FORECASTING SHAREHOLDER VALUE: THE MISSING OBJECTIVE IN BALANCED SCORECARDS [1]

John Schuyler

PetroSkills® and Decision Precision®

ABSTRACT

Most for-profit corporations structure their strategy maps with a "long-term shareholder value" objective at the top. Balanced scorecard (BSC) designs then feature, typically, 15-25 related metrics. Individual metrics are any of: performance measure, value driver, and/or performance predictor. A few rare BSC implementations report historical shareholder return. All BSCs found by the author are otherwise absent the express objective. This paper describes an executive information system built around a stochastic model of the enterprise. The forecast of shareholder value generation is the focus metric and BSC centerpiece. This derives from forecasting *free cash flow* (*FCF*) aggressively obtained, converting to a distribution of *net present value* (*NPV*) and calculating the *expected monetary value* (*EMV*) (or, better, *certainty equivalent*), and then factoring *EMV* to obtain *market capitalization* (Market Cap = shares × share price). The enterprise model is the core means for evaluating and optimizing alternate corporate strategies and for measuring performance.

INTRODUCTION

Balanced scorecards (BSCs), popularized by Kaplan and Norton (2001) since the early 1990s, are widely applied for monitoring and measuring corporate performance and for communicating strategy. Decision analysis class participants have often asked my opinion, over the years, about BSCs. Until a recent shift in my thinking, my usual responses have been negative: BSCs report many criteria, and it seems that the user appears to be a multi-criteria decision-maker. I generally disapprove of multi-criteria decision making, believing that this is often muddled thinking. Why not measure value and progress in the

[1] Presented at the 2006 Crystal Ball User Conference, May 1-3, Denver, Colorado.

context of the organization, especially when the organization is a for-profit enterprise? Nonetheless, scorecard software proliferates as an interface to executive information systems. I endeavored to consider how BSCs might best be used — more properly, in my opinion — to support better decision making in the context of shareholder wealth creation.

On several occasions across a decade, I have written papers about decision policy for the Society of Petroleum Engineers' Hydrocarbon Economics and Evaluation Symposiums. Developing a demonstration BSC for the last paper (2005), furthered my research into modeling and measuring shareholder value. This current paper reports a more-advanced method for, and a demonstration of, measuring shareholder value creation at the corporate level.

I recommend shareholder value creation — rate, trend, and any direction change — as the showpiece of a BSC layout. I will present an enterprise model for the demonstration, this time modeling with Microsoft® Excel and Decisioneering®'s Crystal Ball®. With the companion OptQuest® tool, we are able to optimize management's levers.

Developing this paper has been part of a continuing effort to identify and understand best practice in corporate decision making. Among the improvements over previous models are:

- an enhanced demonstration of *free cash flow* (*FCF*) calculations and *market value discount factor* (*MVD*);
- a stochastic enterprise model generating expected values and mean forecast trajectories inside confidence bands; and
- a more complete and real-to-life demonstration of a BSC format that focuses on shareholder value.

Organization of this paper: Section 2 reviews some of the decision policy elements embracing the decision analysis approach. Section 3 discusses a high-level executive information system design with a centerpiece business model and an attached BSC. The key element of the BSC is the history and forecast of shareholder wealth creation. Section 4 shares some of the model-building experiences in this project that business model-builders may find interesting. One appendix expands upon the detail of risk-aversion, and a second appendix presents a one-page BSC example.

The three key ideas that I hope conference participants and later readers remember from this paper are:

1. The value of the company is based upon its ability to generate *FCF* available to the shareholders. Forecasting *FCF* is the basis for forecasting and measuring shareholder value creation.

2. A stochastic enterprise model is the foundation for both forecasting and performance measurement.

3. If long-term shareholder value is the objective, then the BSC centerpiece should be measuring shareholder wealth creation in order to align decision-making with shareholder interests.

DECISION POLICY BASED UPON SHAREHOLDER VALUE

Shareholders own the company, right? Most BSC strategy maps list "long-term shareholder value" at their apex. Long-term *investors* tend to focus on *value*, whereas *traders* tend to focus on *return*. A value orientation often conflicts with popular financial portfolio theory, which is most deals with returns. Common usages of *return* (or *yield*) are (a) simple gain fraction across a unit period, and (b) the *internal rate of return* (*IRR*). This paper will concentrate on shareholder value creation measured with money. Useful *supplemental* criteria — not part of formal decision policy — include *total return* to shareholders measured as an annualized *IRR*, and *return on capital employed* (*ROCE*).

Shareholders receive returns on their investment principally by two means: dividends and ultimate sale of their stock. In the U.S., dividends are taxed twice. The value of a for-profit enterprise derives from its ability to generate *FCF* available for the shareholders. My prior models (2003) demonstrated the shareholder value-maximizing strategy of a company repurchasing its shares rather than paying dividends.

Management has a recurring decision about what to do with *FCF*. My view of this is somewhat different than the customary finance definition. Most finance professionals assume sufficient money is reinvested in the company to maintain the business. But what if the business should not be maintained? I subtract only *mandatory investments*, for example, company maintenance projects that have such a high rate of return (e.g., above 15% post-tax) that the company would be foolish to pass up these. One benchmark *FCF* profile, then, is the amount of money that can be aggressively extracted from the company and paid to shareholders. The main alternatives for allocating *FCF* are reinvest in the enterprise, pay-down debt, and distribute to shareholders. The guide should be, "What is best for the shareholders?" Or, better, "What would shareholders want the company to do?"

Corporate growth is a fine thing, and it is hoped — often presumed — that this is good for shareholders also. I propose a recurring comparison, though, of at least two strategies: (1) a continuing-business case (maintaining and growing the company), and (2) an orderly business-liquidation case (seeking to accelerate and maximize cash distributed back to investors). The favored strategy — continuing or liquidation — should be the one that provides the greatest shareholder value.

How do we measure shareholder value? My understanding has evolved over the past 15 years. Fresh out of MBA school in the 1970s, I went to work as a planning and evaluation analyst for Cities Service Oil Company. The financial aspects of capital project evaluation seemed simple enough: Calculate *net present value* (*NPV*) discounting at a *weighted-average cost of capital* (*WACC*). We risk-adjusted for high-risk exploration projects and, thus, were then calculating an *expected value* (risked) *NPV* of *net cash flow* (*NCF*). Decision analysts call this *expected monetary value* (*EMV*).

My simple evaluation world fell apart in the early 1990s when a professor, in whose class I was guest lecturing, asked whether a risk-free rate ought to be used when discounting cashflows in Monte Carlo simulation. He referred me to the a highly-regarded textbook, *Principals of Corporate Finance* (Brealey and Myers, 6th edition and earlier, 2000). Could this be right? Brealey and Meyers said that when using Monte Carlo simulation, *NPV* should be calculated using a risk-free rate. A risk-free rate for most financial professionals means a treasury bill or government bond rate. Okay, this aligns with a popular idea: In decision analysis we are risking with probabilities and, therefore, ought not be risking with the discount rate. I've since been on a quest to understand what discount rate and other assumptions *should* be built into corporate decision policy. The key premise in my investigation has been this: The incremental value of a corporate capital investment should, when factored by the fraction ownership in a company, represent incremental value for the *Typical Shareholder*. That is, if shareholders (assumed homogeneous) could specify the corporate decision policy, they would apply their own preferences. I believe, therefore, that the Typical Shareholder's preferences about time value and risk attitude should be scaled up to the corporate level for decision policy.

I have long claimed in my teaching that if we do an evaluation properly, project *EMV* corresponds to incremental company value. However, reconciling *EMV* per share to stock prices is difficult without using a too-high PV discount rate. In preparing the 2005 BSC paper, I realized a straightforward solution. This wasn't anything new, and I'm surprised that it took me so long to seize upon the idea: Stock investors adjust *EMV*s (or *NPV* or some other value proxy) downward in determining market value. That is fair market value—market capitalization—is a fraction of *EMV*. This factoring method has long been the dominant risking method in evaluating collateral for corporate loans (which I did for six years as a petroleum evaluation engineer for a major US bank). Most buyers and sellers of cash-producing assets use a similar calculation.

Summarizing the key ideas of my current thinking:

- In evaluating capital investments, the stochastic project model should forecast company incremental *FCF*.
- Use the PV discount rate only to represent time preference for money. I believe the best discount rate is similar to a Typical Shareholder's home mortgage rate (after adjusting for tax effects). It's a risk-free rate, to be sure, but this is the shareholders' risk-free rate rather than the government's. This discount rate is much lower than typical and is less biased against long-term projects.
- The project's *EMV*, thus determined, represents incremental value to the company, on an *EMV* basis. However, *EMV* does not represent incremental company value in the marketplace.
- *Market value discount factor* (*MVD*) is what I call the factor to convert *EMV* to market capitalization. This solves a nagging problem of end-of-schedule-life terminal values: being a too-high multiple of ending cashflow rate compared to typical price/earnings ratios. Appendix A describes a further enhancement to the use of *MVD*.

Figure 1 illustrates the evaluation process. For a selection of petroleum exploration and production companies that I examined, this *MVD* was about 50%. Why the downward adjustment? Major reasons company investors discount *EMV*s are:

1. business cashflow-generation uncertainty
2. potential for poor management behavior (not managing for shareholders' best interests)
3. market (systematic) risk.

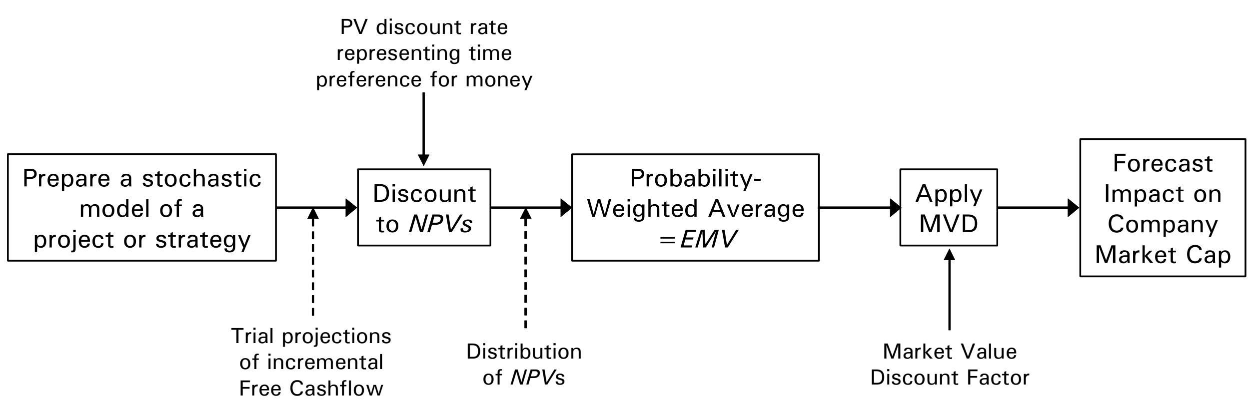

Figure 1: Evaluation Process.

BALANCED SCORECARD DEMONSTRATION

The purpose of the enterprise is to create value for shareholders. Therefore, shareholder value generation should be the centerpiece of the BSC. Figure 2 shows the suggested central chart and its companion. 80% confidence envelopes surround each heavy forecast line. Comparing a forecast (expected value) line to its envelope and median ("P50") line reveals the asymmetry of the forecast: the distribution at any time period is highly positively-skewed. In the lower chart, the Stock Price divided by *EMV* per share ratio is the market value discount (*MVD*) factor. From the chart, an executive can quickly see recent performance, the current forecast, and changing trends. A more-complete BSC layout is shown in Appendix B.

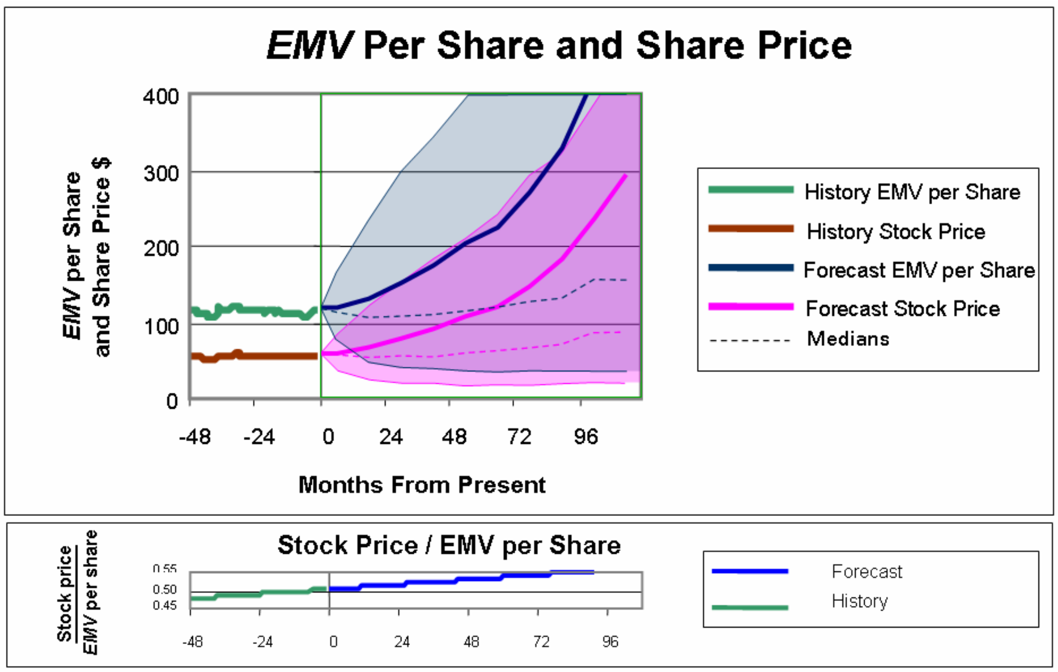

Figure 2: Centerpiece Charts in Balanced Scorecard.

"Are we creating or destroying shareholder value?" is the core question. Beyond the chart's obvious conclusions, we should have ways to determine causes and perspectives. Is the variance from plan caused by internal or external factors? How are we doing compared to our industry peer group? Against the broad stock market? With the enterprise model we can perform what-if analyses to answer such questions. For instance, we can replace actual product prices with the earlier planning price forecast; this price variance shows incremental shareholder value due to actual prices exceeding the mean forecast.

The demonstration model was built to represent a hypothetical oil exploration and production company. The company is gradually producing its existing petroleum reserves. It continues to invest in exploration by geology, geophysics, and wildcat drilling. This is analogous to R&D in other industries. Successful exploration testwells result in field development projects and incremental oil production. Analogous to a units-of-production depreciation method, petroleum accounting recognizes *depletion* as the gradual erosion of their capital investment value through production (as do other natural resource industries). If the company's outlook for exploration economics is unfavorable, then it stops exploration, the principal discretionary expenditure. Funds that would have been reinvested in the company can be used instead for stock buy-backs.

The heart of matter is generating a cashflow forecast with a stochastic model of the enterprise. Monty Carlo simulation allows uncertain input variables to be specified as probability distributions. Perhaps the most important reason for using Monte Carlos simulation is improved evaluation accuracy. Table 1 summarizes two strategies and two calculation methods. The simplistic *base case analysis* uses expected values for all input variables. A conventional (deterministic) discounted cashflow analysis indicates that liquidating the business is the better strategy: liquidation has the higher *NPV*. However, the deterministic model doesn't reflect the dynamics of the situation. Management has considerable flexibility in curtailing the business if conditions or performance worsen, and shareholders have limited downside. The truer value of the enterprise, in terms of *EMV*, is $43 billion. I call the calculation correction *stochastic variance* (*SV*), which is a component of a *variance analysis* explaining the difference between forecast and actual results.

Table 1: Comparing Deterministic vs. Stochastic Results.

	Continuing Business Case	Liquidation Case
Base Case (Deterministic)	*NPV* = $5,786 million	*NPV* = $15,178 million
Monte Carlo Simulation	*EMV* = $42,947million	*EMV* = $16,839 million
Stochastic Variance	*SV* = – $37,161 million	*SV* = –$1,661 million

I believe that good planning and control requires modeling. This model-centric approach applies for both the enterprise and for significant individual projects. Knowing where we've been is of less importance than where we're going: "It's hard to drive by looking in the rear-view mirror." Credible forecasting requires the model have good judgments and data going in. With modern information systems, this model can be updated nearly continuously. Figure 3 shows components of such an executive information system. Key decision variable optimizations can be run when needed and, perhaps routinely, overnight.

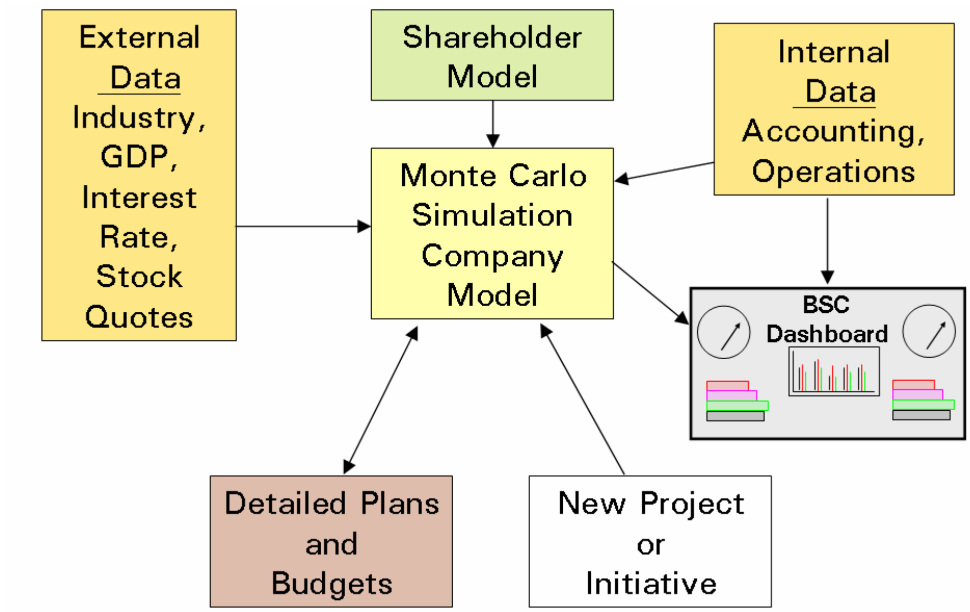

Figure 3: The Enterprise Model as the Centerpiece for Corporation Management

MODELING EXPERIENCE

Building even demonstration models seems to take longer than expected. Thinking about real-world cause-and-effect relationships — system dynamics — is always challenging and interesting. There is always one more feature to add. Using Excel as the modeling platform has well-known spreadsheet strengths and weaknesses. The chief advantages for this demonstration were the Excel charting capabilities (which proved problematic, at times) and in using Crystal Ball with OptQuest.

- The current model features the three types of Crystal Ball cells:
- Nine Assumptions: inflation rate, real price growth, discovery sizes, number of oil discoveries per exploration effort ($million current), and cost to develop discoveries.
- There are hundreds of additional chance events in the model, and these use binary and normal distribution types (using Excel's **`If`** and **`NormInv`** functions) sampled with Excel's **`Rand`** function. Rather than over-complicate the Crystal Ball environment, I used these simple methods for forecasting inflation, real price growth, and oil discoveries.
- Four Decision variables: fraction cashflow to maintenance projects and exploration, the fixed debt ratio, and a fixed cash reserve ratio.
- Two Forecasts: *NPV* for two strategies: continuing business and liquidating the business.

Additionally, there are about 25 single-value input parameters. Some were values that were assumed to be fixed. Others were variables having narrow distributions and/or having only a modest influence on outcome value. All input variables reside on an ‘Assumptions’ worksheet.

Figure 4 illustrates what happens to net cash flow from operations, which, after some strip-offs, leaves *FCF*. Exploration is discretionary and is added-back to the repurchases of stock for the liquidation case. Cashflow from operations and after paying taxes goes first to investments in workover projects (WOs) and new field development. A large portion of *FCF* goes to exploration in the continuing business case, and there is a minimal exploration level. A debt to *EMV* ratio is maintained (typically the maximum allowed by the bank) . Also, the bank cash balance is maintained as determine by a multiple of monthly average revenues. Any excess (shortfall) after maintaining target debt and cash balance is used to repurchase (reissue) shares in the company’s stock. This amount for stock repurchase plus exploration expenditures is the aggressive-extraction *FCF*. Regardless of strategy, the amounts applied to repurchase (or reissue) shares provides the return to shareholders.

I’m assuming that the company maintains a cushion of treasury stock that can be re-sold in the market when cash is tight. In real life, companies instead use their cash balances and bank short-term lending as the shock absorbers. However, I think this model’s approach more cleanly demonstrates what *FCF* is about and how it is the basis for shareholder value.

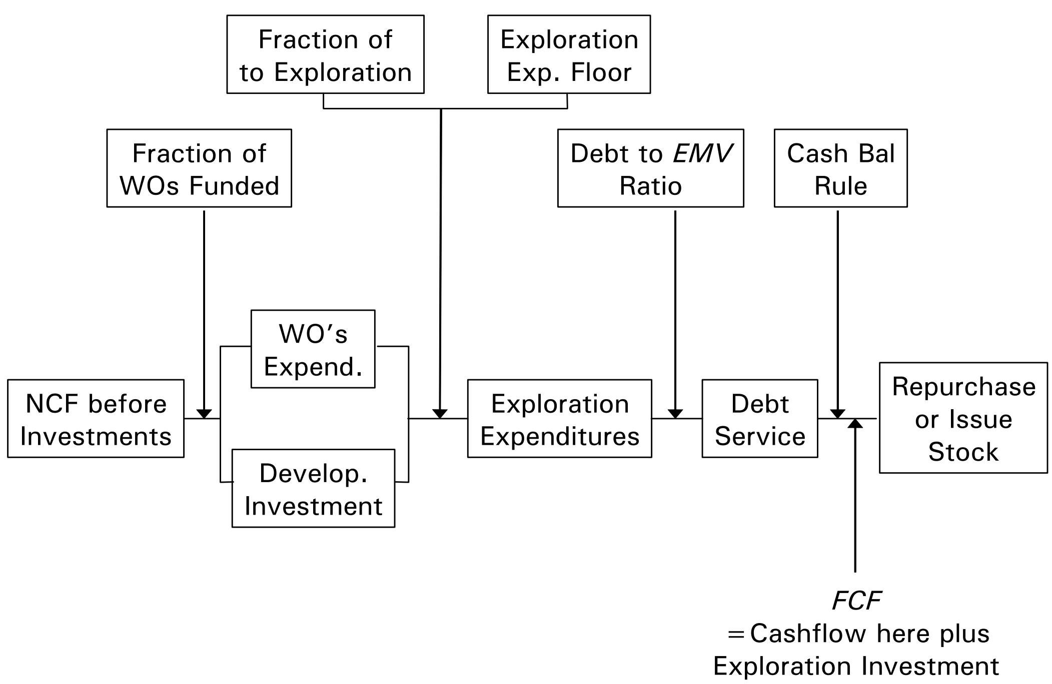

Figure 4: Net Cashflow Production and Allocation

Building schedules in monthly detail provides better timing for discoveries: discovery period, development across many months, and bringing the production on-line. I started the schedule one-year ahead of the chart starts so as to reduce some initial transients. The calculation schedules include about 150 columns x 180 months.

Checksums are almost essential for a model of moderate- to large-complexity. I included several balance checksums for cash, depreciation (depletion), and production. The conservation (of energy, mass) idea in science serves us well for physical quantities. Business uses the balancing analog in business accounting: debits equal credits. I became a better modeler after learning about double-entry bookkeeping.

A major design decision was whether to separate or combine the two main strategy models. I chose to model both strategies in parallel so that charting them together would be easier. Keeping these models synchronized was sometimes difficult through the many revisions. I built this model on-the-fly — once again, foolishly striving for expediency — rather than from a well-formed design and specifications. In the next stage of modeling effort, I resolve to rigorously test and document the model.

SUMMARY

Free cash flow (*FCF*) is something we can measure and design the information system to forecast. *FCF* is the basis for a company to have value. If there's no promise of *FCF*, there's nothing for the shareholder.

A stochastic enterprise model is the means for credible forecasting. We need the *FCF* forecast to see whether the outlook is improving or worsening. If reinvesting in the enterprise does not increase expected value present value *FCF* (i.e., expected monetary value, *EMV*), then *FCF* should be used to repurchase shares or pay down debt.

"Dashboard" software is typically synonymous with the BSC interface. If we agree that long-term shareholder value is the objective, then the BSC focus should be measuring shareholder wealth creation. This will help align decision-making with shareholder interests. The most useful dashboard element will be the timespread chart. Almost any metric, and especially *EMV* and share price, can be presented as a forecast with the historical trace.

I recommend using a *market value discount factor* (*MVD*) to explain and forecast the difference between *EMV* per share and share price. Advanced readers may wish to consider a further embellishment explained in Appendix A.

References

Brealey, Richard A. and Stewart C. Myers, 2000, *Principles of Corporate Finance, 6th ed.*, McGraw-Hill/Irwin.

[Note: A later edition appears to have rescinded the advice about using risk-free rates with Monte Carlo simulation.]

Kaplan, Robert S. and David P. Norton, 2001, *The Strategy-Focused Organization: How Balanced Scorecard Companies Thrive in the New Business Environment*, Harvard Business School Press.

Schuyler, John R., 2003, "Portfolio Management: What is the Contribution to Shareholder Value?," *Proceedings*, SPE Hydrocarbon Economics and Evaluation Symposium, April 5-8, Society of Petroleum Engineers, Dallas, SPE paper no. 82031.

Schuyler, John R., 2005, "Balanced Scorecard Linked to Shareholder Value: Making the Numbers Count," *Proceedings*, SPE Hydrocarbon Economics and Evaluation Symposium, April 4-5, 2005, Society of Petroleum Engineers, Dallas, SPE paper no. 94539.

Appendix A. ENHANCED COMPANY VALUATION METHOD

Two big issues in decision policy are how to best reflect unique (project) and systematic (market) risks in the value calculation. The modeling for this paper has been a stepping-stone in my investigation. This section adds some detail that I suspect will distract most readers from the main messages of this paper, hence its placement as an appendix. Nonetheless, for completeness, I'm including this brief discussion for readers who may be interested.

Discussions in the previous sections use the relationship:

$$\text{Market Capitalization} = EMV \times MVD$$

Most readers recognize and can relate well to *EMV*, and that's why I used this equation until now. However, there is a better value term for multiplying times *MVD* to get market capitalization: Use the *certainty equivalent* (*CE*) instead of *EMV*. This approach breaks out and will better represent the Typical Investor's risk aversion.

For an investor, the *CE* is the cash-in-hand equivalent of a risk. *CE* is the value of a risk to a conservative decision maker, where *EMV* is the value to a risk-neutral decision maker. For small, everyday decisions, these values are about the same. For determining incremental company market cap, the *CE* calculated at the company level would be the internal value of a project. Multiplying times the *MVD* provides the estimate of added company value (as market capitalization).

Risk attitude is often important in decision-making. These are situations where some of the potential outcomes are material with respect to the decision maker's wealth. An investor's risk policy is best represented by a utility function, such as shown as Figure 5. The utility function is used first to translate *NPV*s into utility units (which I label *risk-neutral dollars*, RN$). Calculating expected value with outcomes measured in utility produces the *expected (value) utility* (*EU*). The utility function inversion then translates the *EU* back into units of real money, the *certainty equivalent* (*CE*). This utility function, expressing a feeling about worth for different amounts of money, is a simple and elegant way to express risk attitude. This enables logical, consistent trade-offs between value and risk. In essence, we are making risk-attitude adjusted value calculations. While the utility function chart appears easy enough, we want to use the algebraic equations to get better resolution. Most decision analysts favor the exponential utility curve shape. There are three functionally-equivalent variations, and the particular form that I advocate is:

$$\mathrm{U}(x) = r\left(1 - e^{-x/\mathrm{r}}\right)$$

where *x* is an outcome *NPV*, and *r* is the *risk tolerance coefficient*. This *r* is merely a scaling factor and is typically on the order of 1/5 of an investor's net worth.

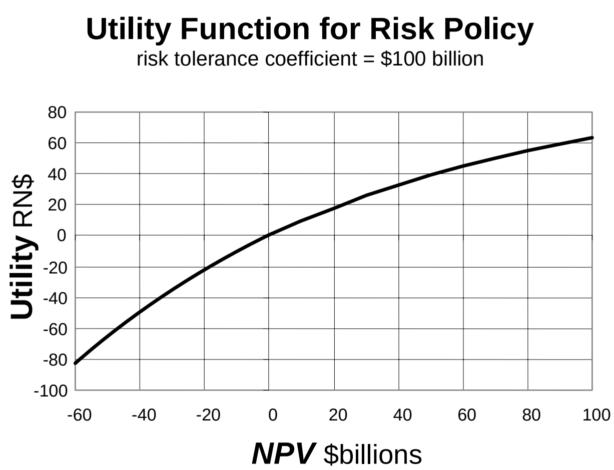

Figure 5: Example Utility Function

Using a decision tree, Monte Carlo simulation, or other method, an *expected utility* (*EU*) is calculated for an *NPV* distribution. The *expected utility decision rule* says the best alternative is the one having the highest *EU*. Since utility is in strange units, it's a good practice to transform the *EU* utility units into *CE* units of real dollars (or other currency). The inverse transform of the previous equation is:

$$CE = -r \ln(1 - EU / r)$$

The conservative Typical Investor's risk policy can be neatly expressed by a utility function. Everyone rational decision maker has one. If we profile a company's Typical Shareholder, we can then scale the investor's risk policy to the company level. For a widely-held company the *r* is huge, perhaps exceeding $100 billion. Day-to-day decisions, with outcome magnitudes well below *r*, won't see much difference between *CE* and *EMV*. For the corporation's incremental decision, $\Delta CE \cong \Delta EMV$. However, when valuing the enterprise as a whole, the considerable value uncertainty affects the shareholder's perception of company value.

Hypothetically, (1) if we have a company whose value is not correlated to the Typical Investor's other portfolio contents, and (2) the company's *r* is scaled up from the Typical Investor's *r*, then this equation will hold:

$$\begin{pmatrix}\text{Company}\\ CE\end{pmatrix} \times \begin{pmatrix}\text{Investor's}\\ \text{Share of}\\ \text{Company}\end{pmatrix} + \begin{pmatrix}CE \text{ of Other}\\ \text{Components in}\\ \text{Investor's Porfolio}\end{pmatrix} = \begin{pmatrix}CE \text{ of}\\ \text{Investor's}\\ \text{Entire Porfolio}\end{pmatrix}$$

Note that the *MVD* is absent this calculation. It's not measuring the market value of the holding. The *MVD* will reflect, principally, adjusting for the quality of cashflow information and systematic risk.

The usual case will be for individual investments to share systematic risk in the market, and this will somewhat reduce the *CE* of the aggregate portfolio. Nonetheless, embedding a conservative risk policy into the corporate decision policy recognizes and accounts for the most of the effect of investors' risk aversion. In some early testing, I'm finding that *MVD* is only slightly affected (reduced 10% or so) by high systematic (market) risk.

I believe this *CE* × *MVD* approach will provide more logical and consistent market cap estimations based upon free cashflow forecasts. Note that *MVD* depends upon the value measure choice. This embellishment contains a lot to digest and these details may detract from the simpler, main ideas in this paper. That's the reason for placing this section in an appendix.

I'm unsure where my investigation will lead when incorporating systematic risk in the company valuation. I'm confident that the above *CE* calculation will withstand scrutiny in representing investors' risk aversion. The *MVD* will contain a relationship to systematic risk, and I suspect this will not be a simple formula with the stock's beta term. We may need to characterize both a company's specific portfolio and the Typical Investor's portfolio. A project's systematic risk works through the company portfolio and on into the investor's portfolio, and it is in this context that systematic risk affects *MVD*.

Appendix B. EXAMPLE BALANCED SCORECARD

Still simplified from the usual 20-25 metrics, Figure 6 illustrates how rate versus time charts are well-suited for monitoring. Perhaps all the charts should include confidence curves around the best (mean) forecast lines, though this is shown only for the upper-left chart.

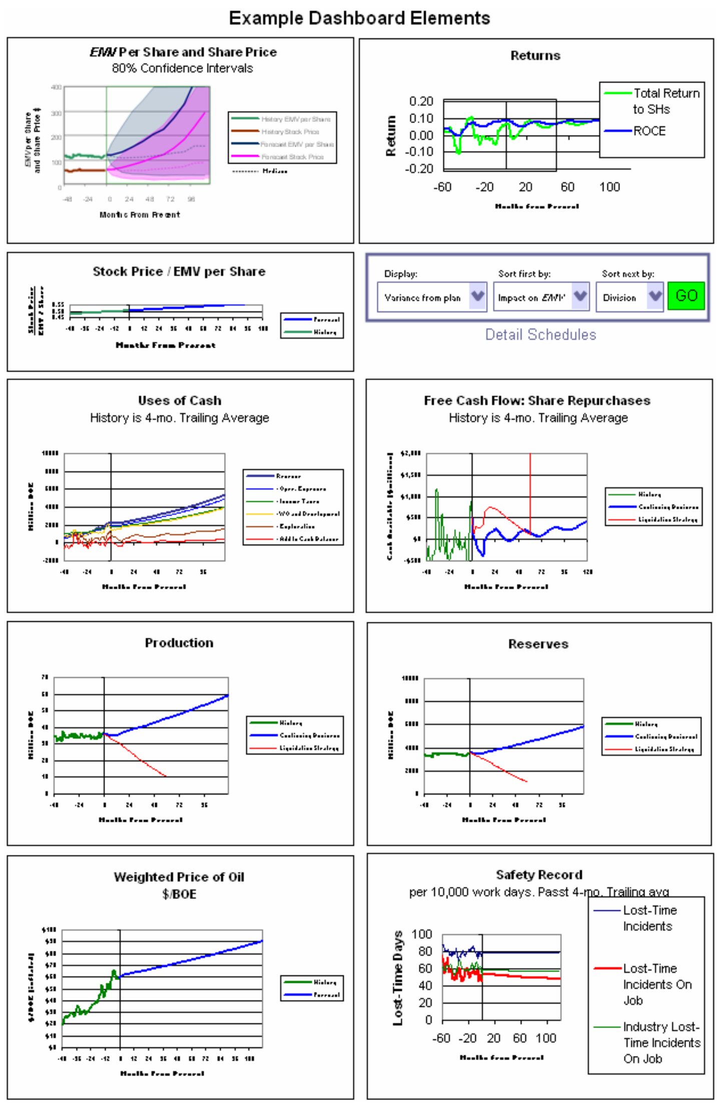

Figure 6: Prototype Balanced Scorecard Embracing the Concepts in this Paper.