Our take on an issue of the day ...

January 2011

In this edition

Temasek's sale of the Singapore gencos exemplified a process that clipped alertly along, especially given the sharp swings in global markets happening in the background. It is crucial to be able to identify and quickly quantify the fundamental sources of value – whether in pool-based or long-term contract-based settings – and to interact credibly and dynamically with relevant stakeholders. It is no small challenge to master the vast, dynamic region with its remarkably low proportion and uneven mix of private sector opportunities. Investors sometimes pay too much for assets in the region, perhaps because of conditions ripe for "winner's curse."

In this first of a series of Insights on valuation, we discuss an often overlooked source of competitive advantage – a modelling framework built for speed.

The authors:

Tom Parkinson
+852 9221 3037
tparkinson@lantaugroup.com

Mike Thomas
+852 9226 2513
mthomas@lantaugroup.com

THE LANTAU GROUP
strategy & economic consulting

Lantau **Pique**

Valuing Power Assets in the Asia Pacific Region: Using Models for Forecasting

The goal of a forecast is to support the analysis of a potential transaction or investment. The model is simply a means to an end. The model's purpose is not to replicate reality under narrowly defined specifications, but to support management decision-making. All models may appear to be black boxes, but not all black boxes are equally useful to the investor. It pays to be prepared by recognizing the intrinsic pros and cons of different models and modeling approaches.

Accuracy versus precision

"Cycle time" is a measure of how quickly new information can be processed into an evolving decision. Faster cycle time creates scope for better decisions. In the hands of an experienced team with a sound understanding of market fundamentals, a simplified modeling structure can greatly speed up the process of developing a view of value and risk. But what detail really matters?

Accuracy and consistency are vital, precision is not – and the distinction is important. A model must provide accurate and consistent estimates of key outputs under a wide range of conditions. But it is generally a distraction to replicate every last detail of dispatch – particularly given that data inputs are necessarily rough approximations of possible future conditions. Extraneous complexity taxes efficiency and is a source of spurious error, especially when deadlines place a high premium on clarity and focus.

Ironically, we've seen situations where a potential investor placed such importance on granular detail that it ignored a material overestimate of value in another, far less detailed, area. An excessive focus on detail may even indicate that the thing being valued is not well understood. Focus on what is most important!

Speed allows faster feedback, thereby simplifying the tuning and debugging process. Speed also creates opportunities to explore issues such as uncertainty, capacity expansion, and behavior of market participants in greater depth. Detail is important, but there are usually ways to structure analyses to cover for deficiencies in detail. In contrast, a lack of speed increases the risk and cost of poor judgment. Time is money, after all.

Maximizing speed requires a more tightly integrated client/consultant approach, which in turn depends on trust and experience. Few things slow a process down more than the need to establish trust. Establish trust well ahead of any need for maximum speed. Flexibility is equally vital. In the middle of a stressful valuation timetable, new information might come to light about constraints, obligations, or shifting market dynamics.

Value creation requires an ability to incorporate new information and then process it quickly, consistently, and insightfully. As important as "factual" knowledge (or the model itself) can be, it is the expertise and flexibility of the overall team that really drive value and insight. Not all teams who use models are equally adept or facile in their use. Life is just that way.

All models are not created equal

Models range from simple, back-of-the-envelope spreadsheets to sophisticated and intricately complex optimization packages. The key drivers, however, are often quite straightforward. The simplest form is a merit-order bid stack (as shown in Figure 1), which estimates the market price from the intersection of supply and demand. Similarly, contribution margins can be estimated directly via simple spreadsheet applications and statistical methods. Real-world constraints associated with system operation can invalidate insights from such simple approximations, however. The question is how far to go in pursuit of sorting out the impact of such constraints.

The available units are then stacked against these load blocks to determine the dispatch. This process makes the inherent assumption that any unit can be dispatched in any hour.

Chronological models

These models have the ability to respect real-world operating constraints that apply across consecutive hours – such as unit commitment, minimum generation levels, ramp rates, minimum on/off times, maintenance outages, and storage reservoir levels.

Clients sometimes expect the use of a chronological model, if only because these are common. The LDC model is less well understood and consequently lead to questions that can side-track an analysis. But an LDC model is ordinarily faster and easier to tailor to relevant questions. Dedicating less computing time to replicating system operation allows for more detailed analysis of the value impacts of uncertainty, capacity expansion and the behavior of market participants.

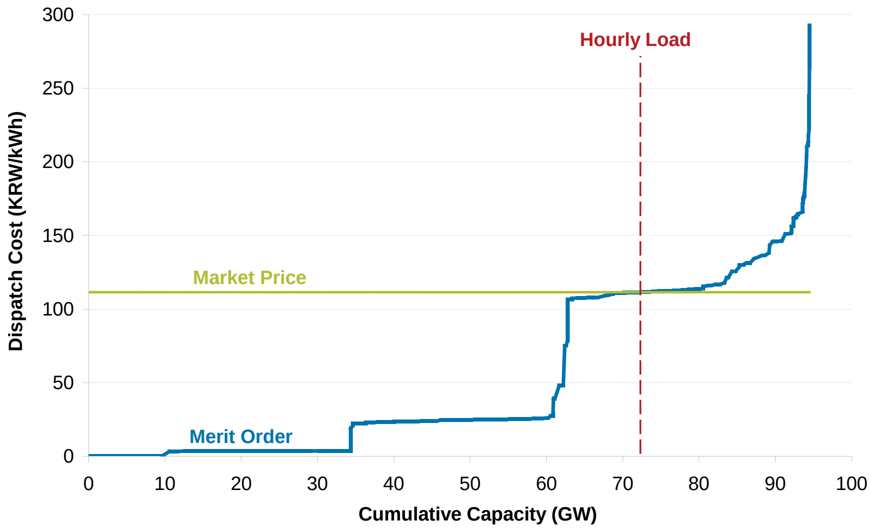

Figure 1: Merit Order and Price Formation in Korean KPX (2015)

The more technically inclined reader – and any serious buyer or seller of power assets should have a strong technical team supporting it – should find value in the discussion to follow. In it, we set out our views on the pros and cons of the two major classes of models used for commercial valuation purposes.

Load duration curve (LDC) models

These models take the 8760 hours of the year and segregate them into blocks of hours with similar loads.

What you want… what you need

In addition to speed, LDC models have a number of attributes that follow from their inherently simpler structure.

Focussed on key drivers

Most chronological modeling detail is spurious – it has little or no impact on the quantitative results. Stripping away irrelevancies makes it easier to focus on the key drivers of long-term prices and profitability – such as relative fuel prices, new technology costs and performance, and behavioral or regulatory influences on capacity expansion. As a result, important data can be represented simply, leading to greater clarity.

Easier to specify and debug

Structural simplicity requires less data, thereby speeding model specification, as well as making it easier to follow the logic and ensure the model works as intended. Indeed, the reality of modeling complex systems is that it is easy to introduce unintended errors, especially when working under time pressure. Use of a simpler structure means fewer places for modeling errors to hide, which means fewer errors will be made. As a result, forecasting results are less likely to reflect spurious modeling errors.

Straightforward

Structural simplicity makes it easier to communicate results and demonstrate how these results derive from input assumptions. For example, the LDC representation yields a related price duration curve expressing the number of hours in which the market price exceeds any value. In turn, a unit's contribution margin equals the area above its dispatch cost and below the price duration curve (see Figure 2).

Flexible

In skilled hands, LDC models can be suitably adjusted to approximate most salient characteristics of actual operating systems. LDC models can incorporate key constraints, such as those associated with fuel contracts or maintenance straightforwardly. For example, examining the effect of different assumptions on maintenance scheduling in an LDC model requires only that we vary peak and off-peak maximum capacity factors. Making equivalent changes in a chronological model can require changing maintenance schedules for every unit – which is tedious at best. Unnecessary complexity breeds unintended consequences and increases the risk of delay.

The straightforward nature of the LDC model is generally a source of strength in real world application. Of course one should always be mindful of the risks of simplification. But such concerns should be considered relative to concerns over the impact of losing speed or getting lost in the weeds. In our experience, chronological modeling is very rarely better for transaction support in the Asia Pacific region.

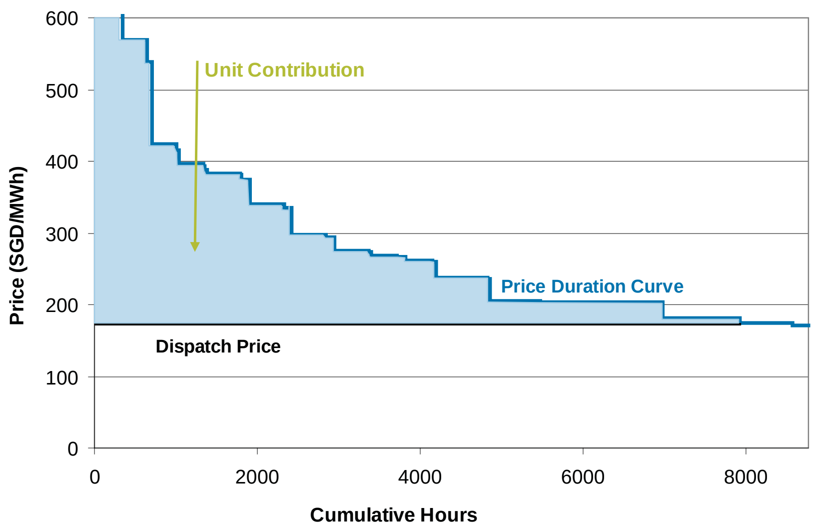

Figure 2: Price Duration Curve and Unit Contribution for Singapore NEMS

Beware of what you wish for…

In large part, the strengths and weaknesses of chronological models are the converse of those for LDC models. They provide a source of comfort to those – such as generation operators or planners – who are used to seeing operating detail and believe such detail is necessary to establish credibility. When run carefully with correct data, chronological models can produce stunningly accurate benchmarks against actual recent system operation. And there are some situations – such as systems with poor load factors and large amounts of inflexible generation – in which modeling ramping and cycling details is necessary to replicate some aspects of system operation. The structural detail in these models may even adjust semi-automatically for certain types of changes in underlying system conditions, allowing less-skilled analysts to produce highly detailed and accurate results. However, apparent accuracy is not a sound basis for choosing a model.

The problem is that chronological models embody a number of potentially underestimated sources of risk.

Optimization error

As detail increases in a chronological model, the underlying optimization problem grows exponentially. Larger problems typically need to be split into pieces that can be solved separately (and iteratively). There are good and bad ways to do this. Some models do not produce optimal solutions under some conditions, leading to inaccurate results. Because of the complexity of these models, determining whether the results actually satisfy the intended optimality conditions can be extremely difficult. The possibility of optimization error is difficult to detect, particularly by less experienced analysts.

Convergence failures

In some cases, the underlying optimization algorithm will lead to optimal results, if allowed to converge. But as the size of the problem to be solved increases, so to does the temptation to find shortcuts to speed them up. When convergence is slow, models may incorporate switches to allow reporting of results prior to achieving full convergence. Operators may even be unaware that the model is not converging. Failure to converge can produce unintended and virtually untraceable errors. Of course, ensuring full convergence may eliminate this risk, but only at the attendant cost of speed.

Data complacency

Replicating historical reality – which can require accurate specification of many data – provides confidence in the going-forward forecast. But this confidence may be misplaced. Replicating future reality to the same degree requires forecasting all of these operating parameters into the future – a realistically impossible exercise. In practice, most inputs are simply held constant, which may or may not actually provide an accurate forecast.

Biased prices

In an LDC model, the market price is the marginal cost of serving an incremental MWh of load. In chronological models, forming an unbiased proxy for the market price can be more difficult. The basic problem is that once a unit is committed, its marginal cost of operation is simply its incremental dispatch cost. Recovering this dispatch cost, however, may not compensate for start-up and no-load costs. Actual wholesale spot markets must augment the marginal cost-based prices in some way to ensure that generators are made whole across each commitment cycle. Chronological models may or may not incorporate appropriate price adjustment mechanisms. It is therefore possible to replicate dispatch precisely, yet provide poor estimates of market prices. LDC models must also account for these start-up and no-load costs, but the adjustment required is more obvious.

The perfect is the enemy of the good

For many commercial applications the LDC model offers material benefits related to speed, transparency, and simplicity/flexibility. They also have their limitations. Our view is that the perfect is the enemy of the good. Given what we know of the risks that affect investors in the Asia Pacific power sector, the LDC framework has real benefits. Short-term precision – as opposed to accuracy – can be a vastly overrated distraction. Longer time scales consistent with investment horizons introduce many key value-related risks whose quantification can be complicated or slowed by focusing on short-term precision. Market participant behavior is another area where precision matters much less than consistency and insightful scenario construction. It makes little sense to optimize one aspect of the overall forecasting problem to many decimal places, while ignoring other aspects.

Where speed matters and trust is high, the LDC model, in the hands of an experienced team, can be the client's best friend – offering faster insight, greater flexibility, and an ability to focus on what really matters.

For more information about The Lantau Group Limited, please visit us at www.lantaugroup.com