1-2008

Probabilistic sales forecasting for Small and Medium-Size Business Operations

Randall E. DURAN
Singapore Management University, rduran@smu.edu.sg

Probabilistic Sales Forecasting for Small and Medium-Size Business Operations

Randall E. Duran

Catena Technologies Pte Ltd, 30 Robinson Road, Robinson Towers #11-04, Singapore 048546, Republic of Singapore

Abstract. One of the most important aspects of operating a business is the forecasting of sales and allocation of resources to fulfill sales. Sales assessments are usually based on mental models that are not well defined, may be biased, and are difficult to refine and improve over time. Defining sales forecasting models for small- and medium-size business operations is especially difficult when the number of sales events is small but the revenue per sales event is large. This chapter reviews the challenges of sales forecasting in this environment and describes how incomplete and potentially suspect information can be used to produce more coherent and adaptable sales forecasts. It outlines an approach for developing sales forecasts based on estimated probability distributions of sales closures. These distributions are then combined with Monte Carlo methods to produce sales forecasts. Distribution estimates are adjusted over time, based on new developments in the sales opportunities. Furthermore, revenue from several types of sources can be combined in the forecast to cater for more complex business environments.

1 Introduction

One of the most important aspects of operating a business is the forecasting of sales and allocation of resources to fulfill sales. While customer relationship management and enterprise resource management systems have helped provide larger organizations with methodologies for forecasting and measuring sales, many small- and medium-size business operations rely on the assessment abilities of one or more sales managers to estimate and manage the sales pipeline. In turn, the outputs of these forecasts as provided to operational mangers are often overly simplistic and do not provide sufficient depth to support resource planning adequately.

The fundamental problem is that sales assessments are usually based on mental models that are not well defined, may be biased, and are difficult to refine and improve over time. Furthermore, defining sales forecasting models for small- and medium-size business operations is often difficult when the number of sales events is small but the revenue per sales event is large. Little or no data may be available to support estimations, and the effect of assumptions and uncertainties is magnified when only a few sales make up the entire revenue stream. Hence, it is not practical to use traditional budgeting processes that set a fixed annual budget when there is a high degree of variability in the expected revenue.

Soft computing approaches, such as Monte Carlo techniques, have been used to determine the general range of possible outcomes where information that is critical to planning is either unknown or unreliable. Savage (2002) focused attention on the problems of using single point estimates in corporate accounting and how Monte Carlo techniques can be used to reflect a more realistic state of corporations' finances. Monte Carlo techniques have also been championed for use in a wide range of

B. Prasad (Ed.): Soft Computing Applications in Business, STUDFUZZ 230, pp. 129–146, 2008.
springerlink.com

applications related to risk modeling and risk analysis (Koller 2000). There has also been research into using adaptive processes and rolling forecasts for budget planning (Hope and Fraser 2003). The objective of the research described in this chapter is to develop an approach that combines computer-based Monte Carlo analysis with adaptive business planning to help provide more accurate business forecast information and support operational resource planning.

This chapter reviews the challenges of sales forecasting in this environment and describes how incomplete and potentially suspect information can be used to produce more coherent and adaptable sales forecasts. It outlines an approach for developing sales forecasts based on estimated probability distributions of sales closures. These distributions are then used with Monte Carlo methods to produce sales forecasts. Distribution estimates are adjusted over time, based on new developments in the sales opportunities. Furthermore, revenue from several types of sources can be combined in the forecast to cater for more complex business environments.

The first section of this chapter presents an overview of the business problem, in context of small- and medium-size business operations. The second section describes the soft computing approach used to address the problem, including an overview of the method and a description of some of the underlying mechanisms. The third section summarizes results of applying this approach in real-world scenarios involving individual, partner, and existing-customer sales situations. The results are then evaluated with regards to the value that they can provide to operational managers, who must base their resourcing decisions on the forecasts. The final section summarizes the benefits of using these techniques and provides suggestions for future areas of research.

2 Overview of the Business Problem

In high value, low volume businesses (HVLVBs) misestimating a single sale can have significant impact on the business. It is also difficult to extrapolate future sales for HVLVBs because the statistical benefits gained through the laws of large numbers are not applicable. There may be few, if any, historical points from which to make inferences, and comparative benchmarks or baselines may not be available. Furthermore, sales forecasts may be affected by events that unfold over time, requiring the forecasts to be changed frequently. Thus, for these types of businesses, planning resource allocation and capital investment according to sales forecasts can be risky. As background information, the following paragraphs provide specific HVLVB examples and describe typical revenue forecasting tools used for planning.

Note that the terms "revenue" and "sales" are used interchangeably, since sales are the primary source of revenue for most small- and medium-size businesses. Likewise, the terms "deal" and "sale" are used interchangeably. Furthermore, a deal or sale "closing" refers to the achievement of contractual commitment that will lead to realization of the projected revenue.

2.1 Business Examples

To help frame the challenges presented when forecasting sales, consider the following two examples. In each example, the businesses' sales are characterized by high value – greater than $100,000 – and low frequency – less than 10 per year – transactions.

They might include business operations that provide highly specialized product and service offerings, such as industry-specific software products, high-end consulting services, and project-focused deliveries.

In the first example, an established company starts a branch office in a foreign country. While the characteristics of product sales may be understood in the domestic market or in other countries, it quite likely that those characteristics will not translate into the new market. Important factors such as branding, cost, and demand may be significantly different in a new market and may not be easily quantified in advance. Hence, uncertainty regarding the expected sales can in turn lead to uncertainty as to how much revenue should be forecast for a new market and how much capital should be invested. As a case in point, many firms have entered new markets, such as China, with unrealistic expectations about how quickly they would be able to generate significant revenue, and they have suffered unexpectedly large losses as a result.

In the second example, consider a startup business that has been operating for only a short period of time. While the general market characteristics may be understood, the willingness of customers to purchase from a newly formed company with no track record may be difficult to quantify. Furthermore, uncontrollable market influences such as natural and man-made disasters may significantly affect the overall economy and, thus, demand for the product. In this situation, deciding when and how much to invest is difficult due to the large number of uncertainties that underlie the assumptions regarding potential revenue.

To simply illustrate the challenge in these examples, imagine that only two sales events are forecast to close within the next six months. In January, one sale – worth \$200,000 – is forecast for March, and the other sale – worth \$500,000 – is forecast for May. Based on the little information available, the business manager then has to determine how much to invest on presales and sales fulfillment capability (i.e. delivery resources). Table 1 shows the potential outcomes when, using a simplistic approach, a half-year spending budget of \$400,000 is allocated and utilized.

Table 1. Results of applying a fixed budget with high sales variance

Possible Outcomes	Resource Investment	Revenue	Financial Effects - P&L	Potential Business Impact
Neither sale closes	\$400,000	\$0	-\$400,000	Need to raise more capital or shut down operation
Opportunity 1 closes, opportunity 2 does not	\$400,000	\$200,000	-\$200,000	Overcapacity: operating losses and reduction of staff
Opportunity 2 closes, opportunity 1 does not	\$400,000	\$500,000	\$100,000	Reasonably matched investment and returns
Both sales close	\$400,000	\$700,000	\$300,000	Under capacity: poor delivery quality, staff turnover, and cancellation of contracts

Note that in this example; only one of the four cases – the third one – produces a fairly positive result. While the fourth case produces the greatest profits, it could have significant negative non-financial results that are likely to impact future revenues. While this scenario is useful for illustrating the fundamental problem, it is important to understand that it has been greatly simplified. The timing of closure, the sale amount, and the likelihood of closing are all estimates that may not be accurate.

Moreover, even businesses that have been operating for several years can also face difficulties forecasting revenue. Consider a mature business whose revenue comes from regular purchases from existing customers, large lump sums from new customer sales, and sales made through external channels such as distribution partners. Different types of uncertainties will underlie each of these revenue sources: sales personnel's estimates may be afflicted by common decision biases (Hammond et al. 2006); repeat business may be inconsistent with respect the timing of orders and amounts; and distribution partners' agreements may be oriented towards annual targets, hence providing little insight as to when revenue, and delivery commitments, will come through. Hence, more mature HVLVBs have similar challenges to new HVLVBs, but with much greater complexity.

2.2 Revenue Forecasting Tools

A common tool used for managing HVLVB sales information is a sales pipeline forecast, which is often embodied as a spreadsheet. This spreadsheet lists potential sales opportunities, their expected closure dates, their percentage chance of closing, the person or entity responsible for closure, and, based on those values, calculates the total expected revenue. While, this technique is helpful for tracking the status of potential sales, its revenue-estimating ability is limited because a simple spreadsheet model does not adequately reflect the true complexity of the sales environment.

With this simple forecasting model, the input estimations related to the percentage chance of closure misrepresents the "all or nothing" nature of sales, especially when the number of sales is small. For example, while ten sales opportunities worth $100,000 each and with 20% likelihood of closing could well (but may not) yield $200,000 revenue, a 20% chance of making a $1,000,000 sale will not. Even the amount forecast is unlikely to be the exact amount realized. The actual dollar amount realized could be greater or smaller, depending on factors such as the agreed final scope of the sale, price negotiations, and fluctuations in foreign exchange rates.

Given the variability of HVLVB sales, it is critical for an experienced sales manager to assess and quantify expected sales revenue. Sales managers will often use the sales pipeline spreadsheet as tracking mechanism, but not as the sole basis for their sales forecast. Instead, they adapt and enhance the calculations provided by the spreadsheet using their own unique estimating methods – i.e. heuristics. Their estimations can be based on knowledge of previous sales results, local purchasing patterns, and the capabilities of individual sales people, and the behavior of customers. Thus, the resulting forecast is highly dependent on the sales manager's personal views and experience, rather than on a system. This relationship is problematic because such personal estimation methods will vary between individual managers, and there is little transparency as to what the individual's methods encompass or are based on. Ten sales mangers could produce significantly different revenue forecasts for the same sales opportunities. Not surprisingly, these highly individual-dependent estimations are not necessarily consistent over time, and can be difficult for others to verify.

Sales forecasts are central to maximizing profits. Profits are achieved by keeping the cost of delivering sales less than actual sales. The cost of delivering sales is largely determined by to the costs of employing resources for sales, presales, and product delivery. It takes time to acquire and develop these resources to the point where they are fully effective. Accurate forecasts are critical to determine how many resources are required and how far in advance they should be secured. Underestimation of sales leads to underinvestment, lower revenue, and lower income at best. Overestimation of sales leads to overinvestment, potential losses, and, in the worst case, business failure. Because the profits of HVLVBs depend on a small number of sales, accurate sales forecasts are especially critical – they can make or break these types of businesses.

2.3 Adaptive Processes

It is difficult, and often not practical, for small- and medium-size businesses to make annual budgets that are fixed and are not affected by the sales volume. Hope and Fraser (2003) argue that this approach is not beneficial for large organizations either, and that more dynamic and adaptive processes are required. Instead of keeping spending budgets fixed, they should be adjusted depending on changes in business over time. As the realization of sales becomes more likely, the budget should be increased, and vice versa.

While dynamic budget adjustment is useful as a concept, its implementation has limitations. Some costs factors, such as travel and capital expenditures, can be increased, reduced or postponed quickly. However, other costs, such as human resource costs, may require significant lead times and may not be easily decreased due to labor laws or overall staff morale considerations. Thus, it is not practical to reevaluate and adjust resourcing on a monthly basis. A more practical approach is to do budget planning based on rolling forecasts – possibly monthly or quarterly – and that may be reassessed on a monthly or quarterly basis, and to take proactive steps to help prepare for the expected outcomes of the forecast. The next section discusses how Monte Carlo techniques can be integrated with such an adaptive forecasting model to yield a Probabilistic Sales Forecasting (PSF) system.

3 A Soft Computing Approach

To improve HVLVB sales forecasting, more complex information must be incorporated into the forecasting model. Partial truths, related to the likelihood of sales completions, must be given better focus. The soft computing approach takes the existing simple sales-pipeline model and enhances it by adding distribution-oriented estimates for the probabilities of sales closures, and allows for adjustment of those distributions over time.

Based on these enhancements to the forecasting model, Monte Carlo simulation techniques are used to assess the potential range of expected sales outcomes. This enhancement transforms the simple pipeline evaluation model from single-value input estimates and a single-value calculated result to a set of inputs and results that reflect the uncertainty of the situation. This approach emulates the implicit mental models that sales managers use to develop their forecasts. But, unlike those implicit estimation methods, it produces much richer forecasting information and provides greater insight into how the estimates of individual sale opportunities relate to the final result.

3.1 Monte Carlo Techniques

In his seminal paper "Risk Analysis in Capital Investment", Hertz (1964) presented a method for creating forecasts and decision models that use range-based approximations of the uncertain variables. This approach has evolved over time as the basis for Monte Carlo simulation techniques. Although software-based tools and computer platforms that implement Monte Carlo simulations have improved tremendously since Hertz's initial work, the overall approach has not changed significantly. The three principal steps that form the basis of Monte Carlo simulations are

1. Estimate the range of values and overall likelihood of those values for each important input variable in the forecasting model, so as to define a probability distribution for each of the variables
2. Simulate random occurrences, based on the defined probability distributions, for each of the input variables estimated in step 1, and measure the resulting value produced by the forecasting model
3. Repeat step 2 many times and record the forecasting results for each simulation. The resulting values represent the likelihood of occurrence for different forecast results, and can be represented as probability distributions for the forecasting model output variables.

These techniques were initially developed for capital investment budgeting purposes and risk analysis and were implemented by custom computer programs. Monte Carlo analysis allows more complex investment scenarios to be modeled more accurately and allows different investments to be compared on their likelihood of different outcomes, rather than just their expected return.

Over the past twenty years, use of Monte Carlo techniques has expanded beyond capital investment, and is used for enterprise risk management, natural resource exploration, and project risk analysis. Furthermore, custom development of computer programs to implement Monte Carlo simulations is no longer necessary. Several tools are available that allow Monte Carlo simulations to be incorporated into spreadsheet models. Mun (2004) provides instruction and examples for applying Monte Carlo analysis using such tools.

The objective of this research is to apply Monte Carlo and adaptive process techniques to the real-world business problems faced by small- and medium-size business operations. Whereas much of the focus of Monte Carlo analysis has been on quantifying risks related to revenue shortfalls, this work focus on the risks of achieving too much as well as too little revenue.

3.2 The Probabilistic Sales Forecasting Model

As shown in Figure 1, the PSF model calculates the expected revenue for the business operation based on multiple sources of sales revenue. The revenue sources considered

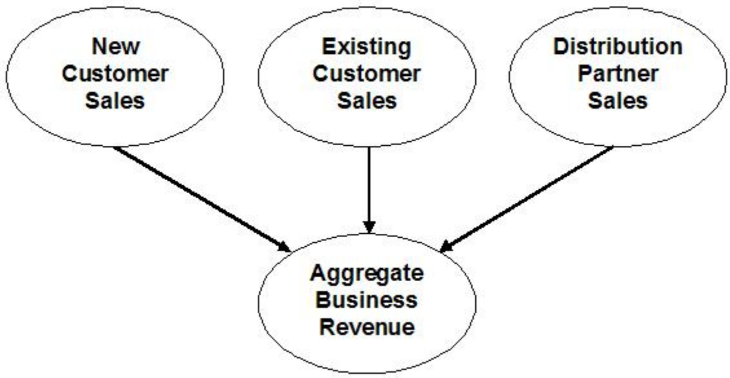

Fig. 1. Relationship of potential revenue sources to overall operating revenue

include new customer sales, existing customer sales, and distribution partner sales. The aggregate business revenue is the sum of revenue from these sources. Depending on the type of business being modeled, it could be the case that one or more of these sources are not relevant; alternatively, other sources of revenue could be applicable. The PSF model is flexible, allowing relevant revenue sources to be included as necessary. Beyond aggregating multiple sources of revenue, the objective of this model is also to factor in the uncertainties and unknowns. Therefore, the likelihood of future sales is also considered.

The key dimensions of the PSF model are the time horizon for revenue acquisition (i.e. sales closure) and the likelihood of achieving various revenue amounts during that time. The PSF model uses a static time horizon of six months. Depending on the type of business, sales closures will often take more than three months to complete, so using a horizon of three months is too short and could easily underestimate actual revenue. On the other hand, it is often difficult to make estimates regarding business with much accuracy beyond six or eight months. Hence, choosing a forecasting horizon of one year would be too speculative as to future events, making the resulting forecast more inaccurate and unreliable.

Six months was chosen as the forecast horizon for the PSF model, as it represents a middle-ground timeframe which is able to capture near-term sale completions at a reasonably accurate level and also allows for sufficient advance notice to support forward planning and implementation of resource adjustments. This strategy is to produce or revise six-month rolling forecasts on a quarterly basis, every three months. However, the forecasts could be generated on a monthly basis, if desired, or updated immediately after major events, such as verbal confirmation of a major new sale, or the sudden and unexpected loss of an existing customer or partner.

The PSF model represents forecast revenue for a six month period with a single output distribution. This result is an aggregation of the potential revenue that could be generated by new sales, existing customers, and distribution partners. The inputs are also probability distributions, and are aggregated using Monte Carlo simulations. The input distributions are constructed based on simple estimates of the sales potential value, specific commitments by customers and partners, and historical data, when it is available. The following will now discuss how this information is factored into the PSF system.

3.3 Defining Base Revenue Distributions

Spetzler and Stael von Holstein (1975) describe detailed techniques for quantifying individual judgments about uncertain quantities. The objective of using these types of assessment and encoding techniques is to express probability distributions in a form that can be easily inspected, validated, and examined for judgmental biases. Unfortunately, the scope of the PSF system does not allow for interview-oriented techniques to be used for determining the probability distributions. As an alternative, normal distributions were adapted in various ways for new sales, existing customer sales, and distribution partner sales. The expectation is that these simple probability distributions would be refined in the future based on experience with the PSF model, or enhanced using interview techniques at a later time. For simplicity, all of the revenue types were modeled as variants of a normal distribution. However, other distribution types may be more appropriate, and this is an area of future research. For the first model implementation, increased precision was traded for reduced complexity. This choice made it easier to implement and maintain, and a simpler model was easier to explain to and review by business stakeholders.

The probability input distributions are estimated somewhat differently for new customer sales, existing customer sales, and distribution partner sales revenue. For new customer sales opportunities, the PSF system models the expected revenue returns as normal distributions that are scaled and adjusted based on four estimated parameters: expected deal size, minimum deal size, maximum deal size, and likelihood of closure. The expected deal size is the “headline” number representing the most likely value of how much the deal is thought to be worth.

Translating the PSF model input parameters into a probability distribution for Monte Carlo simulation is accomplished in two steps. First, the mean of the distribution is defined as the expected deal size. Second, the standard deviation of the distribution is defined as a percentage of the expected deal size, based on the likelihood of closure as shown in Table 2. The method for determining the likelihood of closure is discussed later in this chapter.

As an example, in the early stages a sale the expected size might be of $270,000 with a likelihood of closure at 40%, the baseline distribution would have a mean of $270,000 and a standard deviation of 0.3 x $270,000 = $108,000. As the deal progresses and goes into contractual negotiations – based on refined scope and

Table 2. Rules for determining the standard deviation of new sale input distributions

Expected Likelihood of Success	Standard Deviation
0 – 35%	40%
36-79%	30%
70-90%	20%
90-100%	5%

understanding with the customer – the expected deal size might be reevaluated to be $220,000 and the likelihood of closure adjusted to 80%. In turn, the baseline probability distribution for the sale would be adjusted to have a mean of $220,000 and standard deviation of $220,000 x .20 = $44,000, reflecting increased and more accurate knowledge of the sales situation.

Existing customer sales distributions are estimated based on the mean and standard deviation of historical sales over the previous 18 months. Sale cash flows for consecutive six month periods are calculated, treated as samples of a distribution, and then translated to a distribution to be used for Monte Carlo simulation purposes. Figure 2 shows the monthly sale cash flows, a histogram of their distribution, and the resulting base-probability distribution for an existing customer. Note that the small number of samples is a weakness of this approach. However, using a larger sample – i.e. a longer period of time – may produce even more skewed results if there is a significant growth or reduction trend in progress, which is not uncommon for young businesses. If there are not 18 months of past data available, the customer's potential sales are treated the same way as new customer sales.

For distribution partners the input distributions are determined based on whether the partnership agreement has a revenue target defined. If no revenue target is agreed in advance, expected revenue is estimated based on knowledge of the partners and the opportunities to which they have access. If a target is agreed, the mean is defined as the target amount. The standard deviation to be used for partners, as shown in Table 3, varies depending on the partner maturity and target arrangement. A relatively high standard deviation is used in the case of new partners with no commitment, because of the lack of historical experience and visibility into the potential revenue stream. New partners with a commitment are given a lower standard deviation, to reflect an increased level of confidence in the revenue prediction. Historical sales data is used to determine the standard deviation for more mature partnerships, where there is at least 18 months of historical sales data available.

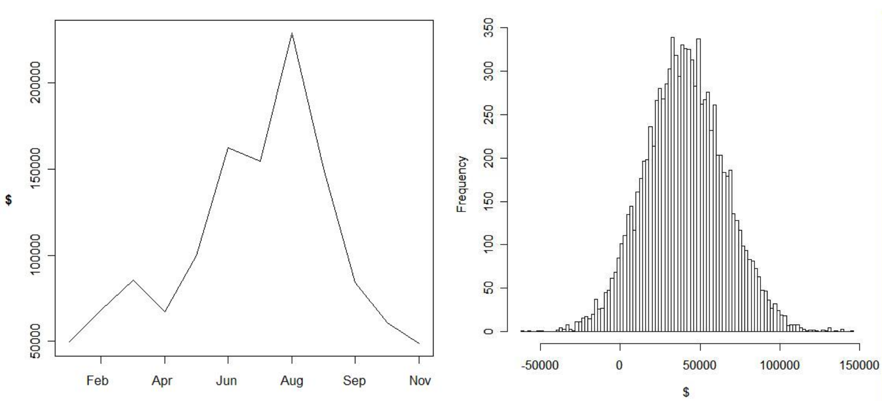

Fig. 2. Existing monthly sales historical data and approximated distribution histogram

Table 3. Rules for determining the standard deviation for partnership revenue

Partnership Characteristics	Standard Deviation
New Partnership, No Commitment	50%
New Partnership, Commitment	25%
Mature Partnership, Historical Data	Based on Historical Data

3.4 Adjustments to the Base Distributions

While normal distributions are a convenient means of modeling the expected amount of revenue that will be generated, actual sale results are not well served by this representation. Therefore, several adjustments to the base input distributions are necessary to make the PSF model more accurate. Specifically, concrete minimum and maximum sale amounts must be incorporated, and the overall likelihood of successfully realizing the revenue must be factored into the simulation model.

Within the PSF model the normal distributions are adjusted by expected minimum and maximum expected revenue potential. The minimum and maximums sale amounts are estimates based on factors such as the minimum size of a sale that is cost-effective for the business to undertake. The maximum may be related to the overall budget or sign-off authority of the purchasing party or, alternatively, the largest size project that the business could deliver with its current or near-term resource capability. During the PSF system simulations, if an input variable fell below the minimum amount, it was treated as a zero value case. If the simulation variable was greater than the maximum amount, it was capped and set as the maximum amount for that simulation instance.

The normal distributions are also adjusted by the sale's likelihood of success. There are many reasons for revenue to not to be realized at all: new customer sales may not close due to competition or change of plans by the customer; partnerships may be severed unexpectedly, and existing customers may choose or be economically forced to end business relationships. To reflect the overall uncertainty of closure, the model incorporates a simulation variable that is a probability distribution of either zero or one, where

$$p(0) = 1 - \text{likelihood of closure}$$
$$p(1) = \text{likelihood of closure}$$

The PSF system multiplies the expected revenue amount simulation variable by the likelihood of success simulation variable, both of which were independent variables.

After applying the minimum and maximum adjustments and factoring in the likelihood of closure, the general shape of the input sales revenue distribution used for the overall business revenue aggregation will be similar to the distribution shown in Figure 3. The possibility of non-closure creates a spike at zero revenue. The lower bound is clipped by the minimum and the higher bound is capped by the maximum. A small spike occurs at the maximum expected revenue amount, since simulation cases that exceed that amount are rounded down to the maximum amount.

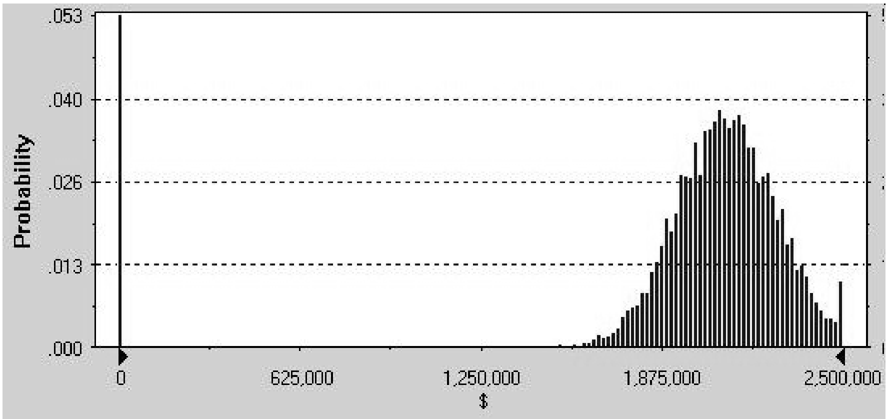

Fig. 3. Adjusted new customer sale probability input distribution

3.5 Likelihood of Success Estimation and Adjustment

While allowing inexact information to be factored into the forecasting process, the PSF system links the likelihood of success to measurable events in the sales cycle. As a new customer sales opportunity progresses, there are different events that will occur that affect the likelihood of its revenue being realized. A company may initially be competing with many others for a particular sale. Being short-listed (or not) for a sale will significantly affect the likelihood of the revenue being realized in the future. Thus, the model's likelihood-of-closure variable is adjusted in accordance with events that affect the sale. Table 4 shows the event-based rules that are applied within the model. While, initially, personal biases may affect the likelihood of closure assigned to various stages of the sales lifecycle, the intention is that the comparison of model performance to actual performance over time can be used to refine the values and eliminate these biases.

Table 4. Rules for determining the likelihood of success for new customer sales

Level	Event	Likelihood of Closure
0	Discussions with Customer	0%
0	Selection of competitor	0%
0	Notification of non-purchase	0%
1	Budget Amount and Decision Timeframe Confirmed	10%
2	Short-listed	40%
3	Pre-engagement Funded	55%
4	Verbal Confirmation of Selection	60%
5	Written Confirmation of Selection	70%
6	Contractual Negotiations Begin	80%
7	Contractual Negotiations Completed	95%

While the specific events and probability values will likely vary from one business to another, the principle behind the adjustment is broadly applicable. For both distribution partner and existing customer sales the likelihood of success reflects the probability that the partnership may be unexpectedly terminated in the next six months. Typically, in context of the lifecycle events of new customer sales, if sales efforts are not gaining momentum, they will be losing momentum. Therefore the PSF system applies rules for assessing the effect of delays in the closure of sales. Using formalized criteria in this way helps make the evaluation process more objective and consistent between different sales opportunities, and minimizes personal biases. The model uses these time-based criteria as a means to reduce the chances of improperly qualified sales or changing situations skewing the results.

For example, where a company may be short-listed for a sale, it may turn out that the customer decides – after many months or even years – to not make any purchase. In this situation, it would skew the revenue projects results to keep the opportunity projected at 40% probability for all of that time. Hence, when one of the events listed in Table 4 occurs, an opportunity closure date is either initially estimated or redefined. If the closure date is past and the sale is not closed, the likelihood of closure factor is reduced to the next lower event level and the closure date is reset to be three months later. This approach allows for the likelihood-of-closure to be increased while the sale is progressing in a positive direction, but reduces the closure likelihood once momentum stops and no progress is being made.

Continuing with the same example, consider that in January the company was short-listed for an opportunity and the closure date was April. If, in April, when the sale was expected to be completed, the company was still at the short-listed stage, then the time-based criteria would cause the likelihood of closure to be reduced from 40% to 10%. This adjustment reflects that the initial estimations regarding the customer decision-making process for closure were overly optimistic, and should be discounted until another positive event occurs. If verbal confirmation of selection was provided in May, the likelihood of closure would be reset as 60% and a new closure date would be estimated. However, if the verbal confirmation does not occur by the end of July – three months after the initially predicted closure date – no further progress would be deemed to have been made and the likelihood of success would be again reduced from 10% to 0%.

The time-based adjustments are applied only to new sales, since in most HVLVBs they have the highest value and the greatest transparency. Adjustments related to changing business conditions over time for existing customer sales are incorporated into the model by using the previous 18 month sales history as the basis of the base probability distribution. Distribution-partner sales estimates are not adjusted on a rolling basis; rather, they are expected to be readjusted on an annual basis based on the previous year's performance and coming year's targets.

4 Results

Having defined the basis of how the PSF system, with the help of sales personnel and managers, determines the input distributions for the business's potential revenue streams, Monte Carlo simulations are then used to calculate the expected revenue for a six month period, producing a aggregated return distribution of revenue. The results of the simulation, which are non-normal, enable managers perform analysis in more depth and answer questions like: “what is the likelihood that the total business unit revenues will be less than $200,000 this quarter?” A sales person could also use the model to get an objective assessment of the total revenue they are likely to achieve at a 50% or 80% level of confidence.

Application of this approach is illustrated using three real-world scenarios related to individual, partner, and existing customer sales in a HVLVB. The first scenario forecasts the sales pipeline of a new sales person, who has great potential, but a limited track record. The second scenario estimates sales that will be generated by a new distribution partner. The third scenario relates to estimating follow-on business with an existing customer. Finally, all three scenarios are combined in the context of a business operation that has its own sales people, distribution partners, and existing sales account managers, who all contribute to the business operation’s revenue.

4.1 Estimating New Sales

The first scenario forecasts the sales pipeline of a new sales person, who has great potential, but a limited track record. The model’s moderately conservative likelihood-of-closure assumptions are used and then adjusted as sales opportunities materialize or fail to be achieved. This model helps objectively measure sales people’s progress over time and predict their overall success. It also can help improve their ability to estimate sales closures based on their historical accuracy.

In June, they have forecast two sales to close in the next six months. The parameters of the two deals are shown in Table 5.

Table 5. Sale parameter determining revenue probability distributions for two new customer sale opportunities

Sale Parameters	New Customer 1	New Customer 2
Min	$600,000	$1,300,000
Expected	$1,000,000	$2,100,000
Max	$1,500,000	$2,500,000
Stdev	$400,000	$420,000
Closure Date	October	November
Event Level	1	5
Event Description	Budget and Decision Timeframe Confirmed	Written Confirmation of Selection
Likelihood of Success	10%	70%

Figure 4 shows a cumulative chart of the expected combined revenue from both opportunities. From the resulting distribution we can see that there is about a 30% chance that no revenue will be generated by either sale and about a 50% chance that the sales will generate more than $1,800,000 of revenue. A simple non-probabilistic model could have estimated of the maximum revenue as the sum of the maximum amounts of the two potential deals, $4,000,000. However, the PSF model shows that there is less than a 5% chance of making more than $2,500,000 from new customer sales.

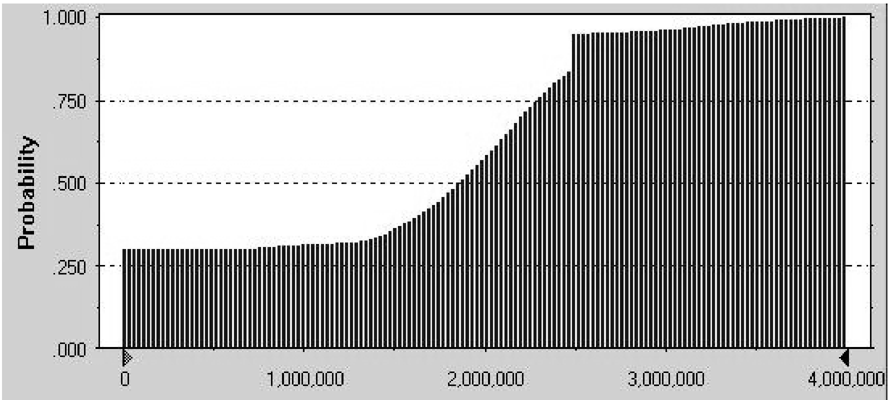

Fig. 4. Cumulative probability chart of expected revenue from two combined new customer sales

4.2 Estimating Partner Sales

The second scenario estimates sales that will be generated by two distribution partners. Compared to solely relying on the partners' commitment, the PSF technique can help account for the potential variance in a partner's revenue contribution and can

Table 6. Sale parameters determining revenue probability distributions for two distribution partners

Sale Parameters	Partner 1	Partner 2
Commitment	No	Yes
Expected	$250,000	$150,000
Stdev	$125,000	$37,500
Likelihood of Success	95%	95%

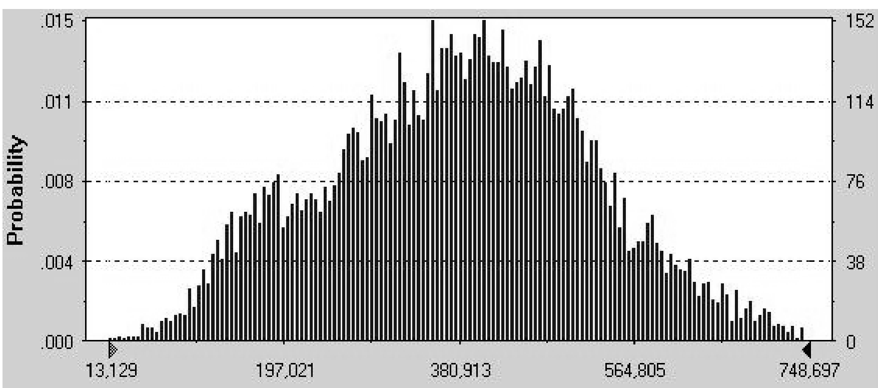

Fig. 5. Histogram of expected combined sales of two partners

help assess whether a partnership should receive more or less investment. Figure 5 shows the probability distribution of the expected combined sales of two distribution partners.

4.3 Estimating Existing Customer Sales

The third scenario relates to estimating follow-on business with two existing customers.

Table 7. Sale parameters determining revenue probability distributions for two existing customers

Sale Parameters	Existing Customer 1	Existing Customer 2
Expected	$277,000	$623,000
Stdev	$75,000	$155,000
Likelihood of Success	95%	75%

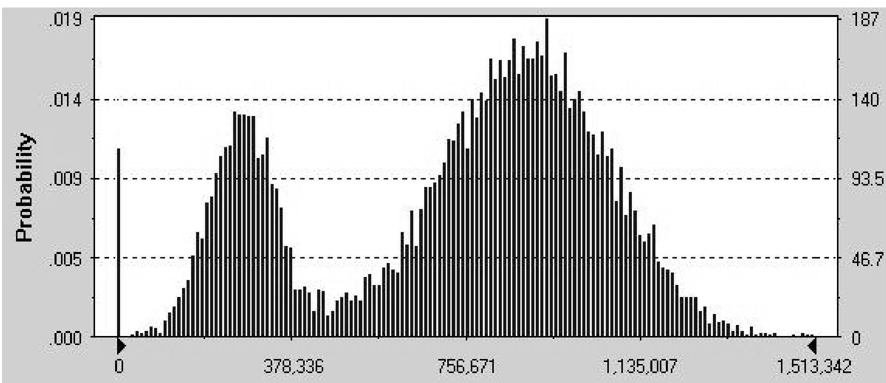

Fig. 6. Histogram of expected combined sales of two partners

The likelihood of success reflects the likelihood that the relationship may be unexpectedly terminated in the next six months. Whereas the business with customer 1 has somewhat higher variance in percentage terms, the relationship is quite solid. In comparison, the relationship with customer two may not be as stable or the overall commitment period may be much shorter, reflecting a higher possibility of the entire revenue stream dropping off during the six month period.

Figure 6 shows the expected revenue from both customers for the six-month period. The result is a distribution that is bimodal and also has a spike at zero, representing the potential loss of both customers.

4.4 Aggregated Forecasts

While the PSF system can be used to generate aggregated sales estimates for each of individual scenarios identified, for operational and resource planning purposes it is more important to look at a forecast that the model would generate for the aggregate business produced by these different scenarios. Therefore, the three scenarios have been combined in the context of a business operation that has its own sales people, distribution partners, and existing sales account managers, who all contribute to the business operation's revenue. By combining the underlying probabilities for each of the sales participants' opportunities within a single Monte Carlo simulation, an aggregate distribution of expected sales for the entire business operation is forecasted. Figures 7 and 8 show the histogram and cumulative probability views of the aggregate revenue forecast.

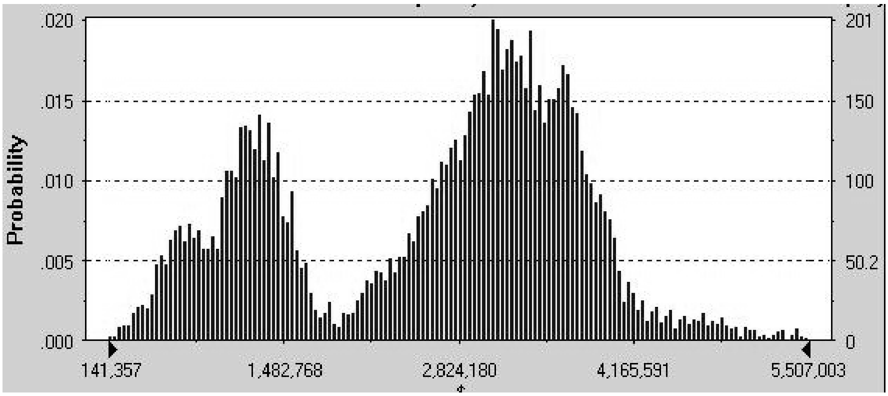

Fig. 7. Histogram of expected aggregated sale channel revenue

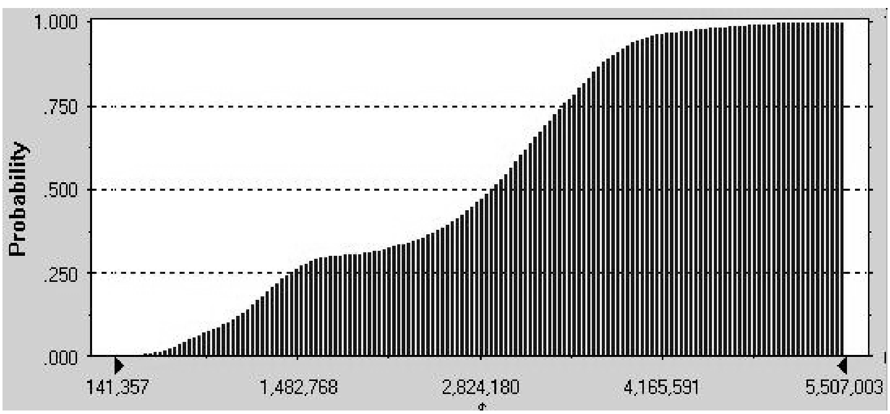

Fig. 8. Cumulative probability chart of expected aggregated sale channel revenue

While the mean, approximately $263,000, could be used as a simple, single-value revenue forecast, the distribution provides a better picture of all the possible outcomes. Using this distribution, further analysis can be easily performed to determine that: while the maximum possible revenue is over $5,500,000, there is less than a 10% chance of realizing more than $3,850,000; there is a 75% chance of achieving at least $1,460,000 in sales revenue; and there is only about a 10% chance that the total revenue will fall between $1,500,000 and $2,500,000. All of these conclusions can be determined by assessing the point at which the curve, as shown in Figure 8, intersects with the respective dollar amounts. This depth of information enables business managers to understand the range of possibilities better and, in turn, more effectively manage the costs and investments that are required to deliver sales.

4.5 Utilizing Probabilistic Forecasts for Resource Planning

Staffing and resourcing is a significant challenge, even when customer demand is well understood. There are many trade offs in balancing operating costs versus readiness for delivery, hiring and development of permanent staff versus the use of contractors, and setting prices to manage the volume of business, and thus, resources required. Adding customer demand as another predominant variable makes resource planning all the more difficult.

The objective of producing sales forecasts in the form of distributions is to help provide more transparency into possible range of outcomes in customer demand, so that resource planners can better determine how best to adjust the other variables under their control. An underlying assumption is that sales revenue can be, more or less, directly translated into resource requirements to deliver the product. This translation often can be roughly calculated as revenue per employee. For example, if the historical revenue per employee is $100,000 and aggregate distribution shows that there is a 50% chance of achieving between $2,700,000 and $3,800,000, then one interpretation could be that an operation manager would need to plan to have between 27-38 staff ready to deliver that business.

There are potentially many ways to utilize these forecasts. Likewise, how the other resourcing variables are adjusted in response will depend on the business and its resource supply options. In the case of the expected revenue shown in Figure 7, if managers were taking a cautious approach they might hire permanent staff to be able to fulfill $2,000,000 of business – taking care of the "small hump" – and begin interviewing and identify potential contractors to use in case the revenues looked like they would fall into the case of the "large hump" part of the curve. Alternatively, if mangers were bullish on future business or were concerned about the ability to acquire additional resources in a tight labor market, they might choose to hire to fulfill $3,500,000 to ensure that they can deliver in at least 75% of the cases.

One of the greatest benefits of the PSF system is being able to regenerate the forecasts quickly and adapt plans as events unfold. Better yet, managers can plan ahead to see how the revenue picture will change if they do unfold and take low-cost steps to help prepare for those outcomes. As events unfold, their overall effect can be better gauged. Relying on emotional instincts when good or bad events occur can be dangerous. When a big sale closes, there is a temptation to over hire, and not consider the potential failed initiatives that might also occur in the future. Likewise, when an expected sale or existing customer is lost, there is the protective reaction to cut costs, possibly to the detriment of other future business. A dynamic model can help reduce the speculative aspects and ensure that decisions are based on the aggregate situation rather than one-time events. To help smooth the effects of periodic revaluations of the forecasts, it could also be useful to adjust the resourcing based on a moving average, so as not to over react to short term changes.

5 Conclusion

By incorporating inexact and approximated information, it is possible to transform a simple, one-dimensional sales estimation technique into a rich and dynamic forecasting tool. Enhancing the forecasting model can help reduce dependencies on hard-to-measure mental models that are used by, and are unique to, individual sales managers. Using a formal model, forecast results can be easily shared, questioned, and further analyzed. Formal models are also more consistent and are less liable to physiological biases, such as sunk costs and the influences of interpersonal relationships. This approach is not expected to deliver precise financial forecasts, but rather to serve as an aid in operational planning. While there is a risk that garbage in may produce garbage out, quantifying the uncertainty of the forecast's inputs is an important step in improving the quality and consistency of the forecast over time. It allows the forecaster's assumptions to be more easily reviewed and questioned by others.

As with all predicative models of complex systems, this approach is not foolproof and should be combined with other information that is not captured within the model to make business decisions. While the PSF model is believed to be sufficiently general and applicable to many types of businesses, it is expected that many of the underlying parameters will need to be adjusted to match specific characteristics of different businesses. Likewise, the model parameters may need to be tuned over time, based on experience using the model and changes to the business' characteristics.

This project has identified a number of areas for extension and future research. One obvious modification would adapt the model to focus on profits instead of sales revenues. Another area of work would to be use more formal encoding techniques for capturing and modeling the uncertainty of sales based on individual's experience, instead of using a one-size-fits-all approach. Beyond optimizing resources, these tools could also be used to perform "what if" analysis on the effects of specific success or failures, and, in turn, determine the appropriate level of aggression when pricing deals, according to their potential impact on resource allocation. Finally, there are many adaptations that could be made to the model with respect to the probability distributions: non-normal distributions – such as triangular or Weibull – could be used as the basis of the variable inputs; alternatively, correlations could be introduced between different probability distributions to reflect interrelationships between different sale opportunities.

References

Hammond, J., Keeney, R., Raffia, H.: The Hidden Traps in Decision Making, Harvard Business Review, January 2006, pp. 118–126 (2006)

Hertz, D.: Risk Analysis in Capital Investment. Harvard Business Review 42, 95–106 (1964)

Hope, J., Fraser, R.: Beyond Budgeting. Harvard Business School Press, Boston (2003)

Koller, G.: Risk modeling for determining value and decision making. Chapman & Hall, CRC, Boca Raton (2000)

Mun, J.: Applied risk analysis: Moving beyond uncertainty in business. John Wiley & Sons, New Jersey (2004)

Savage, S.: Accounting for uncertainty. J Portfolio Management 29(1), 31–39 (2002)

Spetzler, C., von Stael Holstein, C.: Probability encoding in decision analysis. Management Science 22(3) (November 1975)