Managing volumetric risk is exceedingly challenging, yet getting it wrong can damage a company's reputation and make it look as though it is not in control of its operations. *Garth Renne* and *Ken Truesdell* of Brookfield Power suggest some solutions

The one that gets away

Garth Renne and Ken Truesdell of Brookfield Power

★ Energy companies today have access to a suite of well-established tools and techniques to manage price risk, both in the context of proprietary trading and for asset management. The remaining debates in that area mostly concern what might be characterised as narrow technical details: the appropriate stochastic processes to model particular market prices, the best way to deal with uncertain and unstable correlations between products, and so on.

In contrast, alternatives for assessing and managing volume (or volumetric) risk – the risk to financial results stemming from unpredictable fluctuations in production or sales obligations – remain less well developed. This presents risk managers with a challenge when designing measures, limit frameworks, and reports, but also offers opportunities for creative work and unique, company-specific solutions. Brookfield Power[1], a merchant generator, has grappled over the last several years with this issue as its portfolio has grown to approximately 3,800 megawatts of primarily hydroelectric assets.

Apart from its relative infancy, there's a second, and perhaps more fundamental reason why volume risk management is of interest: executed poorly, it can be a greater threat to a company's reputation than its price risk counterpart. Many companies have been criticised (often unfairly) for hedging programs that were perceived as ill-timed from a price standpoint. However, companies that misjudge their volume exposure – possibly leading to unexpected losses on 'hedges' not backed by production – can face far harsher investor reactions. Questions can arise as to whether the company is properly maintaining its assets, or worse, that it doesn't have a firm understanding of its operations.

This article presents some key considerations in assessing volume risk, along with some suggestions for developing limit structures and managing uncertain production volumes. Although the perspective is one of a merchant hydroelectric generator, most of the observations also apply to other common cases of volume risk in the energy industry.

Volume risk: basic considerations

A very simple framework can illustrate some of the basic challenges of volume risk. Figure 1 shows a revenue profile from a single hypothetical – but realistic – plant on one of Brookfield Power's river systems.[2] Clearly, the combination of volume and price uncertainty is a potent mix, with the multiplicative effect generating a huge potential range of outcomes. In practice, this can be compounded by the fact that in many regions, precipitation is concentrated in particular seasons – or even in a few discrete storm systems – meaning that the range of volume uncertainty can behave more erratically than its price counterpart.

The width of earnings-at-risk (EAR) bands for hydro-generating units often surprises risk oversight committees, but a few years of experience usually demonstrates how dramatic the combined effects of volume and price uncertainty can be. Diversification across units and river systems is an obvious source of volume risk reductions, but implementing such a physical portfolio strategy can take many years.

Further complicating the picture is the

1 Brookfield Power is 100% owned by Brookfield Asset Management, which has over US$70 billion of property, power and infrastructure assets under management.
2 The model is based on 25 years of actual daily precipitation and inflow data and a hypothetical run of river plant whose performance is related to inflows using historical relationships. Annual prices are distributed lognormally with a typical daily price shape. No systematic correlation between hydro output and prices is assumed in this simple model.
3 This model assumes a normally distributed physical price at the plantgate, which is correlated with a normally distributed financial price used for hedging. A separate low-probability price spike affects both the physical and financial prices equally.

presence of multiple sources of correlation between volumes and prices. If a unit is located in a heavily hydro-based market like the US Pacific Northwest, low hydrological conditions are associated with higher average prices, potentially offsetting a portion of volume risk. An analogy here is that faced by many farmers: in poor crop years, prices for agricultural commodities are often higher than in years of bumper crops. One consequence of this is that a farmer who sells forward his long-term, expected level of production, may be over-hedging.

While it may seem that these correlation effects are quite weak, given that most markets in North America are not heavily reliant on hydroelectric power, several factors can magnify their influence.

First, most hydro systems have at least some flexibility to shape output, moving energy preferentially into high-priced periods. So, total expected revenue does not fall one-for-one with decreasing energy availability on a percentage basis. This effect itself can be dynamic: often, the flexibility of hydro systems is maximised in moderate water years, with extremely high and low water conditions tending to reduce optionality in dispatch decisions. Secondly, even when hydro plants do not loom large in a region's energy markets, they may supply a significant fraction of ancillary services, the supply of which may also vary in response to water conditions.

Finally – although this is usually more of a consideration with thermal plants – in some markets the loss of a few units can significantly increase the probability of price spikes. It only takes a few instances of a unit being off-line under price spike conditions to create a significant correlation between the unit's performance and its realised sales price.

While volume and price risk are inextricably linked, and combine to form an overall revenue distribution, it can be helpful to view the two sources of risk along separate dimensions. In fact, the two-dimensional aspect of the problem is one of the main reasons why it is so challenging: condensing the situation down to a single risk measure and/or using a single hedging instrument for risk management can inadvertently increase certain risks.

The potential pitfalls of ignoring the two-dimensional nature of volume/price risk are illustrated in Figures 2 to 5, produced from another simple simulation.[3] The top scatter plots represent an unhedged case (physical sales only). Obviously, there is a high positive correlation between revenue from physical sales and both price and volume; the more isolated points in each graph reflect the incidence of price spikes. In this simple setup, the standard formula for a variance-minimising hedge can be applied. The bottom two graphs show what happens if this hedge position is taken. Because the model assumes a very high correlation

Brookfield Power's High Falls station, located on the Lièvre river in Quebec, Canada. Even with new, emerging tools and techniques, the vagaries of weather make hydro assets a challenge for corporate risk management groups

Price level (columns) / Volume level (rows)

$ million	–1 std dev	Mean	+1 std dev
–1 std dev	0.92	1.82	2.68
Mean	1.17	2.30	3.39
+1 std dev	1.41	2.77	4.09

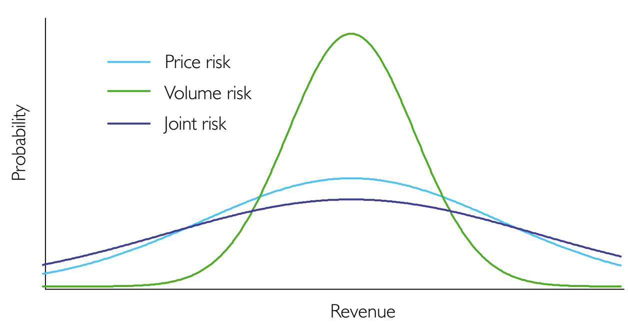

F1. Price and volume risk of a hypothetical hydroelectric station
Source: Brookfield Power

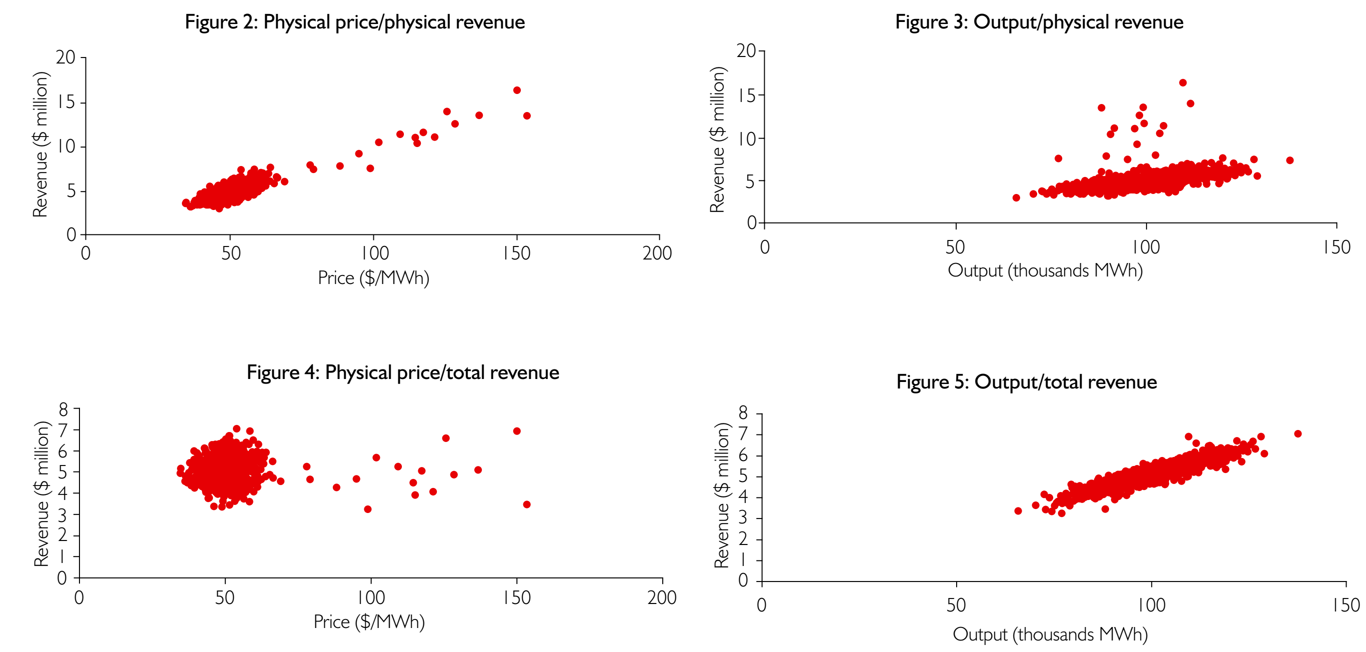

F2 to F5. Results from a simple physical plant/hedging simulation

Selling output forward reduces the standard deviation of total revenue, but is it a good idea? *Source: Brookfield Power*

between the price realised on physical production and the financial price used to settle the hedge, the hedge is very effective at eliminating the overall correlation between prices and total revenue (as seen in the bottom left-hand graph), as well as removing the effects on revenue of price spikes (the bottom right-hand graph). Output and total revenue remain highly correlated – in fact, more highly correlated now that price risk has been substantially controlled.

One might conclude from Figures 2 to 5 that such a hedging program is an excellent idea – certainly, the standard deviation of the revenue distribution is substantially reduced. However, a closer look at specific iterations from the simulation reveals a somewhat different story. Figures 6 and 7 were created by sorting the simulation results from lowest total revenue to highest, and then graphing the variables of interest. Figure 6 simply confirms that all else being equal, higher total revenue figures are associated with higher outputs. Note from the circle in figure 7, though, that one of the worst iterations from the hedged case was a situation in which prices spiked up, and since this iteration was so poor from a revenue standpoint, production volumes were low by implication. The outcome of massive losses on what turned out to be an uncovered forward sale might lead to some discomfort at the next quarterly call with investors. Both the level of prices and the incidence of price spikes show no particular relationship to the rank of iterations by total revenue in the hedged case – meaning that the price hedge is doing its job on that front.

A practical lesson can be drawn from this somewhat theoretical exercise: let upper management know about specific possible outcomes, rather than just reporting summary statistics such as standard deviation. This can be crucial if the risk problem contains a lot of 'moving parts' – volume, price, dispatch and contract flexibility being examples. In addition, risk managers should pay attention to near misses, since in our experience, events such as the worst-case scenario in figure 7 are often foreshadowed by events in which severe losses are just barely avoided.

Volume risk and limit frameworks

Volumetric risk poses a challenge to the design of a risk control framework. Ideally, a control structure limits risk to conform to a company's tolerances, reflects the overall corporate strategy (for example, always being net long to give shareholders exposure to the underlying commodity), and doesn't give traders incentives that differ from corporate preferences. One natural solution would be simply to turn over to a trading book the entire asset position, including any volumetric risk and all hedging activity, and gauge performance by a financially meaningful measure such as net operating income.

In a trading environment, this approach is often unworkable. Uncertain volumes make establishing appropriate transfer price benchmarks contentious – and they also make value-at-risk (VAR) figures difficult to interpret. Meanwhile, EAR models are rarely employable as day-to-day tools for risk control, due both to a lack of standardisation as well as the fact that much of a company's asset-related EAR comes from factors outside the direct control of traders. Finally, because most volume risk is beyond the control of trading groups, having a substantial fraction of their incentives tied to this factor can cloud the relationship between pay and performance.

Brookfield Power has adopted a relatively straightforward approach to this problem: potential output is divided up into tranches, with the trading and marketing group allowed to use certain tools to manage the risk on each segment. A conservative volume is transferred to the traders at market prices and any departure from the default strategy of selling forward all (and only) this volume is subject to a variety of risk limits, including VAR. Above a certain level, trading and marketing is restricted from selling forward outright, but can purchase insurance-type products such as options.

This framework has the nice feature that risk is measured relative to a default position that the company is comfortable with as a long-term strategy (that is, hedging forward the substantial fraction of generation that is highly likely to be available, then selling other volumes on a very short-term basis). Traders do not bear any volume risk on the conservative volumes, but are exposed to the risk that any additional volumes do not show up. This approach of asset-backed optimisation around a core, conservative strategy is explicitly documented in the company's risk policy.

Many companies take a similar approach to risk limits around real assets, translating asset-based positions into delta, gamma, and/or other 'greeks', then placing bounds on these measures in combination with limits on VAR. Of course, volume risk is a more central feature in the case of hydro than for thermal power plants or gas transportation. For hydro, drawing a hard line in the continuous volume distribution at the 'conservative' level is an imperfect solution. Because VAR is measured relative to this guaranteed (from the traders' perspective) volume, while EAR-type measures are based on the complete distribution of output, there can be substantial discrepancies between the two measures. So, while VAR is not Brookfield Power's volume risk measure, and doesn't reflect overall corporate risk, it is a useful, standardised way to ensure that the company's position doesn't stray too far from its core strategy.

Regardless of how well a company's limit framework is designed, volume risk requires that the middle office supplement it with a bottom-up understanding of the underlying physical assets. Key contributors to this are expertise, communication and back-testing. At Brookfield Power, a subgroup of the middle

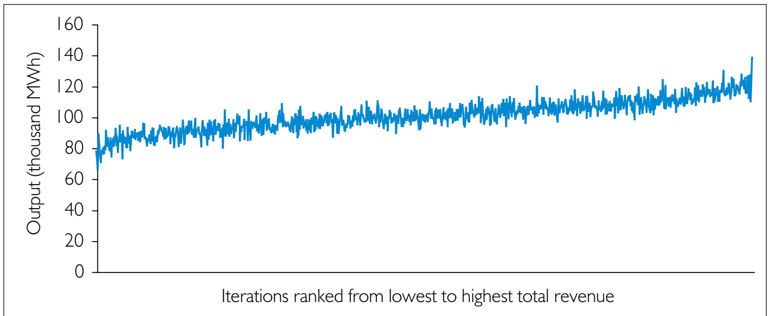

F6. Output by ranked simulation With price risk controlled, output is strongly correlated with total revenue... *Source: Brookfield Power*

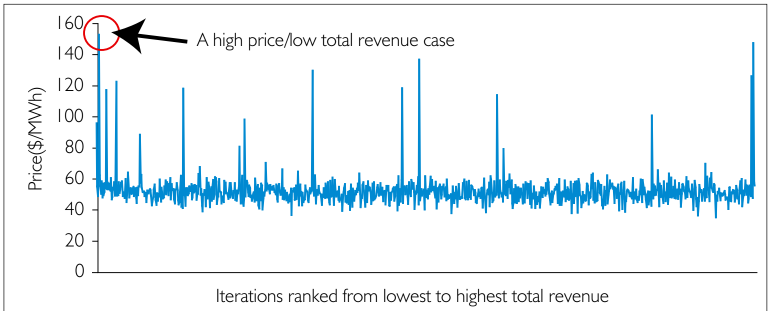

F7. Physical price by ranked simulation ...but the door is opened for some undesirable possibilities (circled) *Source: Brookfield Power*

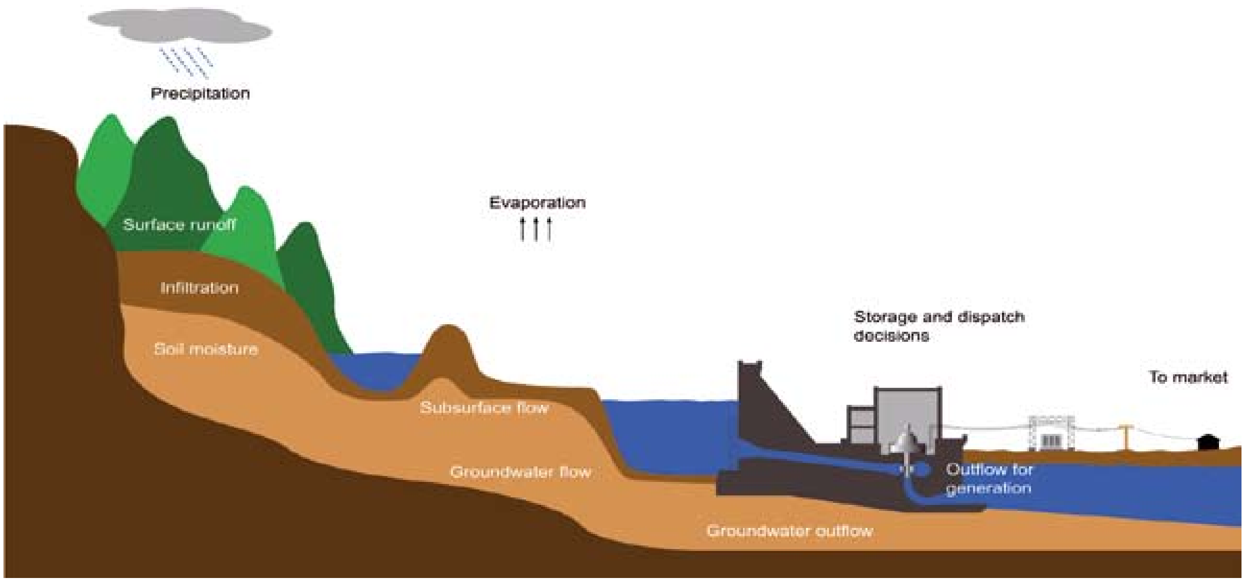

F8. The long chain from precipitation to market
Multiple steps between precipitation and revenue weaken correlations *Source: Brookfield Power and Purdue University, College of Engineering*

office specialises in volume planning to independently determine the volumes of each tranche. Most of these staff have engineering or hydrological backgrounds.

In addition, an understanding of volume risks is created through site visits, along with daily and weekly communications with operations. The operational performance of units and water availability relative to expectations is monitored.

Finally, dispatch optimisation performance is measured relative to a hypothetical 'perfect foresight' case by a model in key markets.

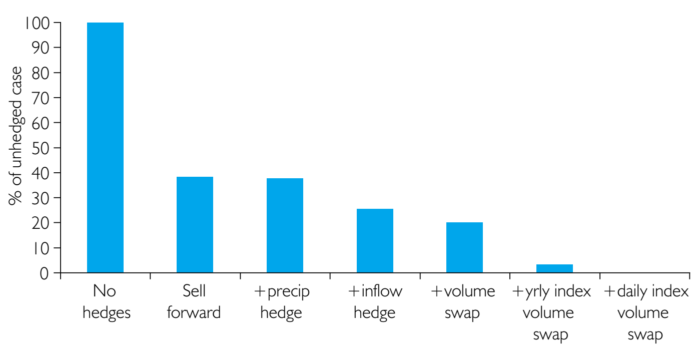

F9. Volume hedging effectiveness Derivatives coupled with forward sales could effectively manage price and volume risk *Source: Brookfield Power*

Even if a company does not have a full EAR measure, a solid understanding of volume risks can be created through similar methodologies.

The potential for derivatives

Given the limitations of price risk management tools in the presence of uncertain production volumes, the logical next step is to investigate techniques to directly offset volumetric risk, with derivatives whose payoffs are linked to precipitation being a natural place to start for hydroelectric producers.

Weather derivatives and output insurance products have been available for many years, and in some cases are now quite standardised, as a quick review of the products available on the Chicago Mercantile Exchange confirms. Because of the need for verifiable, independent data as a basis for final settlement, the most common weather derivatives are based in some manner on data from an independent third party such as MDA Federal (Earthsat) or the National Oceanic and Atmospheric Administration (NOAA) – often heating degree or cooling degree-days (HDDs/CDDs), the level of which can be correlated quite closely with utilities' loads, to take one example. NOAA maintains a network of precipitation gauges across the US, and publishes daily data series, which can in principle be used to design derivatives to manage hydroelectric volume risk.

Unfortunately, the linkage between precipitation and hydroelectric generation revenue is often less clear-cut than in the case of, for example, natural gas heating load and temperature. As shown in Figure 8, time lags between precipitation and reservoir inflows (especially in the case of snowfall), the existence of storage, as well as dispatch decisions (which create a site-specific exposure to a volume-weighted price rather than a market index) all serve to weaken the relationship between precipitation and realised revenue. Determining the ideal factors to apply to different gauges so as to maximise the correlation between the weighted precipitation measure and output is an additional difficulty. 'Mining' the data can find spurious correlations between series that have little or no out-of-sample predictive power. This problem can be particularly acute in the presence of limited data.

Using the same simple model employed to produce Figure 1, Figure 9 shows how a variety of possible derivative structures could aid in managing volume risk for the hypothetical station. All the cases presented ignore transaction costs, and simply employ standard deviation as a measure of risk – a useful starting point, notwithstanding the limitations discussed earlier. The first bar represents the standard deviation of the completely unhedged case, which is normalised to a value of 100. The most basic hedge, the selling forward of expected generation levels, is represented by the second bar. Measured in terms of standard deviation, this strategy eliminates most of the risk of the unhedged case for the simple reason that prices are more volatile than generation in the model – though of course selling forward expected generation creates the possibility of some highly undesirable outcomes when low volumes coincide with high prices!

The cases represented by each of the remaining bars in Figure 9 involve utilising a derivative designed to lay off volumetric risk, combined with a forward sale of expected generation. In the case of a derivative paying a fixed amount per annual millimetre of precipitation, the putative volume hedge has virtually no effect – indicating that the factors shown in Figure 8 are of more than just academic interest. Basing the derivative's payoff on annual inflows to reservoirs is a much more effective option, while a link to actual, annual generation cuts the standard deviation by almost 50% compared to the forward sale alone. A payoff based on the annual index price and actual generation volumes leaves very little residual risk – meaning that in this model, intra-year flows are not particularly 'spiky', and the exposure to the annual price level predominates. By definition, the product considered in the last bar – with a daily payoff equal to the daily index price times the shortfall of generation relative to its expected level – is a perfect volume hedge in this simple model.

Clearly, some of these possible derivatives are unrealistic – even if one could purchase them, their cost would be completely prohibitive. However, there are companies willing to give quotes for some of the precipitation-linked structures, though depending on the specific characteristics of a company's hydro assets, their efficacy may be limited.

Conclusions

It is certainly a time of innovation in the field of volume risk management, with room for companies to develop their own risk control frameworks and strategies. Weather derivatives and other structures are increasingly available as tools for companies seeking to lay off risks from uncertain production or loads. However, before embarking on a volume risk management program, particularly involving such tools, it is vital for a company's middle office to have a solid bottom-up understanding of its volume and price risk distributions, including sources of correlations between those two dimensions.

Also highly recommended is a risk control framework that recognises the challenge of volume risk and the limitations of traditional risk measures such as VAR. But perhaps the most important – yet often neglected – prerequisite is a clear statement of a company's volume risk strategy that is generally understood across the organisation. ER

Garth Renne is vice-president, enterprise risk management at Brookfield Power.
Email: garth.renne@brookfieldpower.com
Ken Truesdell is manager, market risk at Brookfield Power.
Email: ken.truesdell@brookfieldpower.com

The authors wish to thank Patrick MacDonald, storage trader at Cargill Power and Gas Markets, and Ed Mackay, formerly of TransAlta Corporation, for reviewing a draft of this paper.