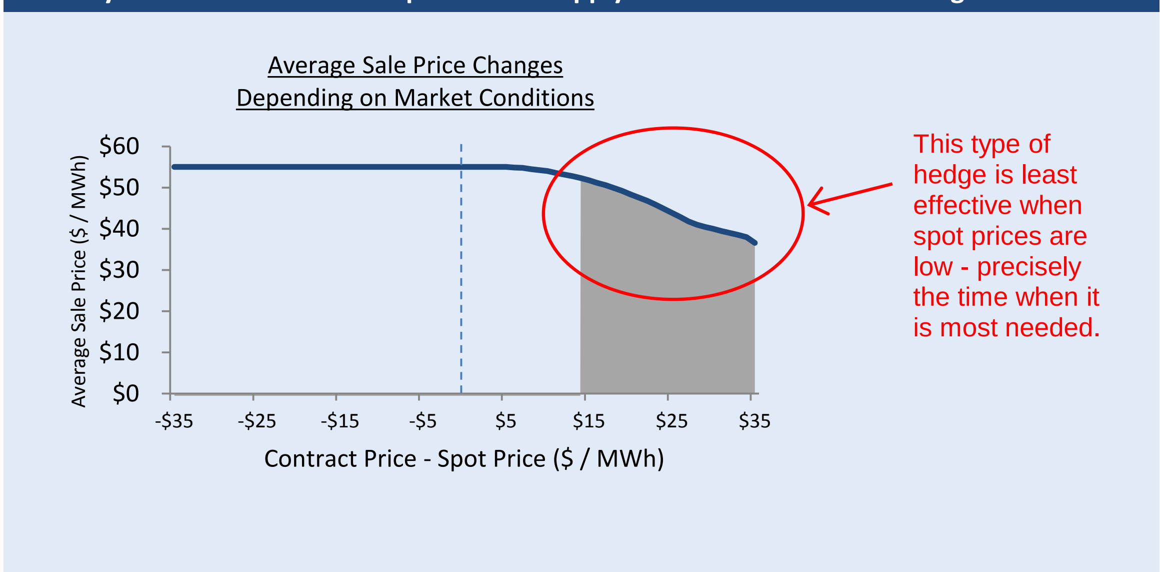

Insight Developed

Customers are often allowed to switch from their default supplier to an alternative supplier, and may choose to do so when market prices have fallen and alternative supply is cheaper than the original fixed price service. Because of customers' financial incentives, the original supplier will lose customers and sales volumes precisely when the hedge is most needed. If customers are reluctant to switch suppliers, sales volumes may remain static under small market price movements. Under large market price drops, however, the hedge may be far less effective than originally hoped.

Revenue Hedging with Put Options

Some wholesale forward markets are sufficiently liquid that options on forward contracts are both available and competitively priced to make them a viable risk mitigation tool. Put options offer protection against low price outcomes while preserving the opportunity to benefit from high price outcomes. The trade-off for this outcome asymmetry is the upfront cost of purchasing the put option. If puts on the exposed commodity are purchased and held to maturity, the EPS risk may be reduced:

Earnings-per-Share (EPS) Guidance Bands with Put Option

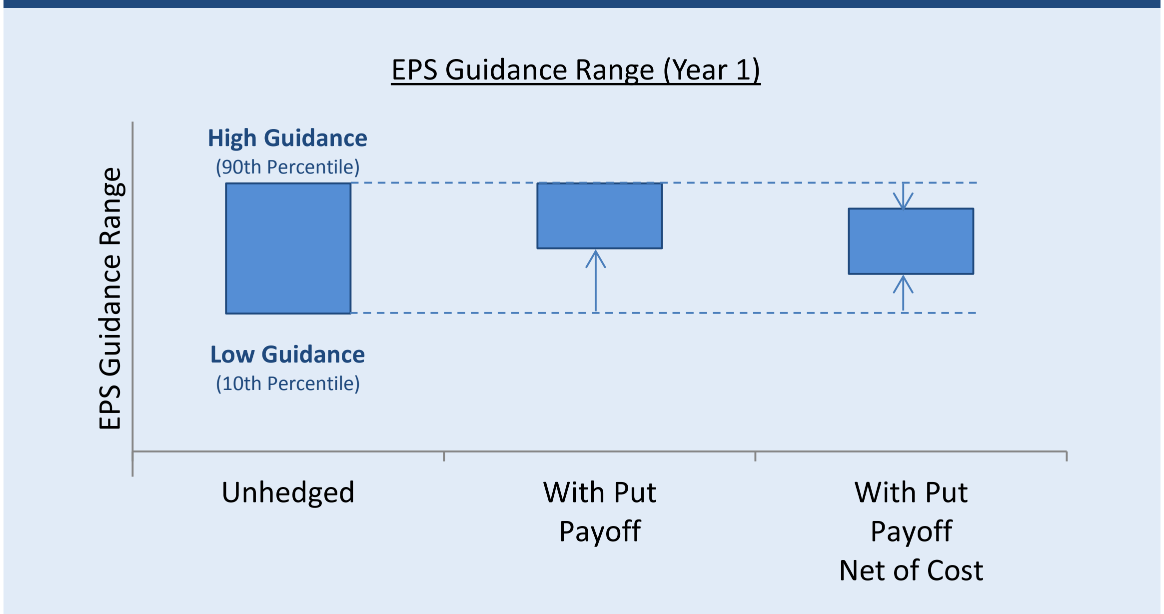

<u>*Insight Developed*</u>

If the put option is held to maturity, low outcomes will be improved and high outcomes will be maintained, with all outcomes scaled downward by the upfront cost of the put option. The initial appeal of the put option – that it preserves upside while protecting against downside – appears to be tempered by the expense of obtaining the option. Options with strike prices close to the current forward price or which are liquidated prior to expiration may actually lead to outcomes not dissimilar to hedging with forwards.

Sample Engagement: Sustainable Debt Levels

Changes in corporate capital structure, such as increasing or decreasing leverage, can create substantial value for a firm's owners. Debt offers considerable tax advantages, which can increase firm value, but can also introduce future cash flow constraints and force management to forgo valuable investment opportunities when revenues fall and cash flow is insufficient to cover interest obligations. Company managers have a strong incentive to demonstrate to both existing and potential debt holders that the level of leverage proposed by management will not create cash flow constraints in the future; they also have a strong incentive to demonstrate to equity holders that excessive debt levels will not drive the company into bankruptcy or destroy firm value by introducing expected financial distress costs. NorthBridge's analytical models can help capture the relationships between debt levels and default risk:

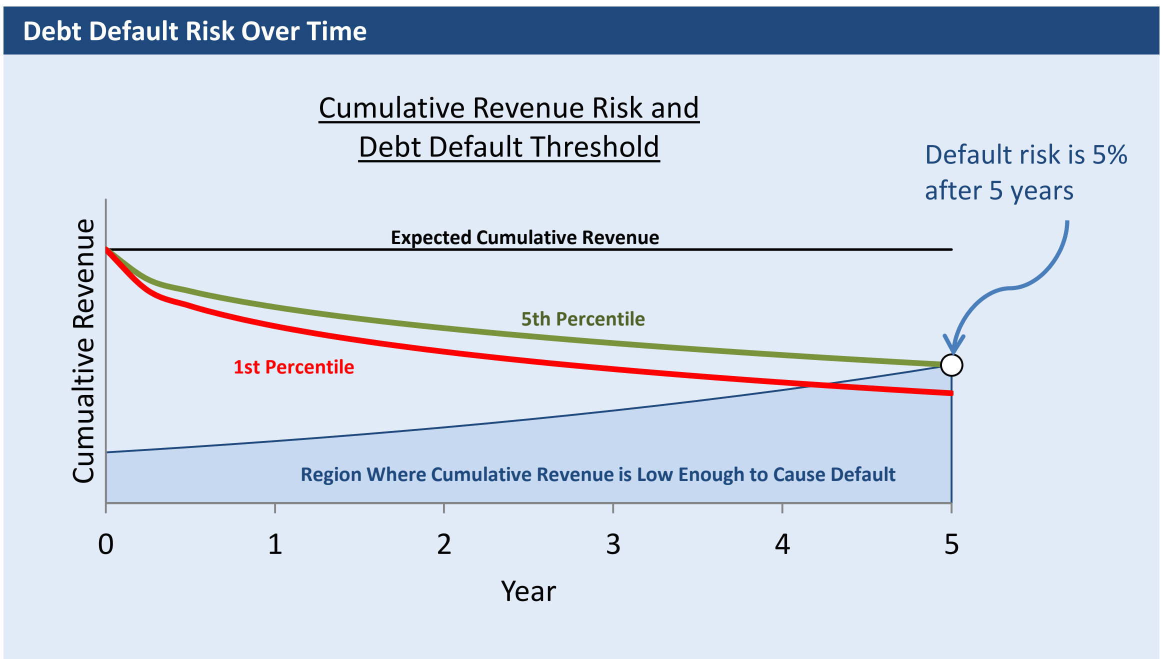

Debt Default Risk at Different Levels of Leverage

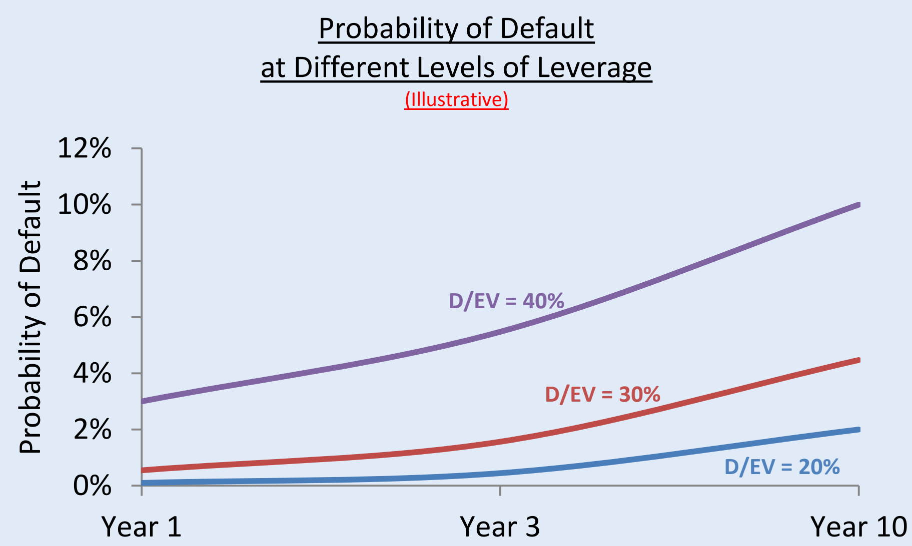

Insight Developed

Commodity prices are often the primary driver of financial health for companies in the energy sector. NorthBridge has helped several management teams identify the likelihood of financial distress at different leverage levels by simulating commodity price outcomes and assessing the probability that the resulting cash flow might be insufficient to avoid financial distress during some future period. Any level of debt greater than zero carries some risk of default, and even companies without long-term debt face financial distress risk if they hold forward positions that could expose them to sizeable mark-to-market obligations. NorthBridge's financial modeling capability allows managers to weigh the cost of default risk (*i.e.*, higher credit spreads) against the tax benefits of higher debt levels to maximize shareholder value.

Sample Engagement: Geographic Asset Diversification

Wholesale electricity markets are not locationally uniform. Though wholesale prices in the long term may be strongly linked to global fuel prices and the cost of building new generating capacity, prices in the short term are influenced by regional supply and demand dynamics and may diverge from prices in nearby regions due to the presence of transmission constraints. The regional aspects of electricity markets may make geographic diversification of assets attractive for power producers.

For example, a power producer with assets in the Mid-Atlantic region may be uncomfortable exceeding a certain leverage level due to uncertainty in cash flow (see prior sample engagement). Short-term price dips could impair the firm's ability to support additional debt. Similarly, a separate firm with assets in Texas may have performed a similar analysis to determine its maximum sustainable debt level. The Mid-Atlantic and Texas wholesale electric markets, however, are sufficiently isolated from each other that the joint cash flow from the combined bundle of assets may support a higher leverage ratio than the assets in either region individually:

Default Risk Reduction Through Geographic Diversification

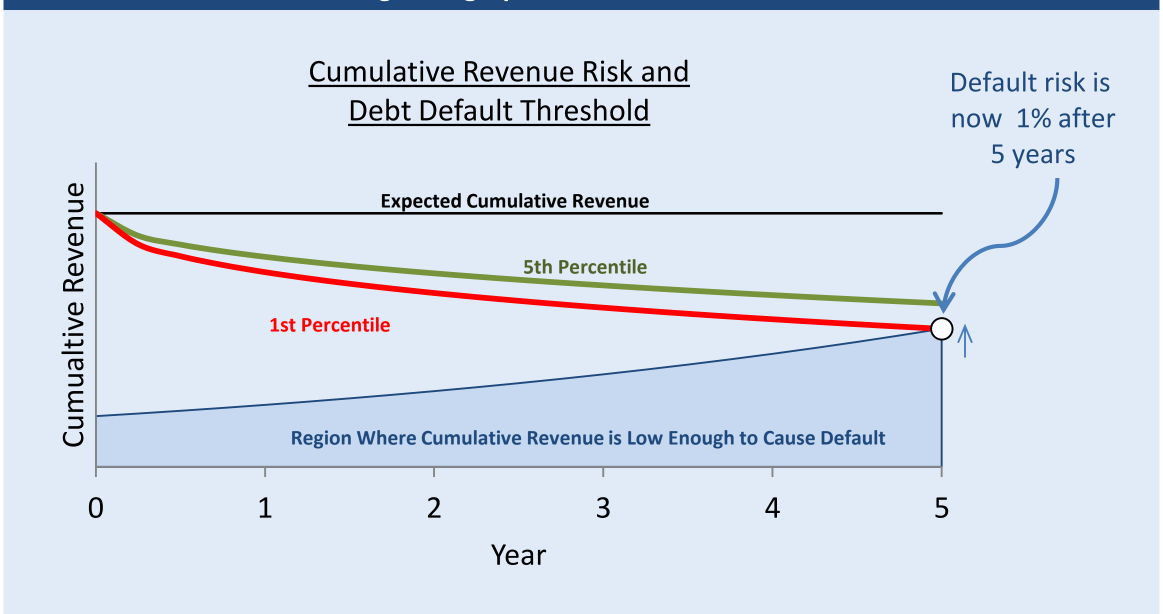

Insight Developed

The chart above illustrates that the likelihood of default at the current leverage level is substantially lower for the bundle of assets than it was for either company individually. Management can use this insight to create value for investors by pooling the cash flows from the assets, increasing debt levels, and capturing the tax benefits of high leverage without increasing credit spreads and other financing costs.

Sample Engagement: Capital Investment Opportunities

The regional nature of wholesale electricity markets can influence the relative desirability of investing in new capacity in one region versus another. When evaluating traditional baseload facilities, regional differences in wholesale around-the-clock (ATC) electric prices and fuel costs are generally the primary consideration. When siting a combined cycle gas turbine (CCGT), additional factors such as the correlation between gas prices and power prices may be a concern. Energy storage investments derive their value from the spread between lower prices during the off-peak hours and higher prices during the on-peak hours, which may vary by region. These regional differences can often be illustrated using contour maps or 'heat maps'. These maps convey regional differences in a way that is easy to visualize and understand:

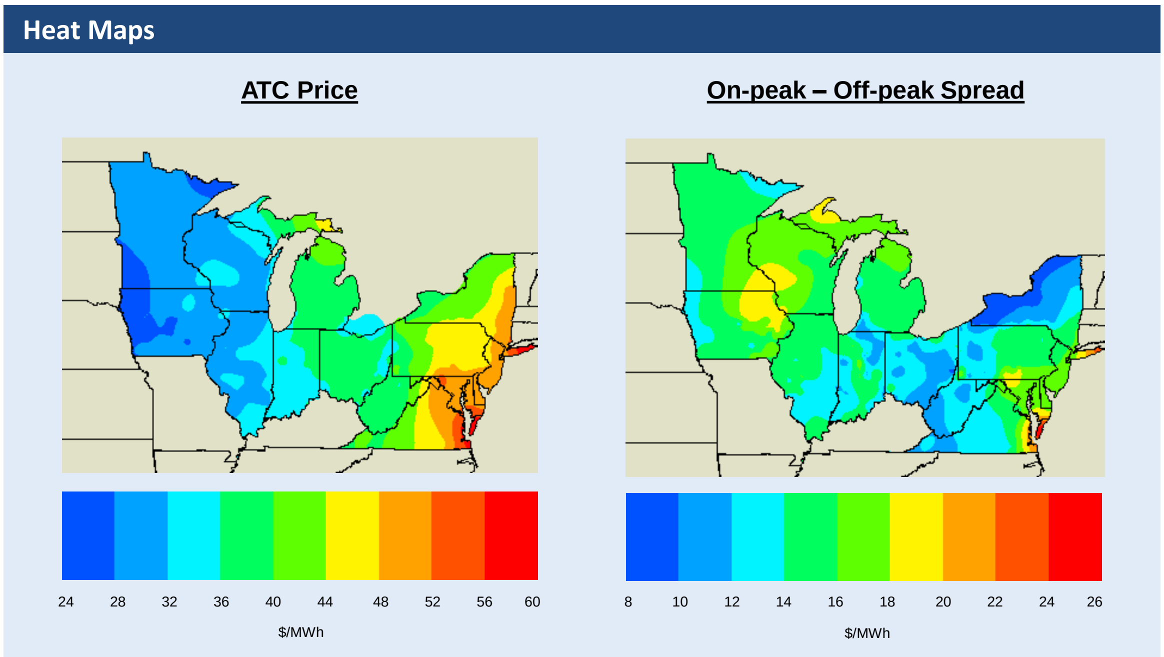

Insight Developed

Markets for both electricity and natural gas have regional characteristics. Investors must not only ask 'Is this investment economic?' but also 'Where is this investment most economic?' A map of absolute price levels indicates that that average prices may be higher in the mid-Atlantic region than in the Great Plains. Spreads between the prices during on-peak and off-peak hours, however, may be greatest in the Great Plains and in some localized pockets elsewhere. This mapping technique can be extended to even more insightful metrics such as project IRR or price volatility.

Sample Engagement: Investment & Retirement Deferral Analysis

Uncertainty in commodity markets does more than create cash flow risk – often it creates opportunities to defer or accelerate management decisions as more information becomes available. Investments in the energy sector often come in the form of large, discrete projects where the scale of the financial investment is a significant portion of the value of the firm. In these cases, the act of investing consumes so much financial flexibility that it effectively precludes the firm from pursuing other NPV positive investments for some period of time. In these cases, managers evaluating the large investment must consider not only the value created by pursuing the investment, but also whether even greater value may be generated by deferring the decision and waiting for more information.

For example, a merchant power producer may be contemplating investing in a new power plant. Forecasts of wholesale electric prices, fuel costs, environmental costs, etc. may suggest that the investment exceeds the cost-of-capital threshold. From a traditional corporate finance perspective, the investment meets the necessary criteria and should be pursued. However, if the firm has a finite capacity to make investments, the investment ought to be viewed as an option where committing to the investment is equivalent to exercising an option prior to the expiration date – it may be 'in-the-money' and have a positive payoff, but there may be a greater expected value by waiting longer:

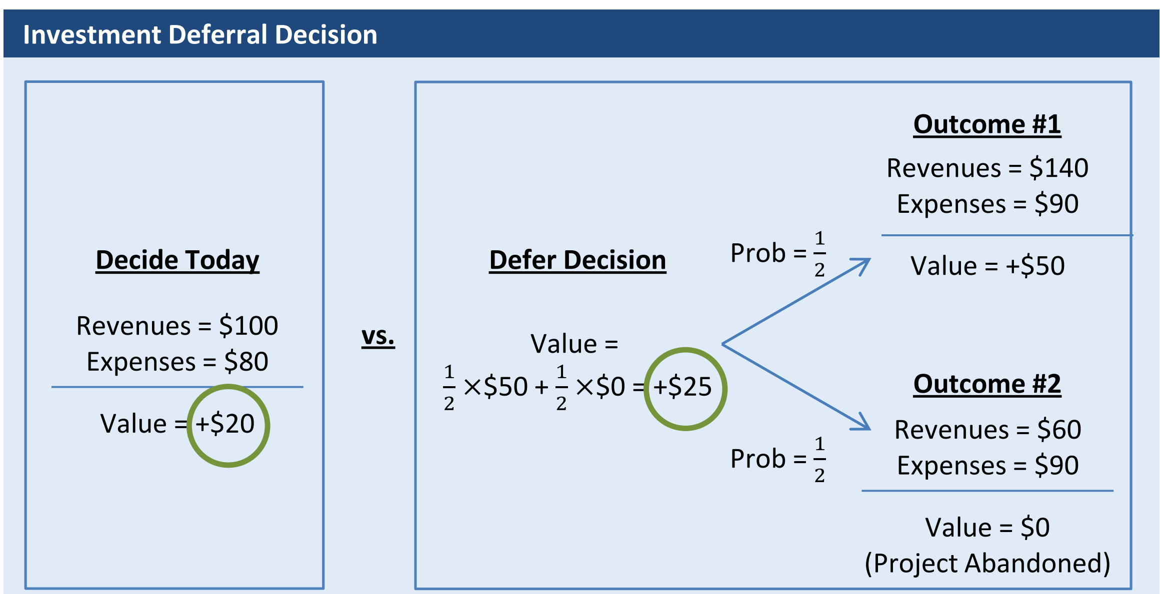

Note: Any actual analysis would examine thousands of scenarios of forward and spot price paths developed using probabilistic techniques.

Insight Developed

Many market uncertainties could lead a company to want to delay committing to an investment. Managers' intuitive understanding that not all (traditionally calculated) NPV positive investments should be pursued is justified when projects are analyzed using the appropriate analytical techniques. With respect to the investment opportunity described above, if managers expect that a major piece of legislation governing carbon regulation is coming close to a vote, there may be good reason to delay the investment until the result of the vote is known. If the result of the vote is to suggest lower carbon prices, the investment may then proceed. If the result of the vote is to suggest higher carbon prices, the project may no longer make sense and might be abandoned. By analyzing the project as an option, we can determine whether or not the value of moving quickly and producing margin quickly outweighs the value of moving slowly and potentially avoiding bad outcomes.

This type of analysis lends analytical support to reasoning that many managers already understand intuitively – <u>investment decisions are not made in a vacuum, but rather must be evaluated in the context of how the decision will either increase or decrease the firm's ability to pursue other opportunities in the future</u>. Other types of capital decisions, such as asset retirements or mothballing, should be analyzed in a similar way.

Sample Engagement: Cost of Providing Full Requirements Default Service Supply When Shopping/Aggregation Is Allowed

Power marketers are increasingly looking to solicitations for fixed-price full requirements default service supply as an opportunity to earn a reasonable return in exchange for insulating retail customers from volatile energy prices, while still allowing those customers to shop for cheaper energy if market conditions change. The optionality provided to customers makes this product difficult to hedge and also makes the appropriate bid price difficult to determine.

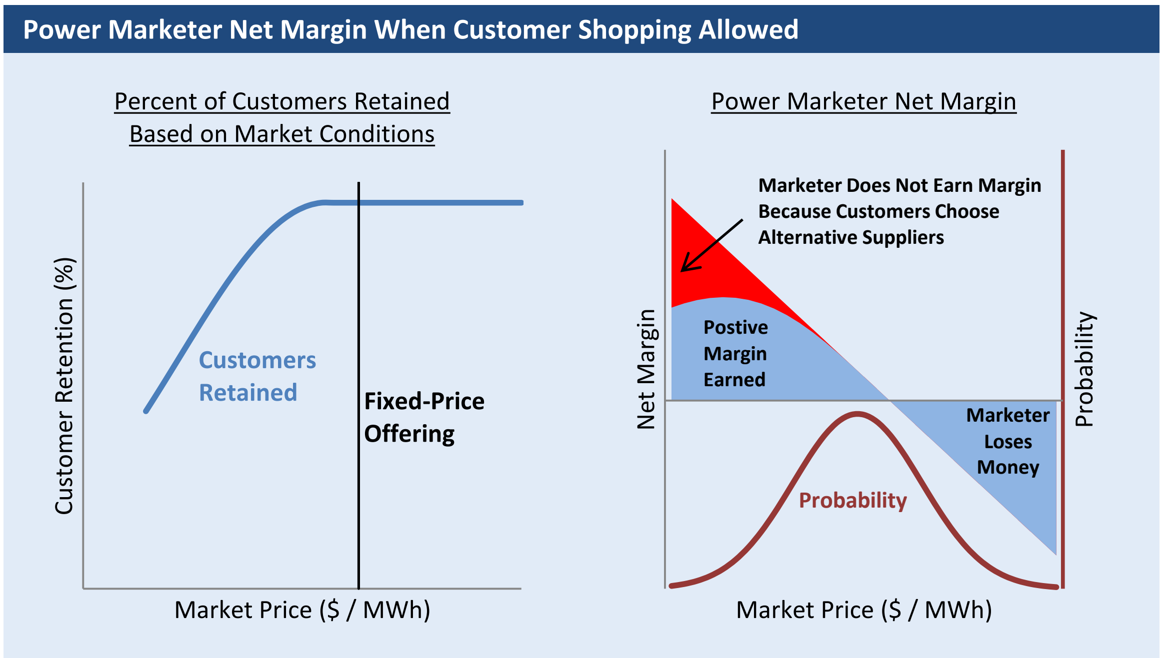

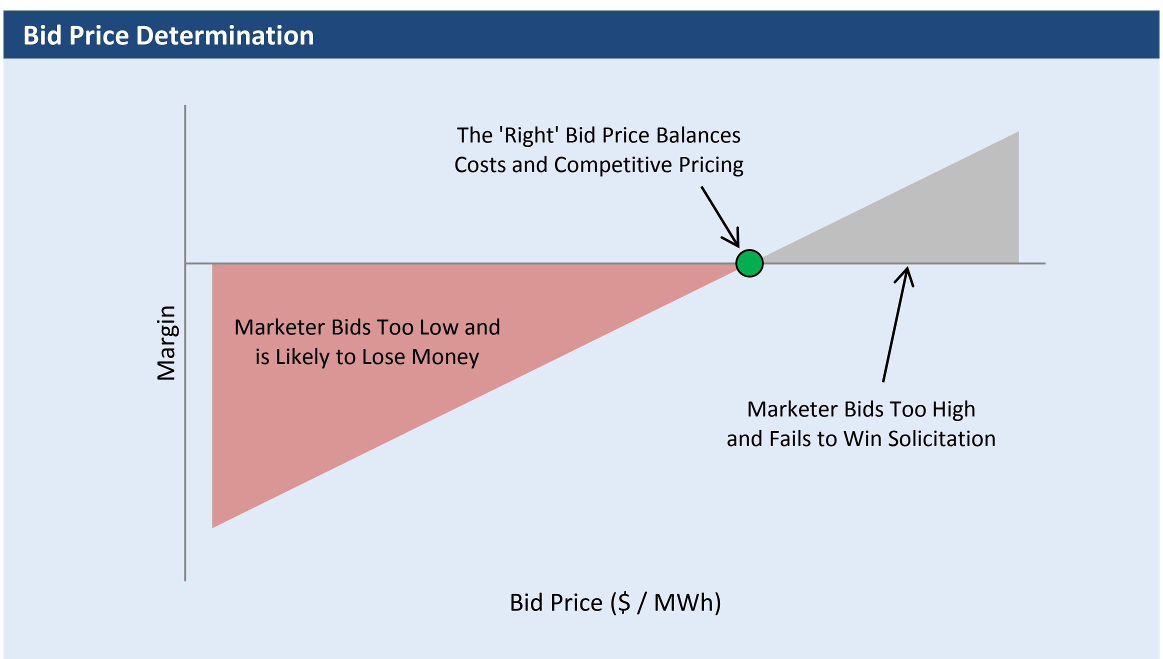

Insight Developed

The product provided by the marketer amounts to an option where the customer will take advantage of changing market conditions precisely when doing so will cause the most financial damage to the marketer. This asymmetry, as well as other characteristics and interrelationships of market prices and loads, has complex and potentially very significant effects on the profitability of providing full requirements supply products.

Power marketers craft their bids to provide this type of supply in such a way that the price customers pay compensates the marketer for the cost associated with providing this optionality to customers. If the bid is too low, the cost of serving the customers will exceed the revenue received. If the bid is too high, the marketer will likely be underbid by another participant in the solicitation. The optimal bid price can be determined such that the marketer is expected to cover its costs and earn a fair margin, but is not expected to be underbid by a competitor.

Sample Engagement: Valuing Unique and Exotic Energy Options

The maturation of wholesale energy markets has ushered in a new era of energy products designed either to help mitigate risk or provide opportunities for advantageous risk sharing. Some products, such as calls and puts on forward contracts have direct analogues in other financial markets. In many cases, transparent markets for these simple (a.k.a. 'plain vanilla') derivatives either already exist or are developing. A firm wishing to buy or sell one of these products can benefit from the transparency of the market to help assess its fair value. But, when the natural market is too illiquid to support robust price discovery, fair value must be determined independently on the basis of forecasted price uncertainty.

For example, consider a firm wishing to purchase a put option on natural gas that will allow the firm to sell natural gas at a pre-specified price and location over a period of several years. Only a few market participants offer such a product and the exact contract specifications may differ between dealers. The first dealer specifies a put option with a strike price of $5/MMBtu for the period covering years 1-3 and asks a price of $0.60/MMBtu. The second specifies an option with the same strike, but the period covers years 1-5 and asks a price of $0.65/MMBtu. Since the product specifications are different, we cannot simply compare the prices to determine which one offers more competitive pricing. The expected value for each contract can be determined using NorthBridge's sophisticated price-path evolution tools:

Expected Value of Put Option (Three Year Period)

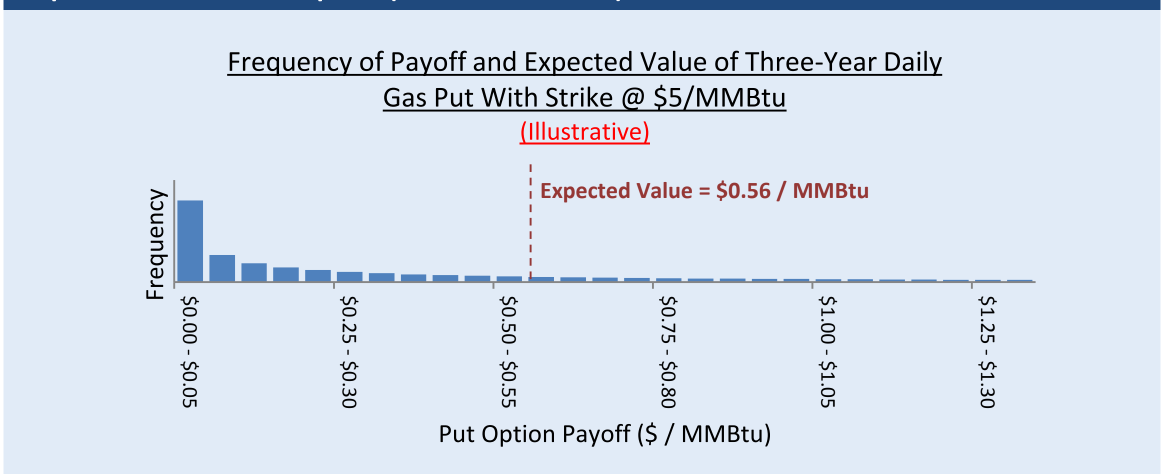

Expected Value of Put Option (Five Year Period)

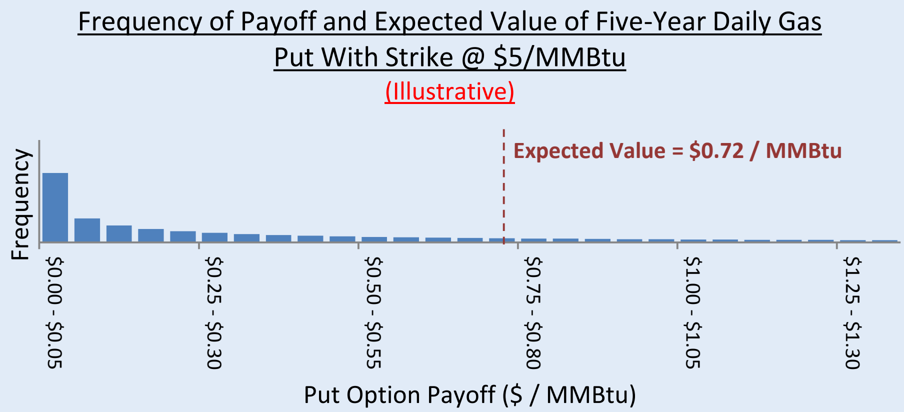

Insight Developed

The simulation of price paths and put payoffs indicates that the price asked by the first dealer is well in excess of the expected payoff of the option, while the second is at or below the expected payoff. This type of analysis does not, in isolation, indicate whether the put option is a good strategic fit for the company's business model, but it does provide the basis for determining if the product is priced fairly.

Sample Engagement: Retail Rate Uncertainty for LDCs Due to Uncertainty in Wholesale Procurement Costs

Local Distribution Companies (LDCs) operate the distribution network that serves retail customers, and provide energy supply to 'default service customers', which are the customers who do not elect to be supplied by an alternative retail supplier. Most electric LDCs in restructured jurisdictions do not own generating capacity,[5] but instead purchase electricity from wholesale providers and then resell that power to their retail customers at a price that is often fixed to some degree. LDCs often purchase 'fixed-price full requirements' products from power marketers, in which all customer requirements are provided at a fixed price. Alternatively, they may manage a portfolio of supply resources and contracts themselves,[6] including purchases of wholesale energy from the forward and spot markets, and either bear the burden of cost uncertainty, or pass that cost uncertainty along to customers. Either way, LDCs must often design and defend a proposed default service supply procurement plan on the basis of the benefits and risks it presents to customers.

When evaluating whether an LDC has proposed a default service plan that best meets the needs of stakeholders, regulators often focus on the cost to customers under the plan, the protection that would be provided to customers if market conditions were to deviate from expected conditions (e.g. if wholesale market prices and/or loads were to spike), the likely opportunities customers would have to shop for lower cost supply from alternative suppliers, and the degree to which customers would be exposed to market price signals. LDCs are often required to show that their proposed plan produces the best cost/risk tradeoff for customers, and NorthBridge's analytical approach and tools, as well as the Firm's experience with default service plans, have proven to be very effective in developing and defending potential default service plans:

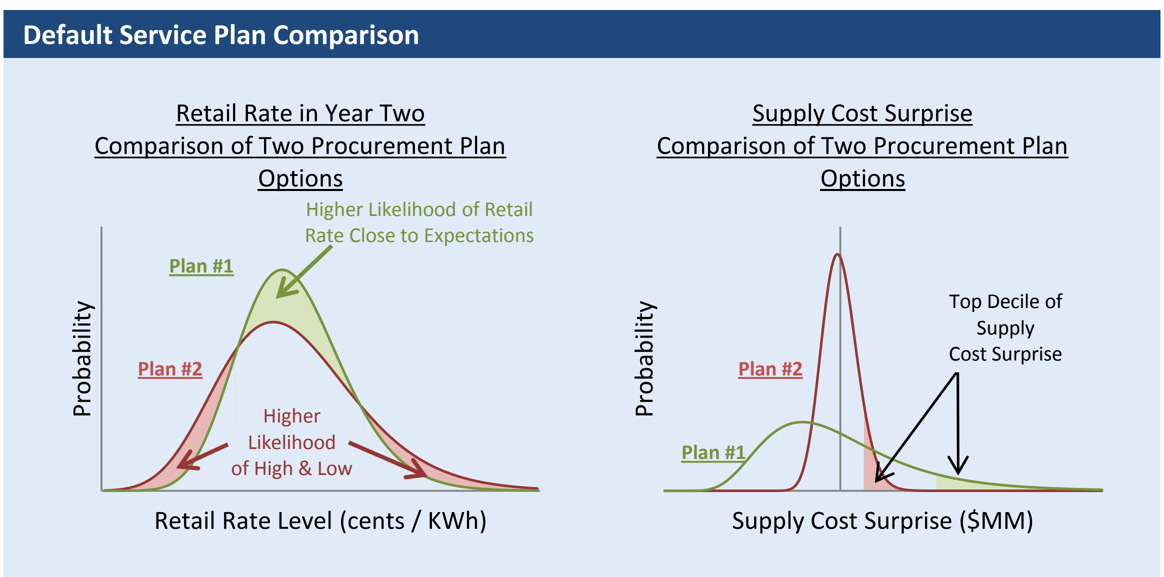

[5] LDCs that are owned by holding companies that also own generating assets generally still procure electric supply from the wholesale market, and the supplier may be an affiliate or an unrelated entity.

[6] They could also hire a third party to manage the portfolio for them.

Insight Developed

Default service procurement plans can differ in many respects, but the major structural differences revolve around product type (e.g. full requirements vs. managed portfolio vs. spot), supply product delivery periods, frequency and timing of solicitations, customer switching rules, cost recovery, etc. Any plan that is approved must pass the somewhat subjective test, 'Do we expect customers will be better off under this plan than under other potential plans?' NorthBridge has assisted LDCs in answering exactly this question by characterizing product pricing, and by simulating outcomes of prices, customer loads, load weighting gross-ups, etc. and then modeling how different default service plan structures would produce different cost/benefit tradeoffs for customers. NorthBridge's consultants have also supported these analyses and presented overall policy arguments as expert witnesses in regulatory proceedings.

Default service plans often involve tradeoffs between competing goals. For example, a plan that guarantees stable costs and rates, ample shopping opportunities, and no working capital concern for the LDC – all at an expected rate level lower than any other alternative – simply isn't possible. Managers and policymakers often should evaluate different default service plans on the basis of quantifiable metrics, such as rate volatility, supply cost surprise, or rate level, and then determine which plan provides a balance of competing objectives that is in the best interests of customers.

Sample Engagement: Project Cost Uncertainty & Disallowance/Under-Recovery/Deferral Risk

The scale of investments in the electric sector is large, and new capital projects often require years of planning and construction. Regulated entities making these investments are in a precarious position. A project that appeared to be economic and competitive when planning/construction began may turn out to be uneconomic by the time it is in operation or at some other time during the life of the asset (*i.e.*, it may have a revenue requirement greater than another alternative that might have been pursued instead). Any possibility that the project will later appear uneconomic, or that the project's overall cost will exceed original expectations, presents a serious financial risk for the utility's investors. Regulators may, unfairly, reason that the utility should have made a different investment and may therefore allow only partial or incomplete recovery of costs. These risks are magnified by the fact that a regulated utility may only have its investment included in ratebase after it is deemed 'used and useful', and that regulators may try to directly or indirectly penalize utilities for any decisions that prove in hindsight to be suboptimal. NorthBridge's capabilities are valuable in helping utilities understand and navigate these risks:

Likelihood and Extent of Disallowance/Under-Recovery Risk

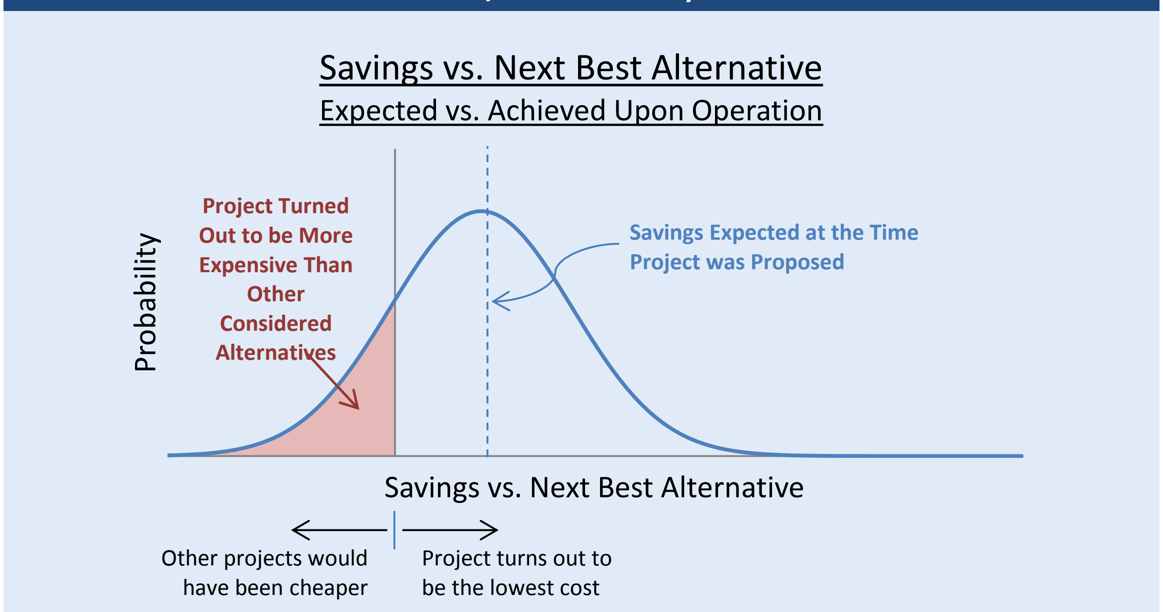

Insight Developed

When project managers recommend an investment in a regulated asset, they do so after having evaluated the project and its alternatives in the context of current forward prices and spot price expectations. However, forward prices fluctuate and may have moved considerably (either higher or lower) by the time the project is completed. If forward prices move in such a way that one of the considered, but rejected, alternative courses of action would have been cheaper than the completed project, full recovery of costs incurred may be jeopardized or deferred.

A key to assessing the financial risk associated with the possibility that an investment may be criticized at a later date is a comprehensive model of how much forward prices, or expectations of future spot prices, may change in the future. Such a model can be used to determine the likelihood that the decision to pursue the project will be contested in the future, as well as the potential magnitude of the financial damage if the decision is contested.

Sample Engagement: Comparative Uncertainty in Investment Returns

When companies evaluate new investments, there are two questions that spring to the forefront: what are the expected future cash flows, and how 'risky' is the investment? An asset's riskiness has direct implications for its discount rate, but precisely identifying the level and type of risk presented by an investment can be challenging when the investment is one-of-a-kind or is otherwise substantially different than the rest of a company's projects (*i.e.*, the company's after-tax WACC may not be appropriate). NorthBridge's commodity price simulation approach allows company executives to visualize how uncertain an investment's total returns are, and how effective certain approaches are in mitigating that risk:

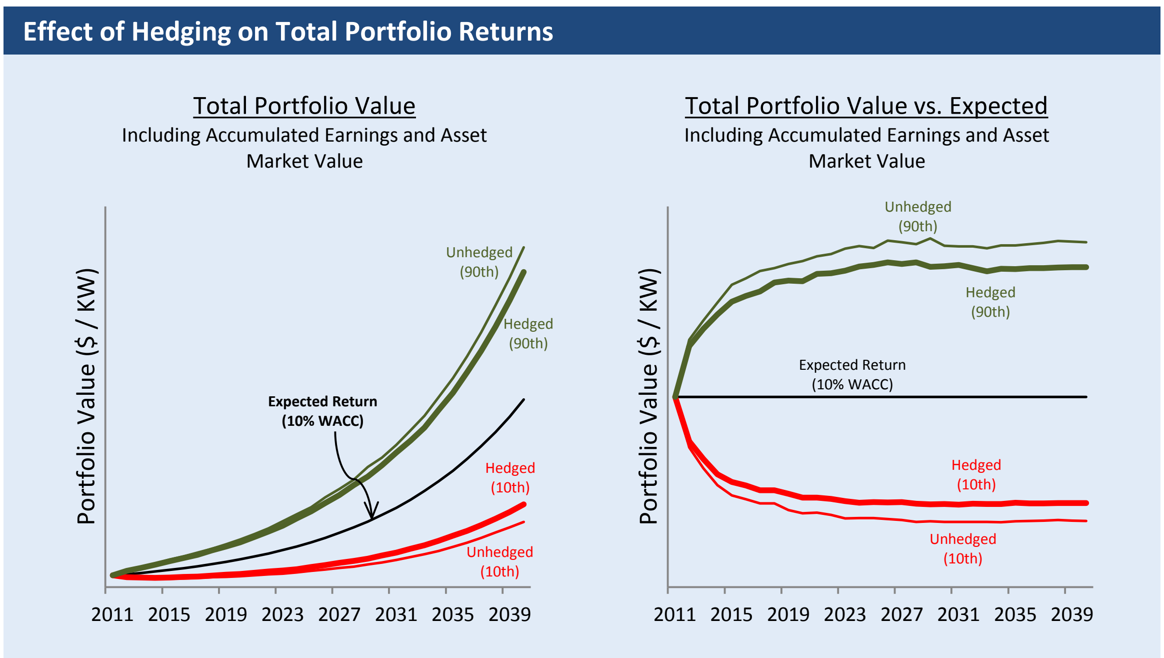

Insight Developed

The 'riskiness' of an asset in the energy sector is heavily dependent upon its exposure to various commodity markets. When a company makes an investment in an asset today, it does so with the expectation that the asset will return cash flows in the future. The total return accrued to an investor at some point in the future is the sum of two pieces: cumulative cash flows returned to the investor (reflecting past performance), and the change in the market value of the asset (reflecting future expectations).[7] Energy assets are considered to be 'risky' because the total return produced by the investment is uncertain.

NorthBridge has developed models to illustrate how uncertainty in spot and forward prices affect both cash flows and conditional future asset valuations, which jointly reveal the riskiness of an

[7] This is analogous to the total return on a bond as observed at any time prior to maturity.

asset's total return. This type of analysis is also helpful when identifying how effective a hedging program will be in reducing the riskiness of an asset and reducing its required return. For example, a rolling three-year hedging program may do an excellent job at securing near-term revenues, but may be largely ineffective at protecting the owner from changes in asset value due to fluctuations in the long-term forward curve.

Appendix: Technical Foundations

Understanding Uncertainty

Quantifying commodity price uncertainty is one of the most challenging problems of computational finance. Equity shares and debt instruments trade in liquid markets where there is generally a reasonable (although not perfect) assumption that the best predictor of tomorrow's price for a given asset is today's price for that asset, and where a price movement of +/- 20% over a week or even a month is considered an extraordinary event. Commodities, on the other hand, are subject to price seasonality due to supply/demand/storage variations over the course of the year (and even over the course of the day as is the case with electric markets), mean reversion, illiquid markets and price scarcity, and often price volatility at levels that are orders of magnitude higher than those for equities.

Traditional quantitative models of price volatility either fail horribly at describing the price dynamics of commodity markets or are so specialized to answer a single question that they cannot easily be used to answer more general corporate strategy questions.[8] For example, naïve attempts to adapt stock option pricing models to commodities fail because two basic characteristics of stocks, first that today's price is the best predictor of tomorrow's price (*i.e.*, no mean reversion), and second that an asset can be purchased today and stored without cost until delivery (spot-forward equivalence), simply do not apply to commodities.

Practitioners bypassed these issues by focusing exclusively on forward contracts for commodities. This emphasis solved the issue of tractability, but failed to address questions that directly involve spot price outcomes. Modern financial models attempted to tie together the joint nature of forward and spot prices in the commodity world by introducing mean reversion. This represents a major advance in replicating a key characteristic of commodity price dynamics, but still leave much to be desired.

Ultimately there are two questions that must be answered before any meaningful work in commodity price simulation can be done:

- What representation of price evolution is sufficiently flexible to mimic the complex behavior of commodity prices as observed in the real world, and
- How does one 'fit' (*i.e.*, calibrate) the model to the commodity in question?

The answer to the first question can be thought of as the skeleton, while the second is akin to the skin and muscle.

NorthBridge has borrowed from, and improved upon, the latest advances in computational finance to answer both of these questions. The remainder of this section is divided into three topic areas:

- Characteristics of commodity price dynamics

[8] GARCH models may fall into this category.

- The structure of a general model to represent price evolution
- Model calibration

Characteristics of Commodity Price Dynamics

Unlike debt and equity, a single commodity can best be described as having a spectrum of prices: a single spot price for immediate delivery and a separate forward price for each future delivery period. The cost of storing commodities makes the distinction between spot and forward price critical. Not all commodities have liquid or visible forward markets, but the lack of a visible market does not imply the lack of a forward price.

Spot and Forward Prices

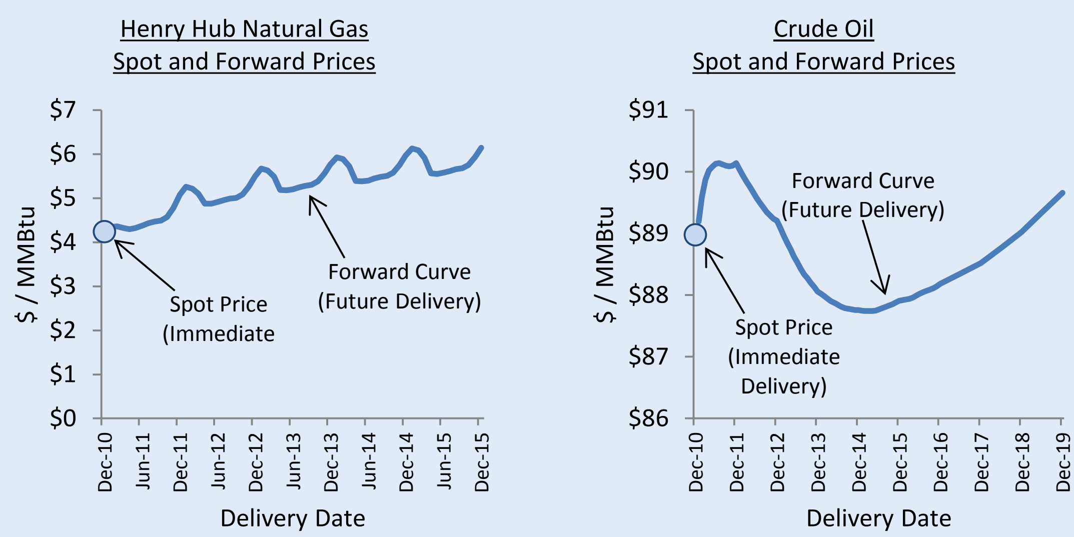

Some characteristics of both spot and forward prices are obvious even to the casual observer. Other characteristics require a more detailed analysis to uncover, but are equally important in a risk assessment.

The most obvious price dynamic characteristics include:

- Volatility (*e.g.*, standard deviation of returns), often generalized as 'uncertainty'
- Mean reversion (*i.e.*, the tendency of price shocks to fade over time), and
- Inter-commodity correlations

Less obvious, but just as important, price dynamic characteristics include:

- Skewness and kurtosis ('fat tails') & higher-order moments of returns
- Range of prices over time
- Volatility of volatility (vol-vol)
- Mean reversion of volatility
- Persistence of volatility (volatility clustering)
- Price/volatility correlations between commodities
- Volatility term structure in forward price returns

- ✧ Skewness and kurtosis of forward price returns

While few legitimate business questions require replicating all of these characteristics perfectly, a sound risk analysis should attempt to avoid producing simulated results that are wildly inconsistent with any of these observed characteristics.

The Structure of a General Model to Represent Price Evolution

The underlying model NorthBridge utilizes to simulate commodity price evolution is based on, at its core, a Monte Carlo random walk. This, in essence, means that we can simulate the price of an asset as varying from one period to the next on the basis of random 'noise' scaled by some amount to represent the magnitude of volatility. The NorthBridge approach, though far more sophisticated, is still based on this principal.

Expressed mathematically, the simple random walk has the form:

$$dP = \sigma_p \cdot dW + drift\ adjustment$$

$dP = Change\ in\ price$
$\sigma_p = Marginal\ volatility\ of\ price$
$dW = Randomly\ distributed\ variable$

The first major modification to the simple Monte Carlo random walk is the incorporation of mean reversion. This has two ramifications. First, incorporating mean reversion allows simulated prices to exhibit the negative autocorrelation observed in real-world spot prices. Second, it provides a natural extension and link between spot and forward prices. Mean reversion allows us to describe a forward price as a spot price that has been subjected to decay back toward a long-run expectation. This adaptation is more of a must-have than a refinement; without it there is simply no way to establish a causal relationship between spot and forward prices.

This is expressed mathematically as:

$$dP = (P - \bar{P}) \cdot h_p \cdot dt + \sigma_p \cdot dW + drift\ adjustment$$

$dP = Change\ in\ price$
$P = Price\ in\ prior\ period$
$\bar{P} = Long\text{-}term\ average\ price$
$h_p = Rate\ of\ mean\ reversion\ of\ price$
$dt = Time\ elapsed\ since\ prior\ period$
$\sigma_p = Marginal\ volatility\ of\ price$
$dW = Randomly\ distributed\ variable$

Spot and Forward Prices

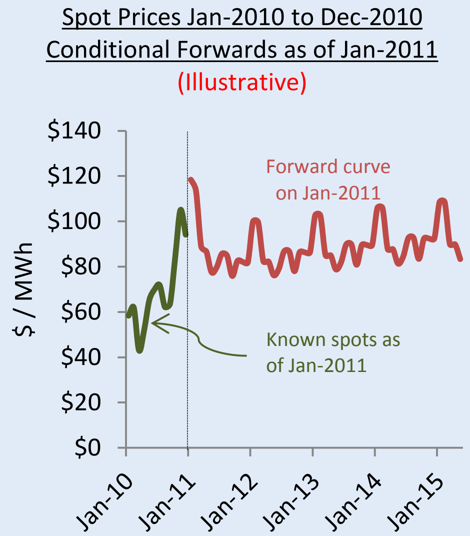

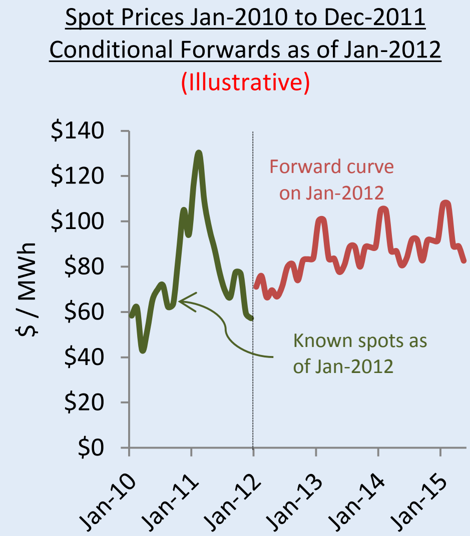

The second major modification is to incorporate what is known as 'stochastic volatility'. One of the most prominent features of all financial markets[9] is that some periods exhibit 'quiet' and stable price environments while other periods are demonstrably volatile and chaotic. Furthermore, this variation between low and high volatility periods appears to be driven by something more fundamental than sampling noise; there are distinct periods or 'clusters' of distinct volatility regimes that can last for years. For example, daily returns on the S&P500 clearly demonstrate both variable and clustered volatility:

[9] Early option pricing models failed to identify or account for this characteristic, often with tragic results.

Volatility Changes over Time

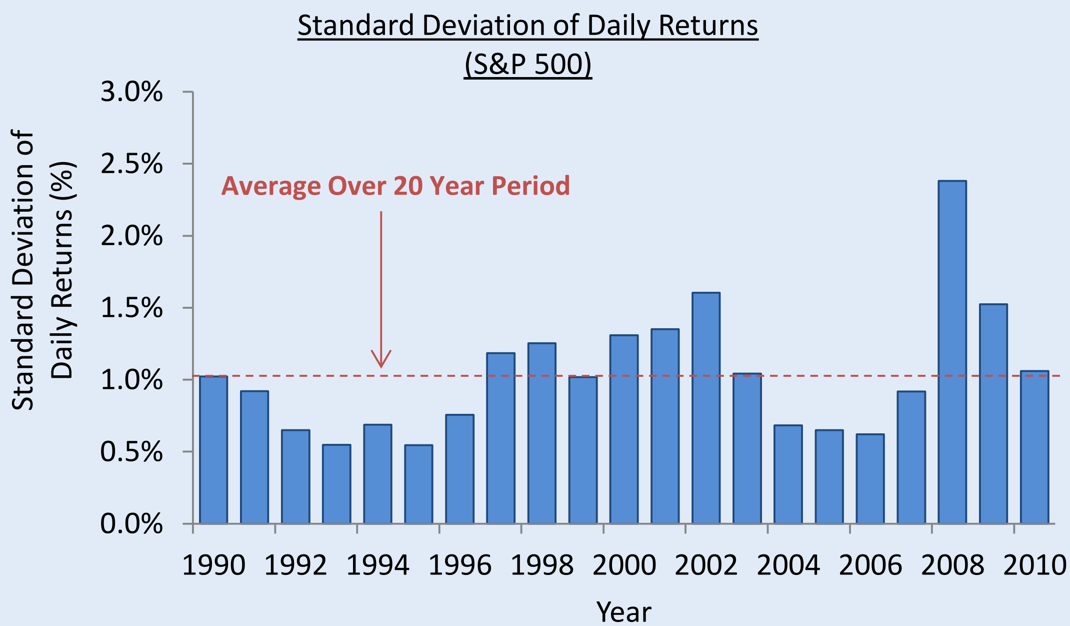

Early models (*e.g.*, Black-Scholes) of how prices evolve assumed constant volatility throughout time. This assumption had the benefit of making the math 'easy', but it also had profound implications on estimates of the likelihood of extreme price movements and the pricing of out-of-the-money options. For example:

Conventional Model of Price Movements

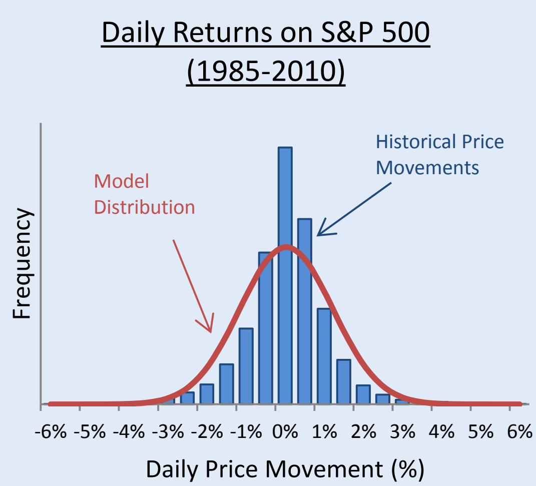

Percentile Outcomes

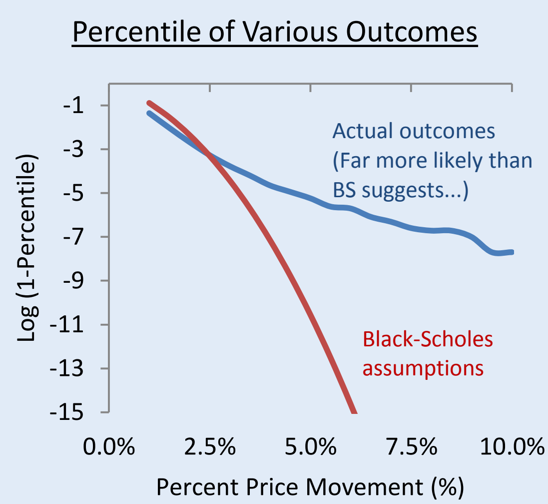

We observe in virtually all markets that extreme price movements (and small price movements) are far more likely than would be implied by the constant volatility assumption. We also observe that price movements of moderate magnitude, in between the small and extreme variety, are less likely in the real world than suggested by the simple model.

Stochastic volatility, the property that volatility is itself variable, is one way to address both problems: volatility clustering and fat tails. The simpler model described above simulates prices as being both variable and as having some degree of reversion to a long run-average. We can do the same with price volatility, in essence modeling volatility as a fluctuating variable that rests underneath price movements and influences their magnitude. We can express this improvement mathematically as:

$$dP = (P - \bar{P}) \cdot h_p \cdot dt + \sigma_p \cdot V \cdot dW + drift\ adjustment$$

$$dV = (V - \bar{V}) \cdot h_v \cdot dt + \sigma_v \cdot dZ + drift\ adjustment$$

$$r(dW, dZ) = \beta$$

$dP = Change\ in\ price$
$P = Price\ in\ prior\ period$
$\bar{P} = Long\text{-}term\ average\ price$
$h_p = Rate\ of\ mean\ reversion\ of\ price$
$dt = Time\ elapsed\ since\ prior\ period$
$\sigma_p = Basecase\ marginal\ volatility\ of\ price$
$dW = Randomly\ distributed\ variable$
$dV = Change\ in\ volatility$
$V = Volatility\ in\ prior\ period$
$\bar{V} = Long\text{-}term\ average\ volatility$
$h_v = Rate\ of\ mean\ reversion\ of\ volatility$
$\sigma_v = Basecase\ marginal\ volatility\ of\ volatility$
$dZ = Randomly\ distributed\ variable$
$\beta = Correlation\ between\ dW\ and\ dZ$

In practice, this produces simulated price outcomes that do a much better job at replicating the types of price movement characteristics that are important. For example, it allows us to do a much more realistic job of estimating the likelihood of extreme price movements:

Improved Model of Price Movements

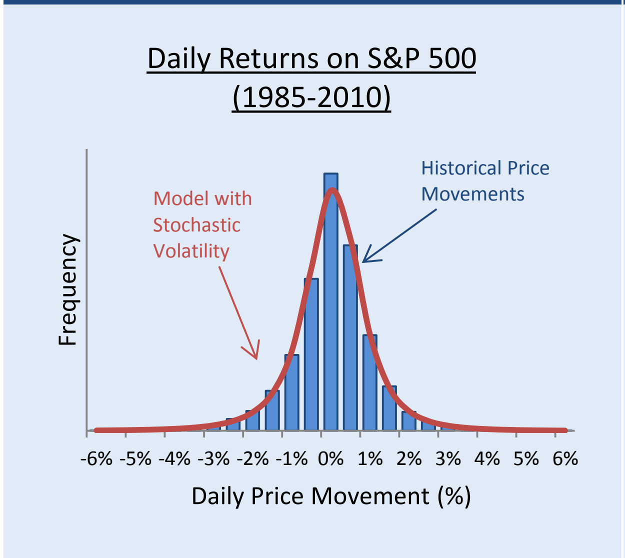

Percentile Outcomes

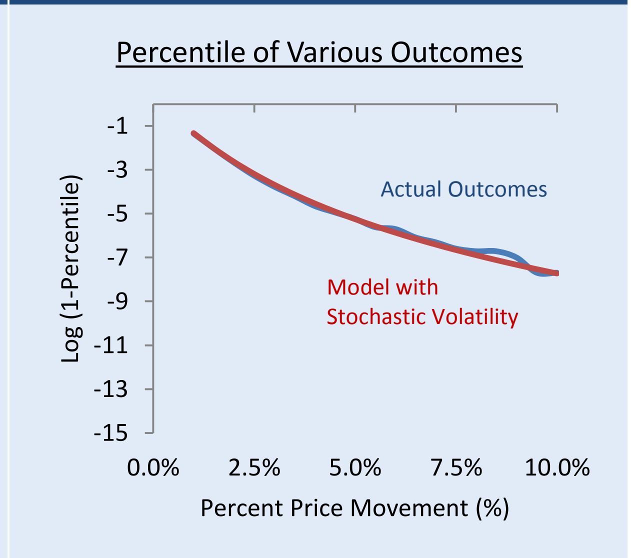

The strength of this approach is not only that it allows us to better reflect the range and likelihood of outcomes at a future point in time, but additionally that the underlying model allows us to simulate and observe all aspects of a commodity's price evolution and its relationship with other commodities. Many business strategy questions require characterizing not only the price of a commodity at some point in the future, but also how the price of that commodity evolves during the intervening time:

Potential Forward Price Paths and Spot Price Outcomes

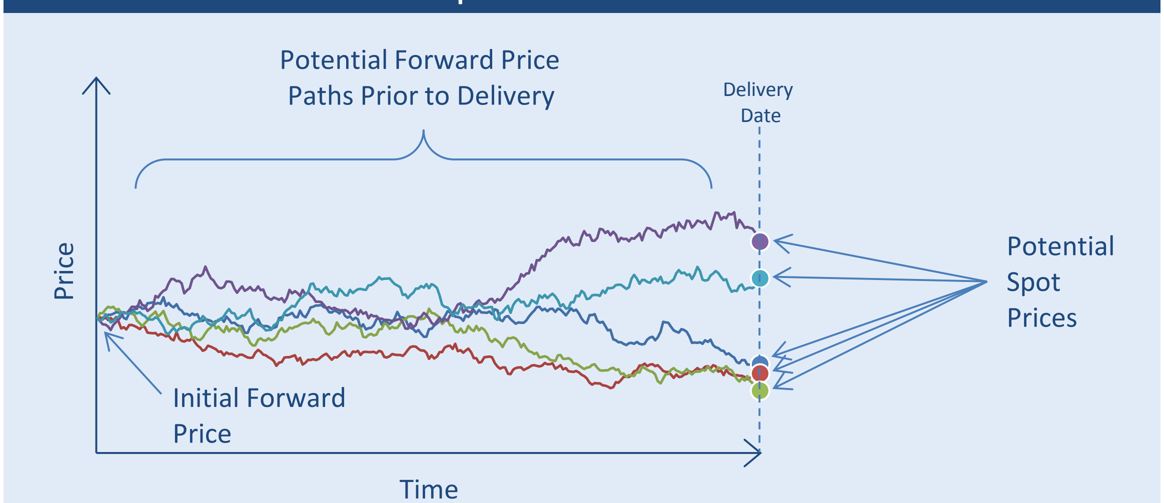

The ability to simulate a price and 'watch' it evolve over time makes this approach exceptionally flexible and adaptable to virtually any strategy decision that requires insight into how commodity prices could vary.

If we were to step back and ask what sort of fundamental drivers affect commodity prices, we might conclude that there are some drivers that create short-term disturbances in the spot price which would not affect long term expectations (*i.e.*, they are quickly mean reverting). An example of such a driver might be a malfunctioning transformer affecting a nodal electric price, or a cold-snap driving up the cost of natural gas. These short-term influences can lead to dramatic changes in the spot price, but may not lead to any observable change in the long-term forward curve.

On the other hand, there may be drivers that have little observable impact on spot prices, but which have a significant impact on long-term expectations. For example, the impact of trends in macroeconomic activity would be nearly unobservable in daily changes in the spot natural gas price,[10] but would probably be the best explanation for any change in forward contracts with long-dated maturities.

Fortunately, we can further adapt our modeling approach to account for these 'multiple personalities'. Rather than modeling a price as a single process with some degree of volatility and mean reversion, we can approximate the multiple fundamental drivers influencing a commodity by simulating its price as a composite of multiple sub-processes. For example, a simulation of natural gas prices may include a short-term process (with high mean reversion) meant to reflect forcing from factors such as weather or pipeline constraints. It may also include a medium-term process (with moderate mean reversion) meant to reflect forcing from intra-year storage conditions. Finally, we may also include a long-term factor meant to approximate non-mean reverting price movements associated with macroeconomic conditions with permanent price implications such as technological innovations or productivity gains.

Below is a graphical illustration showing how we can simulate three separate price processes, and then combine them into a composite price process that contains elements of three different types of price behavior:

[10] This is not to say that long-term macroeconomic influences do not impact daily spot prices; they do. However, any influence from these small, but non-mean reverting, effects will likely be dwarfed by short-term fluctuations, making the long-term component of volatility difficult to discern.

Market Prices as Composites of Multiple Concurrent Processes

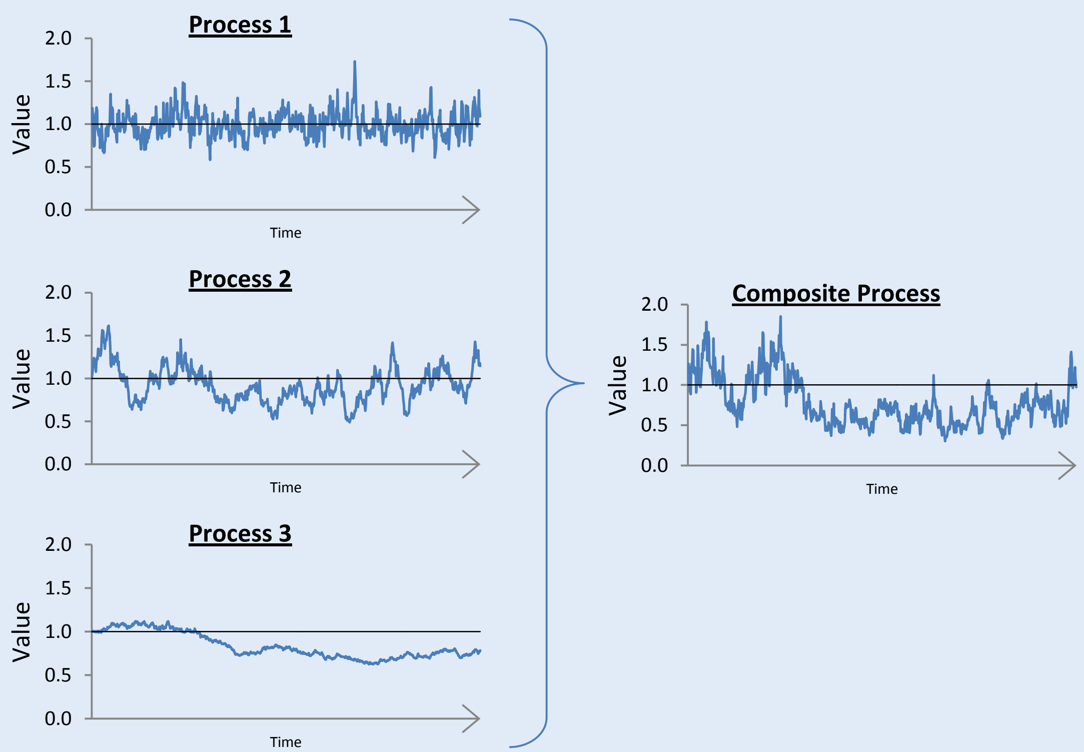

The discussion above describes three processes. In reality, commodity prices are influenced by myriad different drivers, each with unique implications for price volatility and mean reversion. The adaptation of the commodity price model to include multiple concurrent processes better replicates real-world price behavior, but is not a suggestion that we are capable of explicitly identifying and replicating the fundamental processes (or structure) of price movements in the real world.

Some commodity price processes may be well approximated by a single process, meaning that our simulated price paths match all of the important characteristics of the 'real thing' to a sufficient degree of confidence. Other commodities may exhibit more complex price dynamics, driving us to approximate the price process as a composite of two, three, or more independent processes. NorthBridge has developed a simulation/calibration framework capable of incorporating numerous concurrent processes. In general, the practitioner would be best served to start with a model containing a small number of processes, and then include more processes only when the simpler model cannot be calibrated to a sufficient degree of confidence.

Model Calibration

In the prior section, we described the structure of a general model that is sufficiently flexible to replicate the wide array of price dynamics displayed by commodities. As a purely mathematical construct, however, such a model is not useful. The model specification can be thought of as the raw components of a home construction project: lumber, nails, tiles, shingles, drywall, paint, glass, and so on. A fine home contains all these components, but a pile of these components does not constitute a fine home. So the question arises, how does one 'fit' the model so that when it is used to simulate prices, the price paths it produces exhibit the same characteristics as those that define the commodity we are simulating?

In some cases, such a calibration is simple. For example, in naïve option pricing models the only input parameter, volatility, is directly measurable from historical data.[11] All one has to do is calculate the standard deviation of historical returns in order to estimate the parameter to use when simulating price paths.

In more complex model structures, such direct measurement is not feasible; we need to revisit what it means to have found the 'right' model parameters. In the simple example above, the standard deviation of historical returns is not the 'right' parameter because of how it is measured; instead, it is the 'right' parameter because, when it is used to generate price paths, those price paths exhibit the appropriate level of volatility. This perspective can be described in a slightly different way: if we think of the observed historical prices as a single 'scenario' or draw from a set of all possible scenarios generated by the underlying system, then the 'right' model parameters are those that indicate that the observed price path was a more likely outcome that any other set of parameters would suggest. More formally, this means the 'right' model parameters display 'maximum likelihood'.

For example, imagine that we are trying to build a model that replicates the volatility of a single equity price. Perhaps we measure (directly) the historical volatility of the stock price as being 20%. Next, imagine we choose a volatility parameter of 50% for our model and then ask if this satisfies the requirement of displaying 'maximum likelihood'. If we simulate a stock price using the 50% parameter, some price paths generated would exhibit volatility over 50% and some under 50%; the scenarios on average would exhibit volatility of 50%. We would then ask what proportion of the simulated price paths exhibited volatility of around 20%. The answer would be perhaps some, but very few. The model parameter of 50% suggests that the observed historical characteristic would have been very unlikely if the 'true' parameter were 50%. If instead we had used a parameter value of 30%, many more of the generated price paths would have exhibited the level of volatility measured historically. In this case, the model parameter that displays 'maximum likelihood' would actually be 20%. Again, the 'right' answer is not right because we were able to measure it directly, but because it makes the observed outcome (historical volatility) more likely than any other model parameter or set of model parameters.

---

[11] In practice, some modifications are made to the historical observations in order to transform them into a forecast of future variance.

Obviously, for the simplest models, this logic seems excessive when the 'right' values of the model parameters can be determined by direct measurement. However, in more complex models, there is no methodology that allows us to measure the 'right' value directly. Instead, we must resort to finding the 'right' set of model parameters using the maximum likelihood technique.

The algorithm developed by NorthBridge to derive the maximum likelihood parameters involves computationally intensive simulation batches nested within a non-linear optimization loop:

Model Calibration Process

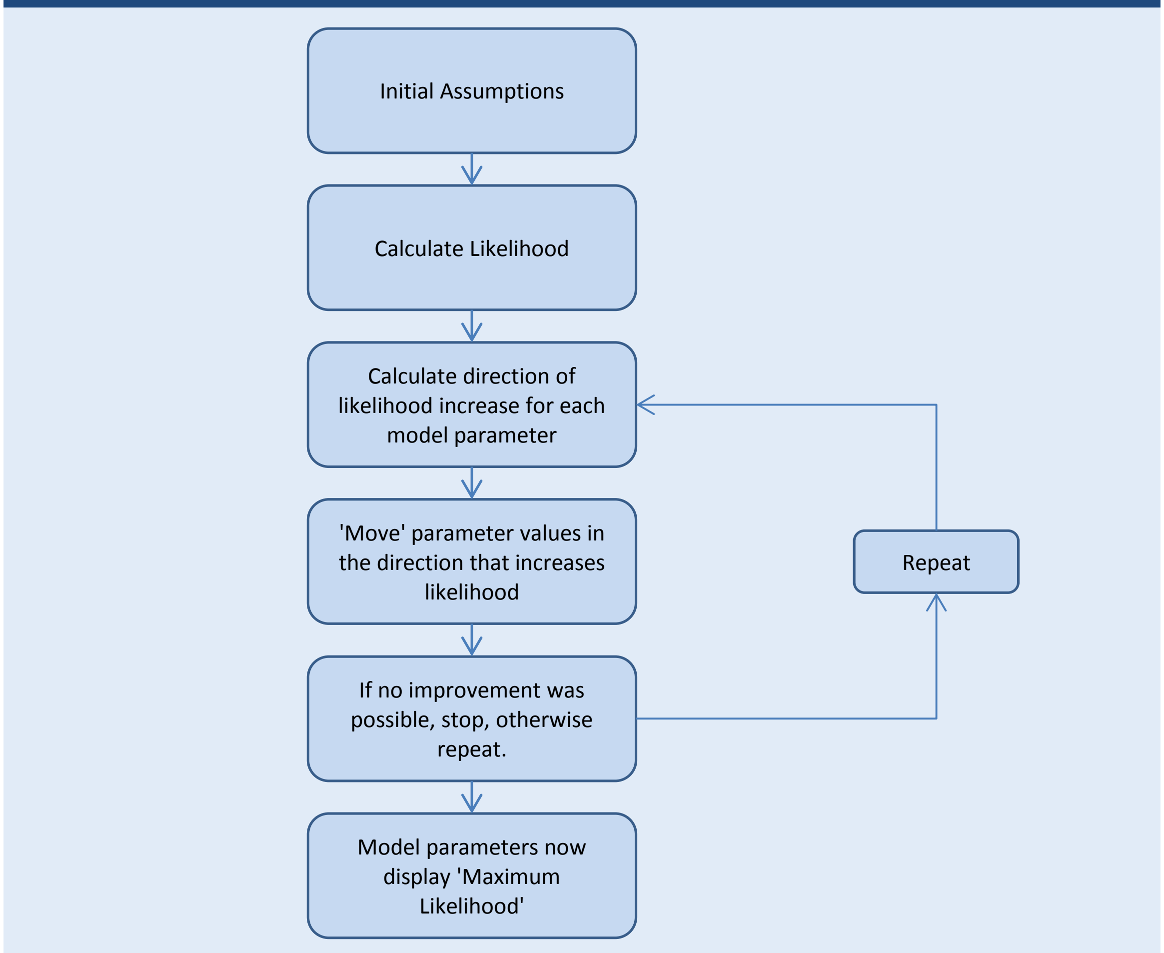

The implementation details of this algorithm are complicated, but the general concept is simple. Start with some naïve assumption about the values of model parameters. Next, calculate the likelihood value of the observed outcomes using those parameters. If we can increase the aggregate likelihood of the outcomes by increasing or decreasing the value of some model parameters, do so. Continue 'tweaking' the model parameters until no further improvement is possible. Of course, this approach can be supplemented with knowledge of additional uncertainties about the future that may not be captured in the observed historical data.

The primary drawback to this approach is that it can be quite demanding computationally when attempting to calibrate the model to large data sets of spot and forward prices, or when attempting to calibrate to datasets of dozens of different commodities (due to the N^2 nature of correlation tables). Fortunately, the process of generating scenarios as part of the likelihood calculation is easily implemented in a parallel algorithm. NorthBridge has developed and refined the calibration tool to work quickly and efficiently, fully utilizing the multiple-core configuration of modern computers.

Since the model described above is calibrated to measured historical price movement characteristics, the selection and validation of the source historical dataset is critical. The model calibration process can be conducted even if only spot data is available,[12] but is more robust if forward price history can be included. The reason for this is two-fold. First, when we talk about commodity price volatility, what we generally mean is price uncertainty relative to expectations. When we observe a spot price movement in isolation, we must make some assumption about how much of the movement was expected versus unexpected. Second, some price dynamics, like mean reversion, are more clearly observable in forward price movements (*e.g.*, by observing the shape of the volatility term structure).

Legitimate forward price history is a valuable input in the calibration process, but this is not to say that all forward price history should always be included. There are several reasons to exclude observations. For example:

- Forward markets are often not liquid and prices quoted may not be current, or may not represent actual trades or be based on actual bids and offers.
- If the market has undergone some structural change, forward-looking price dynamics may be expected to differ from historical dynamics. In this case, one should exclude the price history from the period that is no longer believed to be indicative of future behavior.

[12] Prompt-period forward prices may be substituted for spot prices if necessary.

Contact Information

Contact: David Coleman
Principal
The NorthBridge Group
30 Monument Square
Concord, MA 01742

Phone: (781) 266-2632
Email: dcc@nbgroup.com
Website: www.nbgroup.com