

Managing trade-offs between conflicting goals through a portfolio visualization process

Stephen M. Rasey

ABSTRACT

Portfolio optimization processes help managers understand the costs of achieving performance goals and the trade-offs in the performance of one business measure versus other business performance measures. A good portfolio optimization process makes it possible to negotiate goals and constraints on important key performance measures interactively and collaboratively, at the same time being fully aware of the price being paid to achieve one goal at the expense other goals. This article advocates a gradient search process, instead of a traditional linear programming or mixed integer linear programming methods to build thousands of optimum portfolios from an inventory of investment opportunities. The performance levels of those portfolios are then analyzed interactively with a visualization tool to negotiate collaboratively the trade-offs between goals and resource levels for the corporation or business unit.

The portfolio visualization process described in this article begins by asking asset managers to define different operational strategies and acceptable performance exchange rates for various key performance measures. Combinations of operational strategies and performance exchange rate strategies each create a portfolio strategy and a gradient in a multidimensional resource space. For each portfolio strategy, the gradient search (or greedy algorithm) creates dozens of optimum portfolios, each at different resource levels and with different results in key performance measures. Thousands of optimal candidate portfolios are created and stored, and the coupled visualization tool enables them to be interactively and collaboratively filtered. This powerful combination of technologies allows asset managers to explore the effects of resource constraints and performance goals at specified levels of confidence on dozens

AUTHOR

STEPHEN M. RASEY ~ *WiserWays, LLC, P.O. Box 272462, Houston, Texas 77277-2462 and Custer Resources, Inc., P.O. Box 2740, Evergreen, Colorado 80439; raseysm@wiserways.com*

Dr. Stephen M. Rasey has a B.S. degree in geophysical engineering and a Ph.D. in mineral economics, both from the Colorado School of Mines. He has worked as team leader in risk analysis, strategic planning, and portfolio management for Getty Oil Co., Texaco, and ChevronTexaco accumulating exploration experience in 75 basins worldwide. For the past six years, Dr. Rasey has been with WiserWays, LLC and an associate with Custer Resources consulting on portfolio management and risk analysis process and software development for several international petroleum companies.

Figure 1. An inventory of Opportunities from which investment portfolios will be built.

Opportunity Name	Eval Cost	Prob Success	Net Present Value			Net Equivalent Reserves			
			Fail-ure	P90 Low	P50 Med	P10 High	P90 Low	P50 Med	P10 High
Frisch	20	0.102	-17	107	156	228	97	145	218
Kiner	16	0.449	-13	198	230	268	97	114	140
Barrow	12	0.083	-10	398	543	762	291	437	655
Killebrew	53	0.457	-45	42	63	89	107	128	154
Grove	58	0.475	-48	70	85	103	75	90	108
Harridge	80	0.220	-67	261	332	424	237	308	401
McGinnity	12	0.056	-9	160	411	851	335	586	1026
Roberts	27	0.125	-21	644	829	1105	249	374	561
Carew	50	0.209	-36	187	249	329	205	266	346
Cartwright	115	0.327	-94	383	444	518	244	293	351
Jackson	24	0.056	-17	1303	1861	2659	399	569	813
Flick	17	0.224	-13	42	60	86	23	33	47
McKechnie	42	0.094	-34	209	299	427	137	196	281
Appling	7	0.276	-5	101	145	207	50	72	103
Leonard	2	0.051	-2	22	31	44	18	26	36
Palmer	42	0.079	-33	955	1365	1949	792	1131	1616
Hoyt	3	0.432	-3	7	10	14	16	23	32
Rickey	3	0.075	-2	52	74	106	24	35	49

More Opportunities...

More Measures...

of key performance measures against thousands of potential portfolios generated by the gradient search. Managers can negotiate corporate and unit goals and budgets by knowing what portfolios are possible and what trade-offs exist between the unit's key performance measures. The visualization process allows managers to quickly select several portfolios they find equally acceptable and from those portfolios determine (1) which projects are common to all or most of the selected portfolios, indicating high priority; (2) which opportunities are members of only a few of the portfolios, indicating they are valuable but substitutable; and (3) which opportunities are absent from all the chosen portfolios and are thus low priority. The ability to quickly categorize the opportunities from a collection of superior portfolios gives managers high confidence in their investment decisions.

THE INVENTORY OF OPPORTUNITIES

Portfolio analysis begins with an inventory of investment opportunities. In an exploration program,

these opportunities can be described in a table, where each row is an investment opportunity, either currently owned or one that can be acquired. Each column describes the range of rewards (such as net present value (NPV), total liquid resources, total gas resources, projected earnings, etc.), the budget resources to acquire and evaluate the opportunity (cost of drilling, seismic data, land, and development cost if successful), the probability of success, and the range of potential outcomes for key performance measures. An example is given in Figure 1.

A manager seeking to build a robust portfolio of funded opportunities will typically need to satisfy goals in many key performance measures simultaneously, such as high NPV, low capital cost, low operating cost, high reserves additions, high production, high earnings, and high return on capital employed.

The trouble is that advancing one goal might advance progress toward some goals but may harm progress toward other goals. For instance, any reduction in capital expenditures (Capex) would help NPV if the same projects could be implemented. Reducing Capex might negatively impact reward

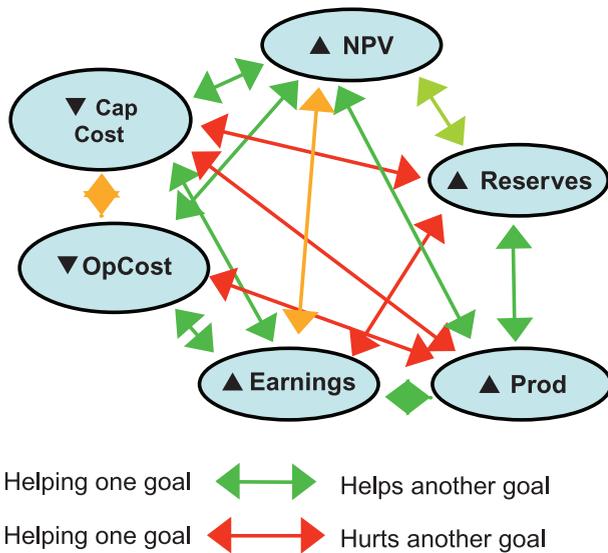


Figure 2. Actions taken to reach one key performance goal can make goals in other measures harder to achieve. NPV = net present value.

measures such as production and reserve adds and could require increased operating costs. Increasing reserve additions will lead to higher future production and future earnings but requires higher Capex, increased depreciation, and lower earnings in the short term. These conflicting goals are illustrated in Figure 2. This balancing act is complicated. Portfolio optimization models help navigate this web of interactions and find valuable courses of action.

COMMONLY USED MARKOWITZ-STYLE LINEAR PROGRAMMING PORTFOLIO OPTIMIZATION

A common method for portfolio optimization is to use a linear programming (LP) model derived from the work of Markowitz (1952) on asset diversification on portfolio returns. A general description of an LP model is to maximize or minimize some objective function subject to constraints on many combinations of variables (i.e., business goals, budgets, and execution resources).

In upstream petroleum investments, an optimization model can look something like the following for a business seeking to maximize NPV

and portfolio NPV, composed of the fully risked NPVs of its funded projects:

$$\text{Maximize } \$ = \sum \text{NPV}_i \times w_i$$

Subject to constraints on the following:

$$\sum \text{variance}_i \times w_i \leq \text{maximum allowed variance (risk);}$$

$$\sum \text{Capex}_{i,y} \times w_i \leq \text{maximum Capex}_y \text{ by year } y;$$

$$\sum \text{oil, gas reserves}_{i,y} \times w_i \geq \text{minimum oil, gas reserves}_y \text{ by year } y;$$

$$\sum \text{oil, gas prod}_{i,y} \times w_i \geq \text{minimum oil, gas prod}_y \text{ by year } y;$$

$$\sum \text{earnings}_{i,y} \times w_i \geq \text{minimum earnings}_y \text{ by year } y;$$

$$w_i \geq 0;$$

$$w_i \leq \text{some maximum working interest available for project } i \leq 1.0;$$

where Σ is a sum over all projects i from 1 to the number of projects n .

NPV_i is the fully risked expected value of net present value for project i and the decision variables w_i are the working interest to be used for project i . Many implementations of this kind of model are present. A more complicated variation allows projects to be scheduled by year.

What these LP models have in common is that managers must choose one goal or some linear mix of goals to maximize or minimize and then choose constraining limits on a large number of resource constraints and other performance goals. For each run of the model, for each set of constraints, one optimum portfolio is obtained consisting of a set of w_i variables. To check sensitivities on the solution, change one or more of the constraint limits and then rerun the model to get another portfolio of opportunity funding decisions.

This type of model treats maximum Capex available in each year as one of the known constraining limits. In practice, however, Capex is not a given; instead, it is a critical resource to be negotiated for in the strategic and tactical planning process. Even when the Capex level is agreed on, it frequently changes a short while later.

Another important caveat in the LP type of model is that the decision variables (w_i) or working interest by project are assumed to be continuous between zero and some maximum value. The

Markowitz model was developed for the stocks and bonds securities industry where w_i is the fraction of the portfolio value in opportunity i . The model assumes the cost and time to change w_i (i.e., the commission cost of buying and selling of shares) can be ignored. In the upstream petroleum business, the cost of changing working interest in a project is nontrivial. The working interest of exploration and production opportunities, when it is changeable, is commonly restricted to a few specific values offered.

To be more realistic, some portfolio optimization LP models restrict w_i to be non-continuous values representing a small number of possible working interests available in any project: 0 (not funded), 15%, 25%, etc. This variation requires a version of optimization called mixed integer linear programming, whose complexity is beyond the scope of this article.

DUAL VARIABLES: THE EXCHANGE RATES OF LINEAR PROGRAMMING

An underappreciated aspect of LP optimization is that, in addition to the optimum portfolio solution provided, LP also provides dual variables that show the marginal cost of each constraint on the objective function.

For every constraint there is a dual variable $\delta_j = \partial\$/\partial(\text{goal}_j)$, where $\partial\$/\partial(\text{goal}_j)$ is the partial derivative of the objective function (\$) with respect to a change in the value of goal j , the right-hand limit of constraint j . In other words, a dual variable is a performance exchange rate for how much the objective function can change for a change in the goal. Comparing the dual variables allows trade-offs to be assessed between different performance goals and resource limits.

Example of the Usefulness of Linear Programming Dual Variables

Suppose for an inventory of N projects, we want an optimum portfolio that maximizes NPV, subject to constraints on only eight goals: Reserve adds are greater than or equal to 500 each year for years 1, 2, 3, and 4 and earnings are greater than or equal

to 800, 900, 1000, and 1100, for years 1, 2, 3, and 4. The normal Markowitz type of LP optimization process to find the efficient frontier of portfolios is to run the LP under these constraints with an additional constraint of

$$\sum \text{variance}_i \times w_i \leq \text{maximum allowed variance}$$

For each level of risk, the Markowitz model produces one optimum portfolio. Using different values of maximum allowed risk, the Markowitz model generates different optimum portfolio solutions. Plotting these portfolios creates the well-known efficient frontier curve (Figure 3a).

Suppose the model is run for the same reserves and earnings goals but at different levels of risk. In Figure 3a, we show the six optimum portfolios; the fourth portfolio, at a risk of 1000 and an NPV of 2600 is highlighted in yellow. Figure 3b and c show results for two different performance measures (reserves and earnings) for the fourth portfolio for the four years with constraints plus an unconstrained fifth year. Examination of the reserves chart shows that, in year 4, the dual variable (blue bar in Figure 3) reaches \$4.50 per barrel oil equivalent (BOE), yet it is zero for year 1 and under \$1 per BOE for years 2 and 3. This indicates that the year 4 reserve constraint was a significant factor in the optimum solution. The LP optimization reduced NPV by \$4.50 for each BOE it needed to achieve the year 4 minimum constraint of 500 MMBOE reserve adds.

In year 1, the optimum solution was able to deliver 800 MMBOE in reserve adds when the constraint was greater than or equal to 500. This is a slack constraint, one that is more than satisfied by the solution. Dual variables for slack constraints are always zero. In other words, changing the constraint from greater than or equal to 500 to greater than or equal to 501 will not change the solution.

Using the LP model, we chose projects to meet the 500-MMBOE goal, and in doing so, we sacrificed NPV for more reserves in year 4 at the rate of \$4.50 NPV per year 4 BOE. That might or might not be an acceptable value to asset managers, but it is sufficiently anomalous compared to years 1, 2, and 3 that it calls into question the validity of the

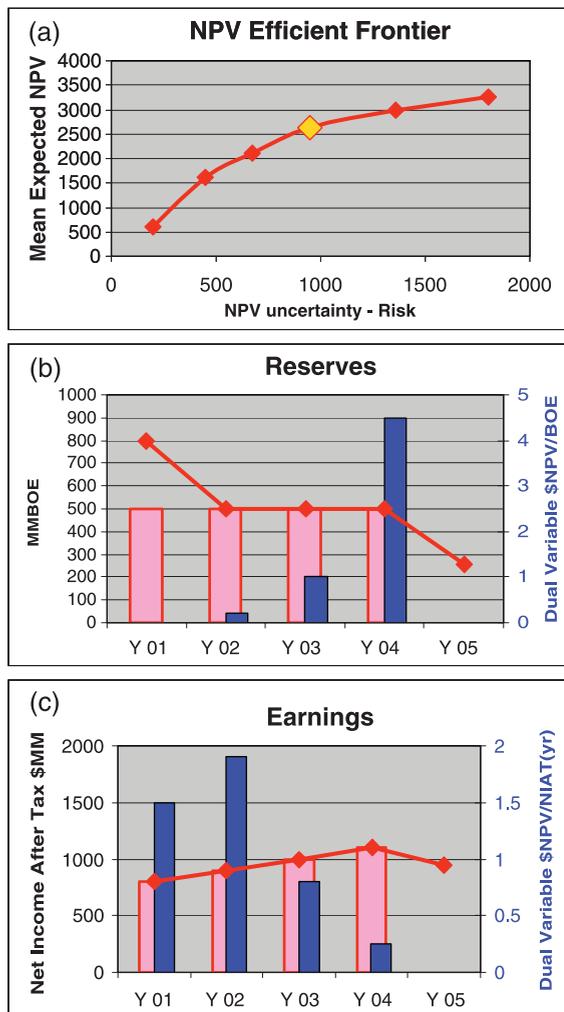


Figure 3. Results for the fourth portfolio (see yellow diamond in [a]) in the example problem. (b, c) The pink bars show the defined minimum reserve and earnings goals by year as specified by the eight constraints. The red lines show the expected mean reserves and mean earnings by year for the fourth portfolio point from the optimum solution. The blue bars show the values of the dual variables for that solution point. NPV = net present value; NIAT = net income after tax.

LP over the entire portfolio. Is it overemphasizing the relatively unimportant year 4 goal? Would a strategy to pay significantly more for reserves in later years than in previous years really make sense?

Likewise, in the earnings plot (Figure 3c), the optimum portfolio solution satisfies all four years of earnings goals, but the dual variables for the first two years are the most expensive. In year 2, \$1.90 of NPV was sacrificed to make that last marginal \$1.00 in earnings. Whether that is an appropriate trade-off is a judgment call for management, and it

is a caveat in the optimum portfolio of which management must be made aware.

Adjusting the Dual Variables

Suppose we decide that a portfolio solution that pays \$4.50 of NPV for a marginal year 4 BOE is unreasonable and is not part of a strategy we wish to pursue. Conversely, we recognize that we can get more year 2 resources quite cheaply in terms of NPV, only \$0.10/BOE (Figure 4). Is it possible to put in the LP model some way to constrain the values of duals variables? No, not directly. To reduce the value of the dual, which is the cost of the constraint, we must make the constraint easier to satisfy. In this case, we must reduce the minimum required mean reserves for year 4. If we want a value of \$2 NPV/BOE, how much do we reduce the greater than or equal to 500 MMBOE year 4 constraint? The answer requires trial and error.

If the LP model is rerun (Figure 5) (for each level of risk) with the year 4 reserves constraint greater than or equal to 350 instead of 500, then NPV rises by about 250 for the most levels of risk. In Figure 5b, the year 4 dual variable drops to 1.00, and the year 2 reserve dual drops to zero when constraint goes slack. The year 3 reserve dual variable increases from 1.00 to 2.75. That one change in the reserve constraint also changes the dual variables in the earnings. Now, years 1 and 2 earnings goals and year 3 reserve goals seem to be driving the optimum portfolio.

All the dual variables we have so far analyzed have been associated only with the portfolio point

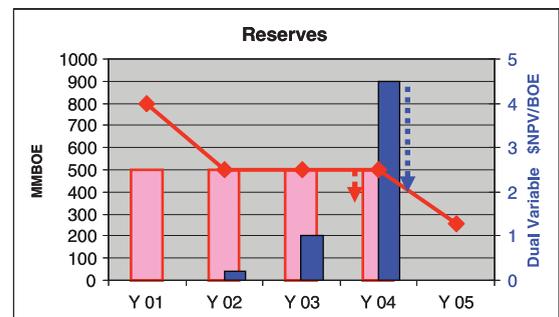


Figure 4. Can we find a solution where the dual variable for year 04 reserves is less than or equal to \$2.00 NPV/BOE? NPV = net present value.

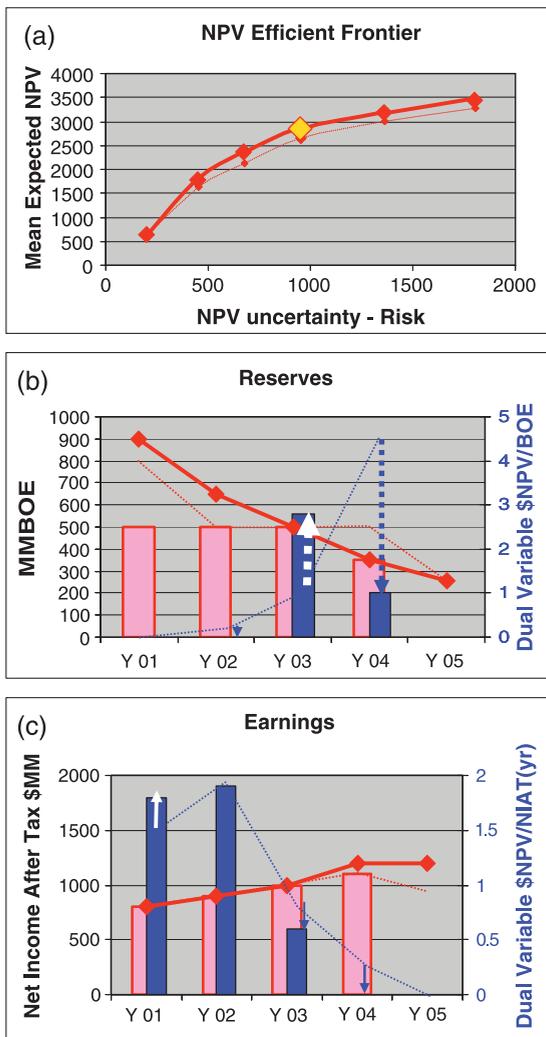


Figure 5. Relaxing the year 4 reserve constraint increases mean NPV at most levels of risk. All the dual variables for reserve and earnings adjust. NPV = net present value; NIAT = net income after tax.

at the NPV risk = 1000. At different portfolio risk points, you get different dual variables and different concerns about their limits and changes.

Making use of the dual variable of an LP portfolio optimization has value. However, controlling the dual variables in a typical Markowitz LP model is a tedious process.

BUILDING PORTFOLIOS USING GRADIENTS OF PERFORMANCE EXCHANGE RATES

A better way to employ the concept of dual variables as key performance exchange rates is to

use a gradient search technique, a concept derived from the field of capital budgeting and knapsack optimization. In some ways, knapsack techniques are better suited to managing oil and gas assets because they are formulated to use integer decision variables (1, 0 = fund, not fund) instead of the continuous w_i variables from LP. A simple optimization process, called the “greedy algorithm” (Bjorndal, 1995), is to rank all opportunities according to a formula, which is a ratio of (a mix of rewards)/(a mix of costs), sort them in descending order, and pick them in that order until the budget is used up. The key performance exchange rates (i.e., the LP dual variables) define the mix of rewards and costs. This simple process is guaranteed to be optimum for the exchange rates used (see Appendix 1).

If the projects were ranked by NPV/cost, then chosen in order, A, B, and C, then at the total cost of A + B + C, we would have the maximum possible (optimum) value for NPV and also generate some value for earnings and reserve adds (Figure 6).

Shortcomings of this simple rank ordering process include the following:

1. The point where the total cost of picked projects equals the total budget is indeed an optimum solution. However, if you run out of money between project n and $n+1$, then it is not an optimum solution and finding the best combination that fits the budget becomes more difficult and cannot be done by rank ordering.
2. Additional performance constraints, such as minimum reserves adds or earnings in years 1, 2, 3, and 4, cannot be added.

The first shortcoming can be avoided in practice. If optimum rank order of 20 projects requires 371 days of budget, why bother to find a worse combination of projects to fit into 365 days of budget? The budget and Capex approved will undoubtedly change over that time, anyway. Knowing what the best projects are for the next X dollars, and next 1.2X and 1.4X, is a good thing for asset management. In the oil and gas business, our knapsacks are pipelines in more ways than one.

The second shortcoming is more important. It is vital to pick projects for a portfolio that satisfy

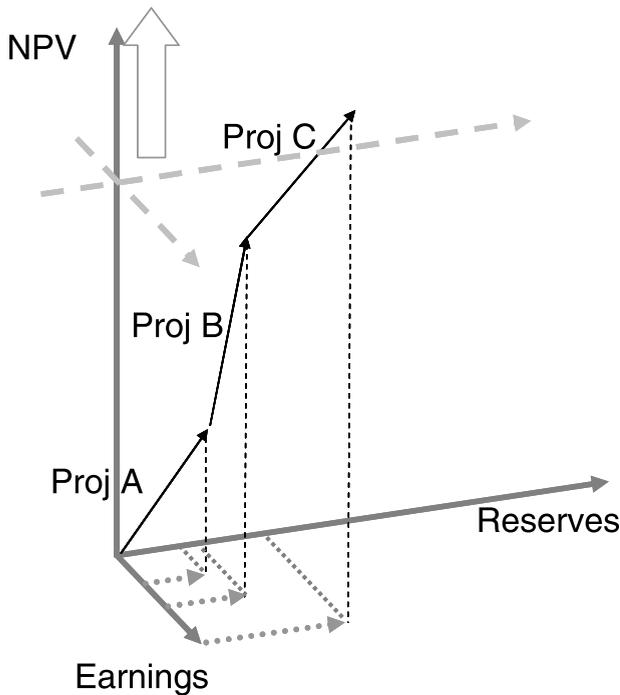


Figure 6. A three project portfolio to maximize NPV (net present value), where nothing else matters.

more than one performance goal. Portfolios must meet minimum performance requirements (at some specified confidence level) for earnings, reserve adds, NPV, production, return on capital employed, and other measures. We must also honor restrictions on various resources such as Capex, man power, rig-days, etc. By using ranges of performance exchange rates between the desired goals, in combination, many different gradients are created where one goal is traded for others. Using the gradient search on each of those gradients quickly creates thousands of portfolios whose performance exchange rates are defined and acceptable to management. The performance for each portfolio is then calculated and stored in the database. The visualization tool is used to filter out portfolios with unacceptable performance in one or more of the key performance measures. The visualization process vividly illustrates the trade-offs between those performance measures.

Figure 7 shows a modified version of Figure 6 using a rank order of projects not based on NPV alone, but on a gradient, a direction normal to the plane, whose definition is constant = NPV +

1.5 BOE + 0.8 earnings. This will likely yield a different set of projects from that in Figure 6, one that delivers more reserves, more earnings, and less NPV for the same total cost. It will be an optimum portfolio for the performance exchange rates (i.e., dual variables) of 1.5 NPV/BOE and 0.8 NPV/\$Earnings.

GRADIENT PORTFOLIO BUILDING PLUS VISUALIZATION EQUALS INTERACTIVE PORTFOLIO MANAGEMENT

We show several sample displays, which are visualizations of portfolios in a Spotfire DecisionSite[®] application. Each of these portfolios, a set of funded projects, is represented as a point in X,Y scatter graphs (the four on top and middle two on the right) or as a line on a profile plot (one of the bottom four charts), connecting one portfolio through many measures along the x axis.

This visualization tool allows interactive application of filters to the population of portfolios. Portfolios that pass the filters are selected portfolios and are in color. Deselected portfolios are those that fail one or more filters and are displayed in

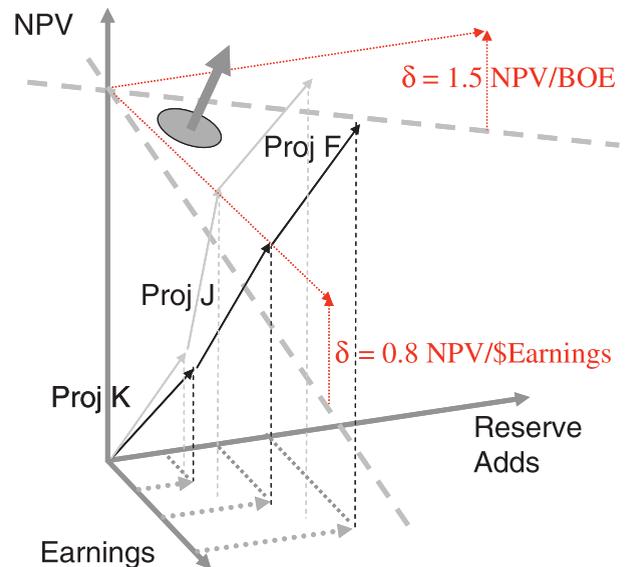


Figure 7. Three-project portfolio picked on a gradient, where 1 BOE = 1.5 NPV and \$1 Earnings = 0.8 NPV (net present value).

13,500 Portfolios displayed here, each represented by a point on linked X,Y charts, or a Line on Profile Charts

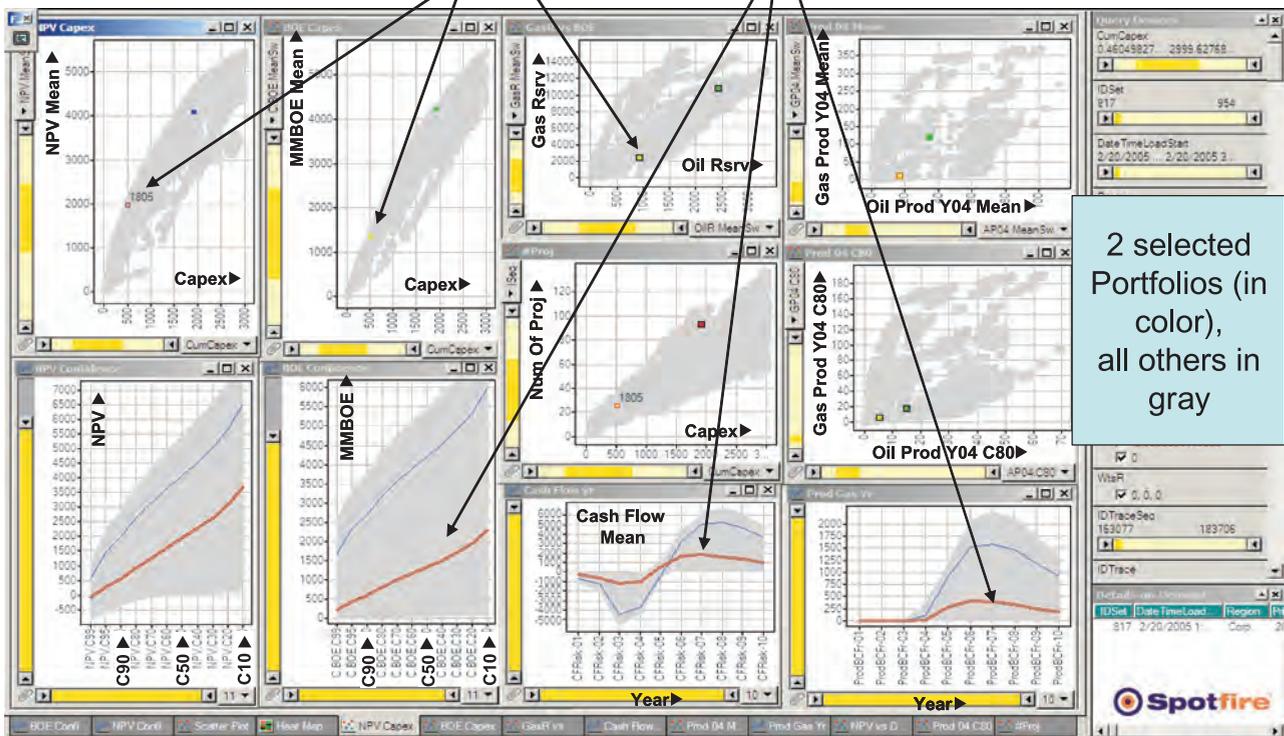


Figure 8. The 13,500 portfolios built on 90 strategies. Portfolios are points on some graphs, lines on others. A total of 52 portfolio measures in one screen. NPV = net present value.

gray in the background to give context to the selected portfolios. Marked portfolios are a subset of the selected portfolios that have been identified for further information or work and shown in yellow points or red lines. Selected and marked portfolios visible in one chart are visible in all other charts.

In Figure 8, the lower left profile plots are particularly useful for portfolio management and are not available in standard LP processes. These show the probability distribution of NPV and MMBOE for 11 percentile points along their distributions, starting with C99 (the 99% confidence level, also termed P99 low), through C95, C90 (P90 low), and so on, to C10 (P10 high). In this visualization approach, it is possible to calculate the true probability distribution of any measure from each candidate portfolio, then display as many points from the distribution as desired for analysis. It is not necessary to assume the distribution to be normal, log normal, or anything in particular. The visualization illustrates the true shape of distribution. With this

information, it is then simple to place a constraint (a visualization filter) at specific probability and performance levels such as requiring that the NPV 90% confidence level exceed zero: NPV C90 greater than 0.

Figure 9 is an example of crossplotting the NPV mean against the 99% confidence of NPV (NPV C99). This reveals the trade-off of between mean NPV and the potential downside (worst case) at the NPV C99. Ideally, the highest NPV means are preferable, but controlling the worst case (NPV C99) is also important. This plot shows that it is possible to have the highest mean and highest NPV C99 in the same portfolios (upper right of the selected portfolios).

Another advantage of this visualization process is that the whole number of funded projects can easily be displayed, tracked, analyzed, and constrained. The middle scatter plot from Figure 8 is the number of projects versus the total Capex. This is also shown in Figure 10a.

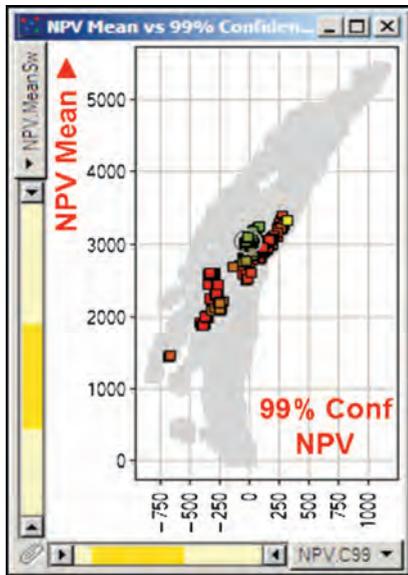


Figure 9. A crossplot of NPV (net present value) mean versus the NPV C99 for projects at a Capex level of 1100–1200.

The number of projects that can be successfully managed is a measure to consider in portfolio management. Diversification can lower risk, but once the number of projects invested in (even at small working interests) exceeds the ability to manage them properly, greater diversification will increase portfolio risk (see Appendix 2).

The gradient search process creates and stores portfolios quickly (~2 per second). Yet to have 10,000 portfolios ready for review means that portfolios must be created in advance using the strategies and gradients already approved by management. The portfolios are calculated and are stored in a database together with the values of the per-

formance metrics. These are then loaded into the visualization tool for analysis. In this workflow, portfolio measures of any level of complexity, and nonlinearity can be calculated ahead of time, stored, retrieved, plotted, and then analyzed interactively.

Figure 10b, shows a plot of mean NPV against the maximum negative mean expected cash flow of the portfolio. Portfolios delivering higher mean NPV with the smallest negative cumulative cash flow would be preferable to other portfolios. On this plot, the best portfolios are on the upper right edge of the cloud.

A PRACTICAL PROCESS FOR BUILDING STRATEGY-DIRECTED PORTFOLIOS USING GRADIENT SEARCH

Strategies are of two main types: operational strategies and value strategies.

Operational strategies are the everyday project management rules. Which projects must be done before others? What opportunities, in which plays and play types and in what countries, should we consider for the portfolio? How many projects can be implemented in a play in a given period? How much money can be invested in one country before political risk becomes unacceptable? How many wells can we expose ourselves to in one play at any time?

Value strategies are how we trade off one goal or resource constraint for another. What portfolio measures are most important to the company and

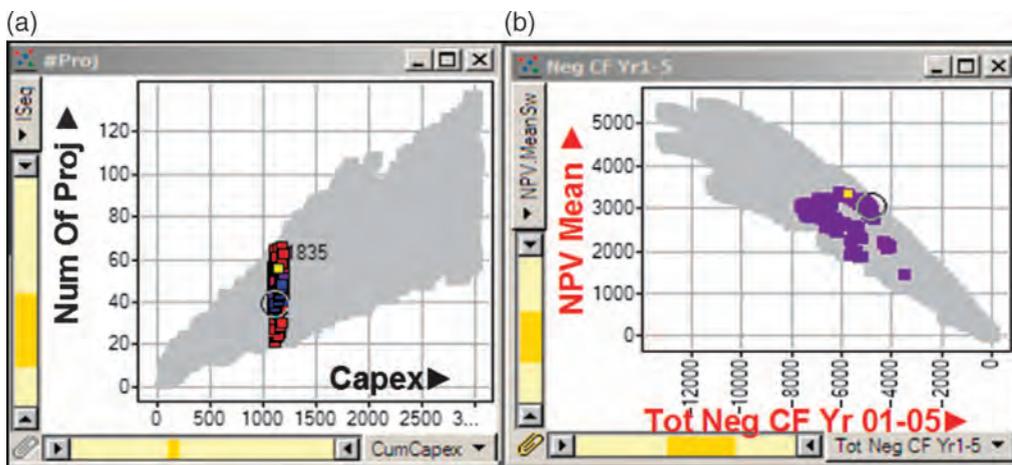


Figure 10. Two plots from with portfolio at 1100 to 1200 Capex. (a) Whole number of projects funded versus Capex. (b) NPV (net present value) mean versus cumulative negative cash flow (CF) over the first five years.

what performance exchange rate would we accept for a greater value in one portfolio measure and for a reduction in others? How much Capex should we spend for each barrel of reserve adds? Value strategies are the gradients made up from the performance exchange rates (the dual variables of the LP process) discussed previously in this article.

For each operational strategy, we combine all our value strategies. Each combination of operational strategy and value strategy will be called a portfolio strategy for the rest of this article.

In this combined gradient and visualization process advocated here, we quickly build portfolios based on each portfolio strategy set of operational rules and gradients. For each portfolio strategy, many portfolios are created. Using the greedy optimization technique, first, the single best project to fund is identified, thus giving a one-project portfolio (A). Then, the next best project is identified according to the greedy optimization technique; added to portfolio A, this results in a two-project portfolio B. portfolios A and B are built on the same strategy but have different levels of resource usage and different performance results. The process is repeated so that, for each portfolio strategy and for each portfolio of N funded projects, the next best project to fund is always present. If that next best project is added incrementally to the portfolio of N projects, this will give a portfolio with N + 1 projects. This is repeated until the limit of resources, time, or projects is reached.

This process may seem heuristic, yielding good though not optimum portfolios. The greedy algorithm, in fact, yields optimum portfolios at every step of the way. For a more detailed explanation, see Appendix 1.

Once the optimum portfolios are created, the values of the desired measures are calculated for each portfolio. These may be simple mean expected values, statistical variance, fully simulated probability distributions, convolutions, or other customized and complicated performance ratios such as return on capital employed. All calculated results are stored in a database along with the portfolio identifier, the funded projects (at the resource level at which they were picked), and the strategy that built each portfolio.

A new portfolio strategy is obtained by changing the value strategy or the operational strategy. A new sequential set of portfolios can then be built based on that strategy, the portfolio metrics calculated and saved to the database. This is repeated for all strategies to be tested.

This approach creates a huge number of optimum portfolios in a surprisingly short time. The greedy algorithm can create portfolio sets quickly. A software process we developed called BlitzPort (Rasey, 2005) is able to generate the portfolios, calculate all portfolio metrics, including high fidelity probability distributions via convolution on six key performance measures, and store the results to a database at the rate of 1 strategy and 150 portfolios per minute. Although some forethought is needed to define the strategies to test, it is possible to run 1000 strategies and 150,000 portfolios overnight. It is also possible to run 60 what-if strategies and 9000 portfolios over the lunch break of a portfolio meeting.

ANALYZING STRATEGIES AND PORTFOLIOS IN THE VISUALIZATION TOOL

Strategy A: Simply Maximize NPV

As an introduction to the concept of building portfolios with strategies, we will start with the simplest strategy: We can do any project at any time and we want to maximize mean expected NPV for the available Capex. This is a ranking of opportunities on bang for buck (Woolsey, 1975). We will call this “strategy A.”

Figure 11 is the Spotfire display showing the portfolios of strategy A to maximize mean expected risked NPV versus Capex used. The scatter chart at the upper left shows that at every level of Capex, no portfolios from other strategies (in gray) that give higher mean NPV are present. Each of the 150 portfolios calculated for strategy A are optimum for the conditions of that strategy. One portfolio, at about Capex = 500, NPV = 2000, is marked and shows as a yellow point on all scatter plots (Figure 11a–f) and a red line on the profile plots (Figure 11g–j). We can see from these plots that a maximize NPV

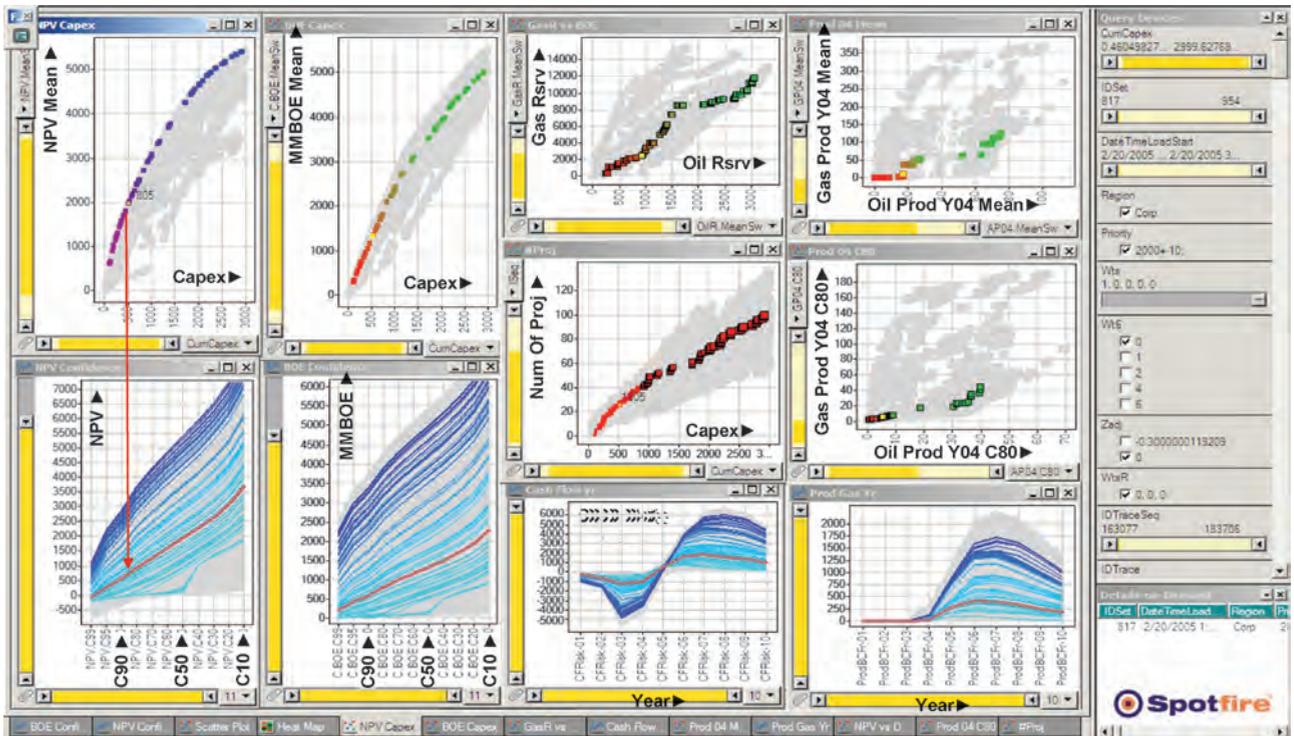


Figure 11. Portfolios from strategy A: Maximize NPV (net present value) versus Capex only. The charts are (top row, left to right) a, b, c, d, (middle row) e, f, and (bottom row) g, h, i, j.

strategy comes to about 90% maximum of MMBOE (Figure 11b). The production of gas in year 4 is very low (Figure 11d), especially in the marked portfolio. The C80 values (80% confidence levels, or P80 Low) for year 4 oil and gas production are both quite low.

Could we define a strategy to build portfolios that produce more gas for the same Capex? This could be obtained by applying a performance exchange rate where more gas production, or specifically gas production in year 4, is equated to an appropriate NPV.

Strategy B: Add a Gas Production Exchange Rate

Adding in the constraint to maximize NPV with an exchange rate of 1 mcf/yr = 3 NPV has a significant effect (Figure 12). The gas production for year 4 is near the largest found, the 80% confidence value of gas production is also high, but the trade-off was a lower mean NPV (Figure 12a) and lower mean MMBOE (Figure 12b), and NPV C10 dropped from 3600 to 1100. How much is that gas production worth? What are we willing to pay for

more gas production? That is a strategy question for management and that is why the gradient search uses the marginal price to pay for more of a performance measure as a performance exchange rate input variable defined by management and not an uncontrolled output variable.

Consider how this visualization display of thousands of portfolios compares to an LP optimization model that provides 1 portfolio at a time. If we had started with the constraints, sum Capex is less than or equal to 500, and sum year 4 expected gas production is greater than or equal to 200 mcf/yr, we would have found an optimum portfolio but would we have any idea of the price we paid in terms of NPV and total MMBOE to satisfy this constraint? If the LP uses for constraints exactly the same values for Capex and gas production year 4 as what was selected in the visualization, then according to the theorem of complementary slackness, we can prove that the portfolio found via LP and the portfolio generated via the gradient search and selected in visualization will be the same. Only if we looked at the dual variables of the solution and ran sensitivity cases. Yet, in the visualization process,

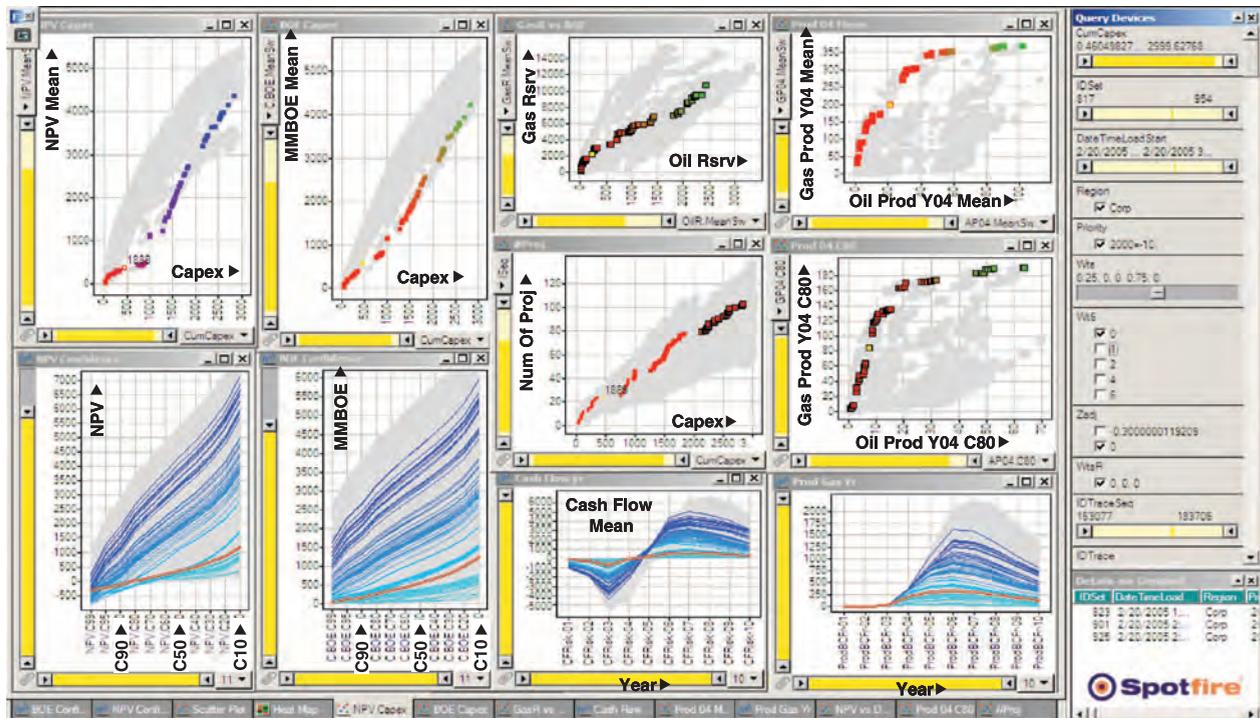


Figure 12. Portfolios from strategy B: Maximize (NPV + 3 mcf/yr in year 4). The charts are (top row, left to right) a, b, c, d, (middle row) e, f, and (bottom row) g, h, i, j. NPV = net present value.

the price paid in terms of NPV for 200 mcf/yr of year 4 mean expected gas production is immediately obvious to the decision makers.

Strategy C: Decrease the Gas Production Exchange Rate

Strategy C uses the gradient, where 1 extra mcf/yr of production is worth \$1 of NPV (or one extra bcf of production in year 4 is worth \$1 million in NPV). The strategy C portfolios are plotted together with the strategy B results (for comparison) in Figure 13. The new portfolio from the strategy C plot is higher on the NPV versus the Capex chart and MMBOE versus the Capex chart (Figure 13a, b); more oil production and less gas production than the portfolios from strategy B are expected, but still much higher gas production than the portfolios from strategy A shown in Figure 11.

Strategy D: Further Decrease the Gas Production Exchange Rate

Strategy D has portfolios built on a gradient, where we will give up only \$0.33 NPV for an additional

mcf/yr of production in year 4. Figure 14 shows the portfolios of strategy D added to those of strategies B and C. These portfolios are close to the maximum of NPV/Capex we found in strategy A. Only a little more gas production compared to the strategy A "NPV only" case (Figure 11) is present, much less than the other two strategies B and C, where we paid more for more gas production.

VIEWING MULTISTRATEGY PORTFOLIOS IN BANDS OF CAPEX

Figures 11–14 used a few selected strategies to illustrate how changing the value strategies (i.e., the gradient) changes the resulting optimum portfolios. This involved looking at four strategies at all levels of Capex. Although instructive, this is not the normal process for analyzing portfolios in practice. A more natural and pragmatic way to view the portfolios is to turn on a large selection of strategies, perhaps focusing on some common operational strategies, and view the portfolios that fit in bands of budgeted resource or Capex and goals.

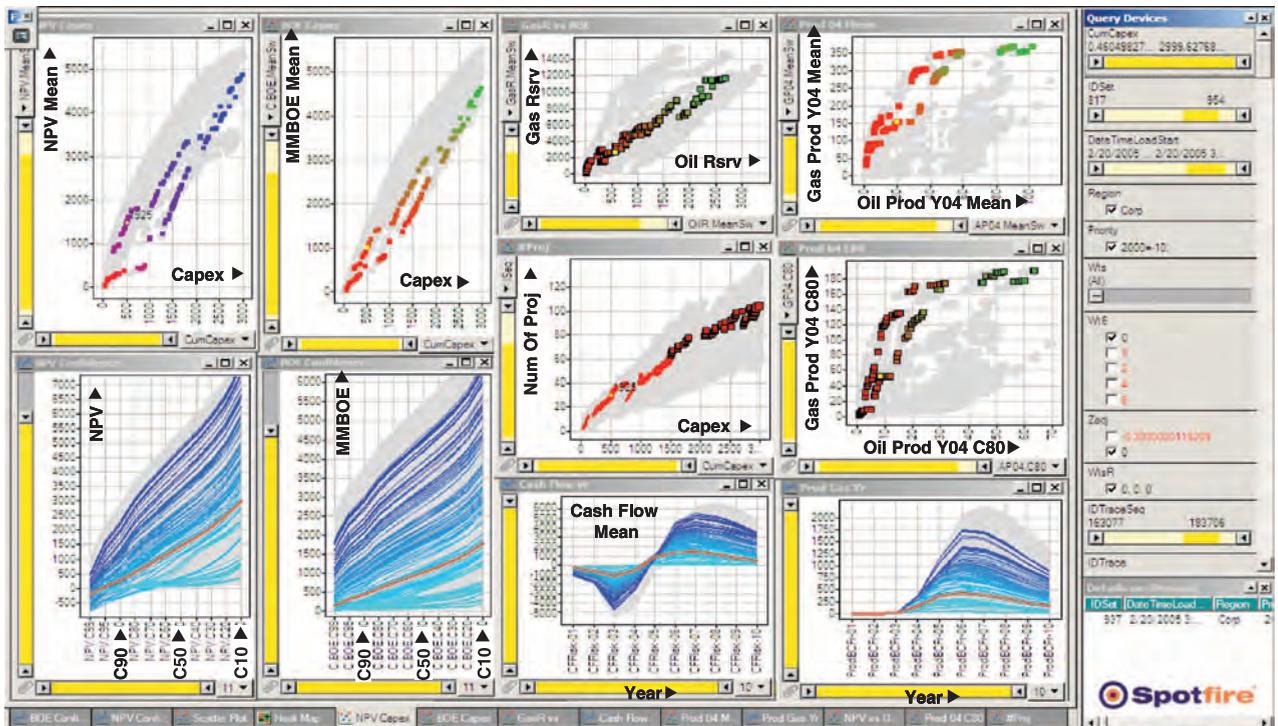


Figure 13. Portfolios from strategies B and C: (NPV + 3 mcf/yr 4 prod) and (NPV + 1 NPV/yr 4 prod). The charts are (top row, left to right) a, b, c, d, (middle row) e, f, and (bottom row) g, h, i, j. NPV = net present value.

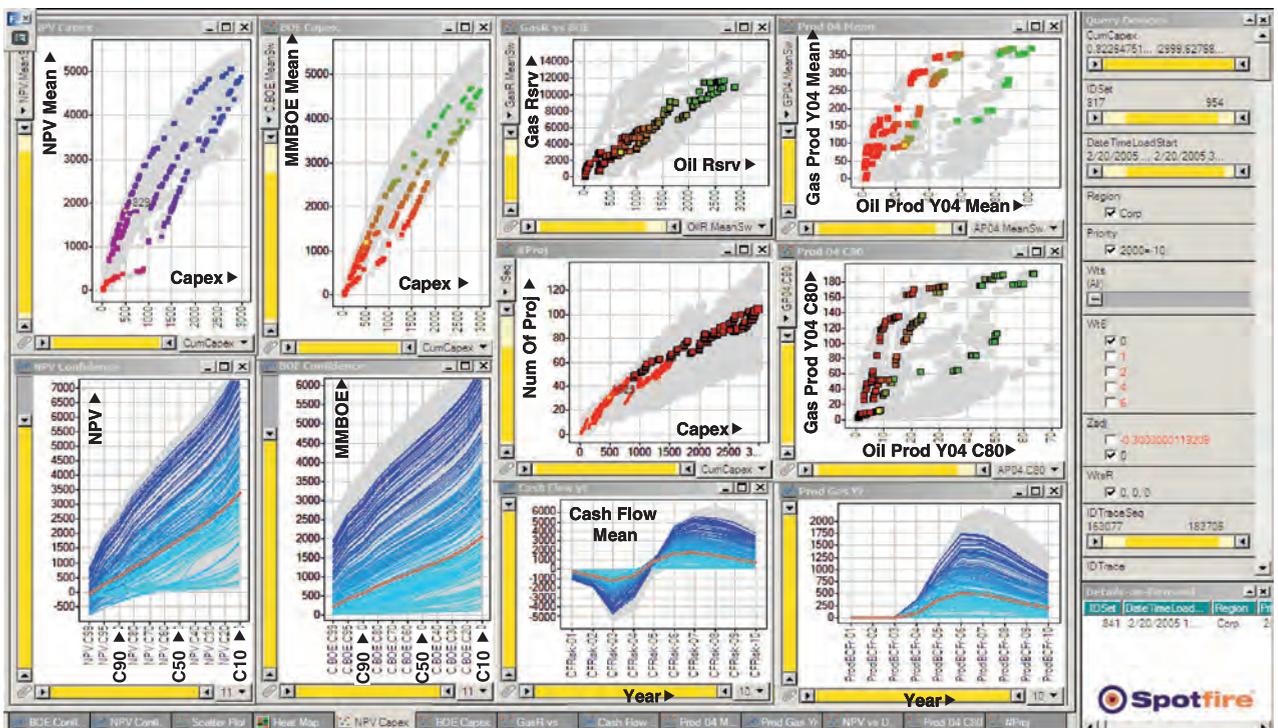


Figure 14. Portfolios from strategy D: (NPV + 0.33 mcf/yr 4 prod) added to Figure 13. The charts are (top row, left to right) a, b, c, d, (middle row) e, f, and (bottom row) g, h, i, j. NPV = net present value.

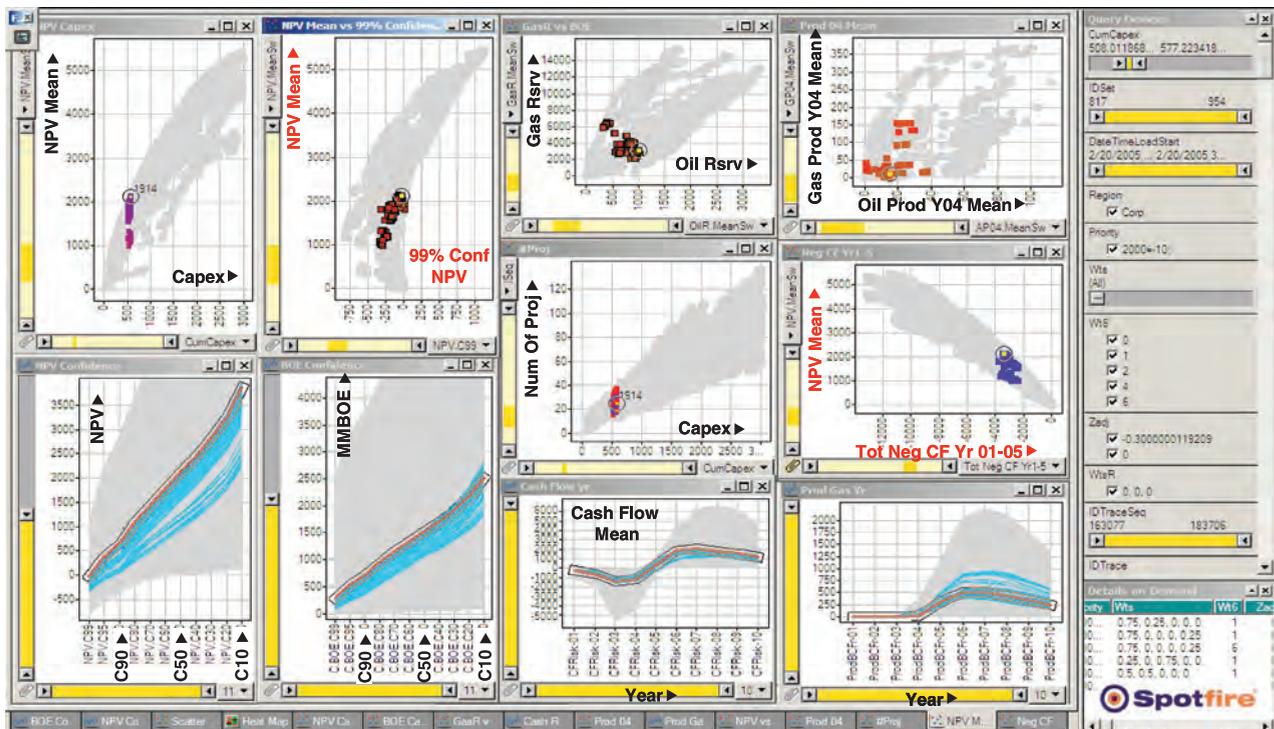


Figure 15. All portfolios, from any strategy, using between 500 and 600 Capex. The charts are (top row, left to right) a, b, c, d, (middle row) e, f, and (bottom row) g, h, i, j. NPV = net present value.

Portfolios with Capex between 500 and 600 from All Strategies

The visualization layout used is slightly different than in prior figures. Figure 15 includes a scatter plot of NPV mean versus the NPV C99, (99% confidence NPV, Figure 15b), a quantity also visible in Figure 15g. This chart, also shown in Figure 9, can show trade-offs between mean NPV and the worst case NPV. The best place to be is in the upper right part of the cloud (highest NPV with high NPV C99). Another new chart (Figure 15f) is the mean expected NPV versus the total negative cash flow over 5 yr; it was described in Figure 10.

For the portfolios selected here in Figure 15, one portfolio from strategy 1914 is highlighted with yellow points and red lines and shows the best combination of mean NPV, NPV C99, a low number of projects, and near the best edge for NPV versus the total negative cash flow, and it tends toward the oil-heavy side of the reserve mix (Figure 15c). Unusually, although this portfolio has one of the best NPV C99 results, it also has one of the best

NPV C10 (P10 high NPV) values; this portfolio (strategy 1914 and Capex ~570) will be referred to later.

Portfolios with Capex between 1000 and 1100

When the band of Capex is raised to between 1000 and 1100 for all strategies, this can be viewed live in the visualization tool and is seen by the motion of the clouds of portfolios as the Capex query device slider (upper right of Figure 16) is moved to higher or lower levels of Capex. This interactive process is difficult to capture in static snapshots such as Figure 16.

Within the given Capex band, portfolios that deliver between 1100 and 3200 in mean portfolio NPV are present. Portfolios with a NPV C99 (P99 low of NPV) from -550 to +270 are present. Portfolios with more than a 99% chance of positive NPV are present. Portfolios with the highest (mean expected) NPV of 2800 to 3200 will require a total negative Capex between -4000 and -7500, but it is clear from Figure 16f that -4000

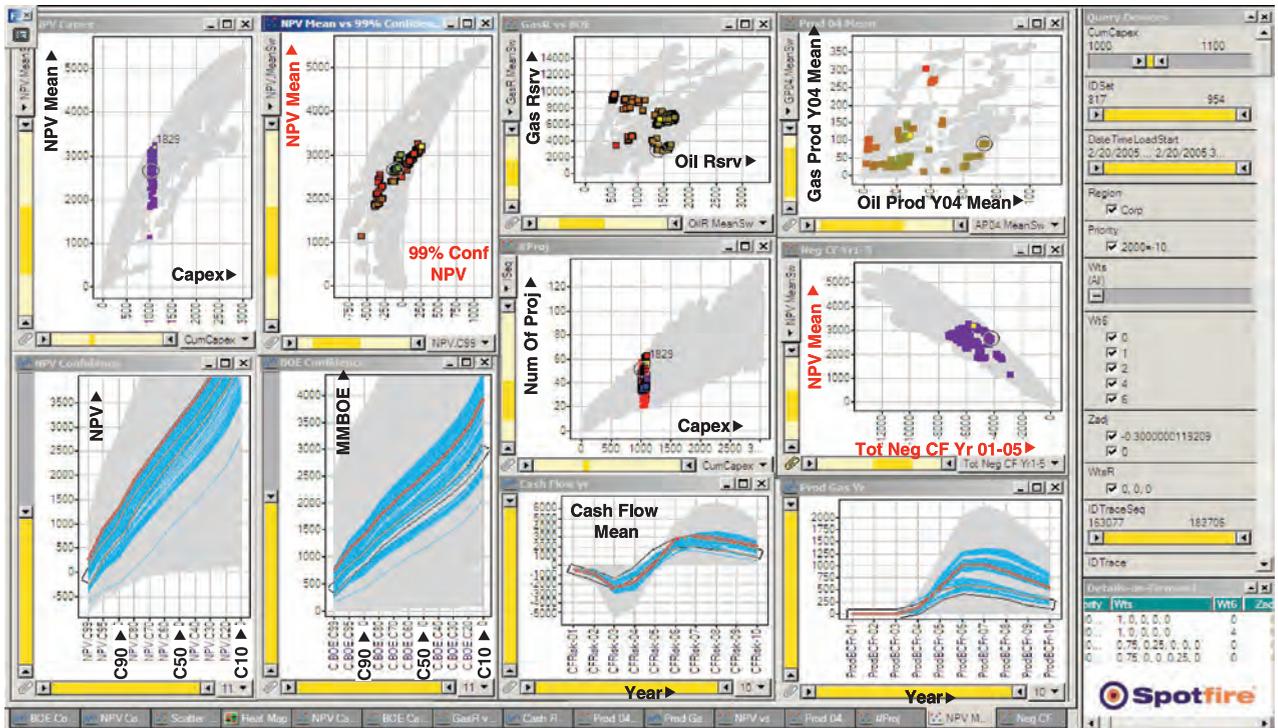


Figure 16. All portfolios, from any strategy, using between 1000 and 1100 Capex. The charts are (top row, left to right) a, b, c, d, (middle row) e, f, and (bottom row) g, h, i, j. NPV = net present value.

of the total negative cash flow is the best case if a mean NPV greater than 2500 is to be achieved.

The portfolio marked as a yellow point and red line in Figure 16 is from strategy 1829. It also has high mean NPV and high NPV C99, but it requires a comparatively high number of projects to manage, almost 60. In some companies, this could be a problem, and placing a filter on the number of projects may be considered.

Also visible in Figure 16 is a second kind of highlight, what Spotfire calls “brush linking,” where a circle around a point or a box around a line links the same portfolios in all charts. Here, in this chart, we brush linked the point with the maximum year 4 oil production (Figure 16d) to see where that portfolio falls against other portfolios.

PICKING THE BEST PORTFOLIOS FROM HUNDREDS OF STRATEGIES

The portfolios used in a visualization exercise can come from any source. They can be portfolios created one at a time in a LP or mixed integer linear

programming model. They can include portfolios submitted by each of the managers. However, the real power of the visualization tool is unleashed only if you have thousands of portfolios to sift through, exploring the full possibility space. Portfolio management using visualization requires the automated creation of thousands of portfolios, calculating and storing into a database the statistics on each. If managers prepare the ground with sensible operational and value strategies, then a computer can do all the work to create portfolios for each portfolio strategy and Capex level which are all optimal, feasible, and practical.

Finding the portfolios that best fit a variety of goals can now be done through visualization interactively, at a conference table with decisions makers negotiating goals and resource levels. Instant feedback is possible as you apply filters (constraints) to portfolio measures or lasso a group of portfolios with common attributes on one plot but wildly different results in other measures. Find the dozen portfolios that best fit your negotiated goals and you find the best strategies. This process is illustrated below by considering the example of how

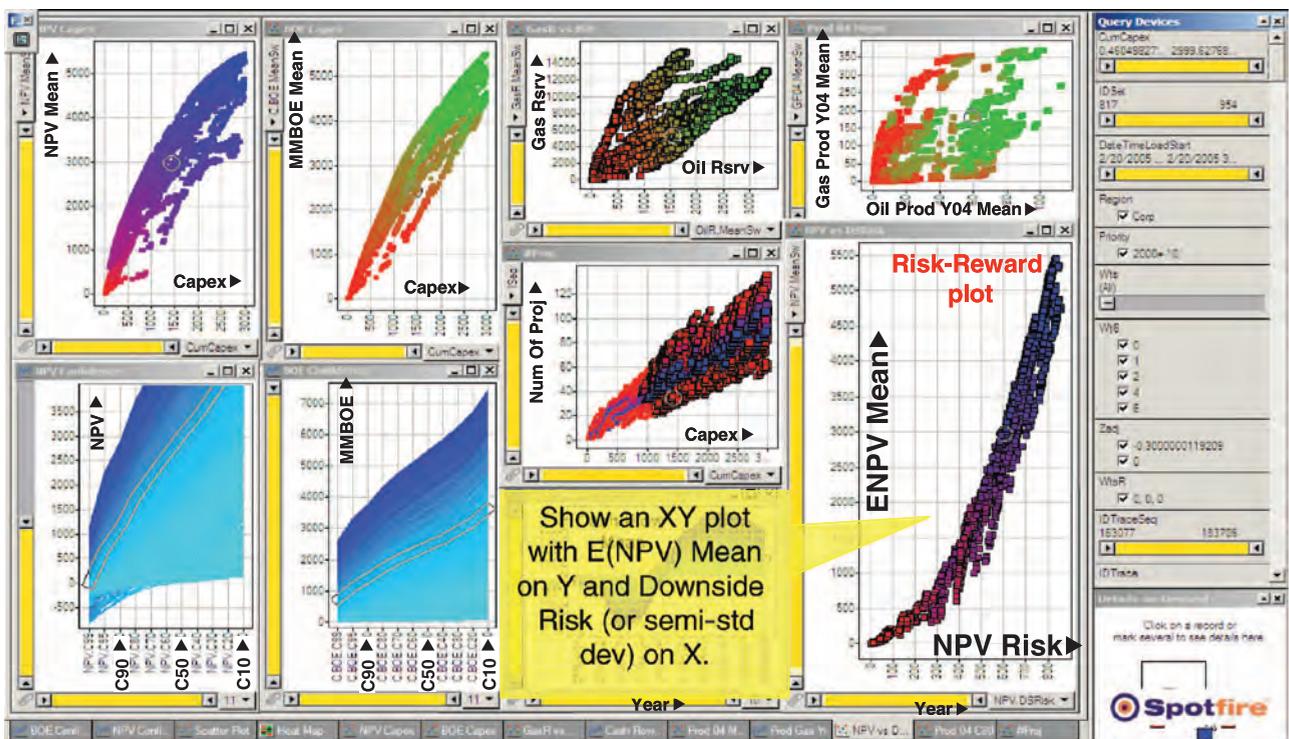


Figure 17. Initial state, with all portfolios visible. Lower left chart is a mean expected NPV (net present value) versus the NPV downside risk.

to select those portfolios with Capex less than 1500 and the highest NPV at a 99% confidence from a database of hundreds of strategies (Figure 17).

Step 1. Filter the Portfolios to Show Only Those with Capex Less Than 1500

Begin by picking a budget level the portfolios must honor. In the course of analysis, several budget levels may be investigated. Negotiating the budget levels is one of the key tasks in portfolio management, so it is vital to analyze more than one Capex level. In the visualization tool, the Capex limit is changed by moving a query device slider (Figure 17).

Step 2. Mark a Few Portfolios with the Highest NPV C99

Although we are interested in maximizing the overall NPV and production, we might first want to see what portfolios with the highest NPV C99 or NPV C90 values and what their probability distributions for NPV look like.

Step 3. Review Where the Marked Portfolios Appear in Other Portfolio Measures

The marked portfolios with high NPV C99 turn out to have high means and relatively high P10s (Figure 18). This is surprising given the mix of low and high risk of projects evident in the opportunity inventory seen in Figure 1.

DECIDING THE SMART OPPORTUNITIES TO FUND

Portfolio management is about deciding which projects to fund and having confidence in funding the right projects. Portfolio managers need to determine, not what is the single best portfolio to select, but which of the individual opportunities are best to fund, preferably under a variety of strategies.

Continuing the example described above, the next step is to use the visualization tool to filter down to a final 6–12 portfolios that best meet the negotiated goals. Figure 19 shows six portfolios with

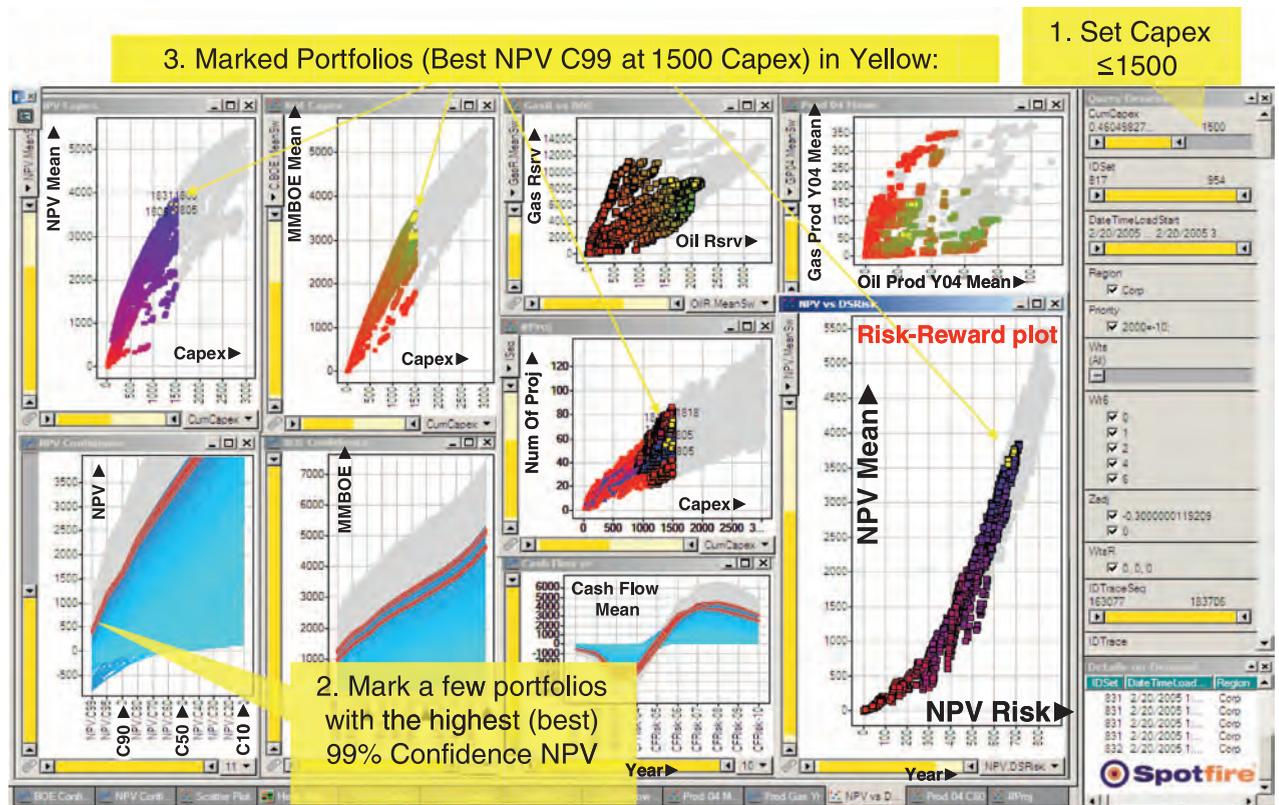


Figure 18. Figure 17 after applying filtering steps 1, 2, and 3 with a few portfolios marked. NPV = net present value.

a variety of goals from high NPV for the Capex used to delivering high production of oil or gas or both in year 4. All six portfolios, while delivering different NPV, plot on the efficient edge of expected NPV versus cumulative negative cash flow (chart at the lower right of Figure 19).

We then extract from the database a project census table (Figure 20) in which each row represents an individual investment opportunity and each column represents a marked portfolio point labeled with the strategy number. The Capex used and the mean expected NPV are also included (Figure 20). The intersection is the cumulative portfolio Capex or resource level, where that opportunity (row) is funded in that strategy (column). Lower numbers show projects picked early in the gradient search meaning these projects are most important, or most helpful, in reaching those goals with the available resources.

In the census table (Figure 20), the columns are ordered left to right in descending portfolio mean risked NPV. Each column is labeled with the strategy number that created the portfolio and

is the same as the labels on the visualization scatter points.

Strategies 1819, 1837, and 1833 all delivered high mean NPV for the Capex used. These strategies also sought out high reserves, but except for 1833, do not deliver much early production in year 4. Strategy 1931 is designed for high oil production in the fourth year. Strategy 1925 is designed for high gas production in year 4 with comparatively high oil production in year 4, but it had to give up NPV. Strategy 1829 delivers less gas production in year 4 than strategy 1925, has almost as much oil production, but improves on NPV. Strategy 1833 has only a little less gas production and oil production in year 4 than 1829 but has close to the best NPV for portfolios using less than or equal to 610 Capex.

A single optimum portfolio could be picked and its project list could be shown to decision makers as the answer. Instead, it is much more powerful to pick several portfolios that all have acceptable portfolio performance, then display the funded projects in a table such as Figure 20. It is self-evident

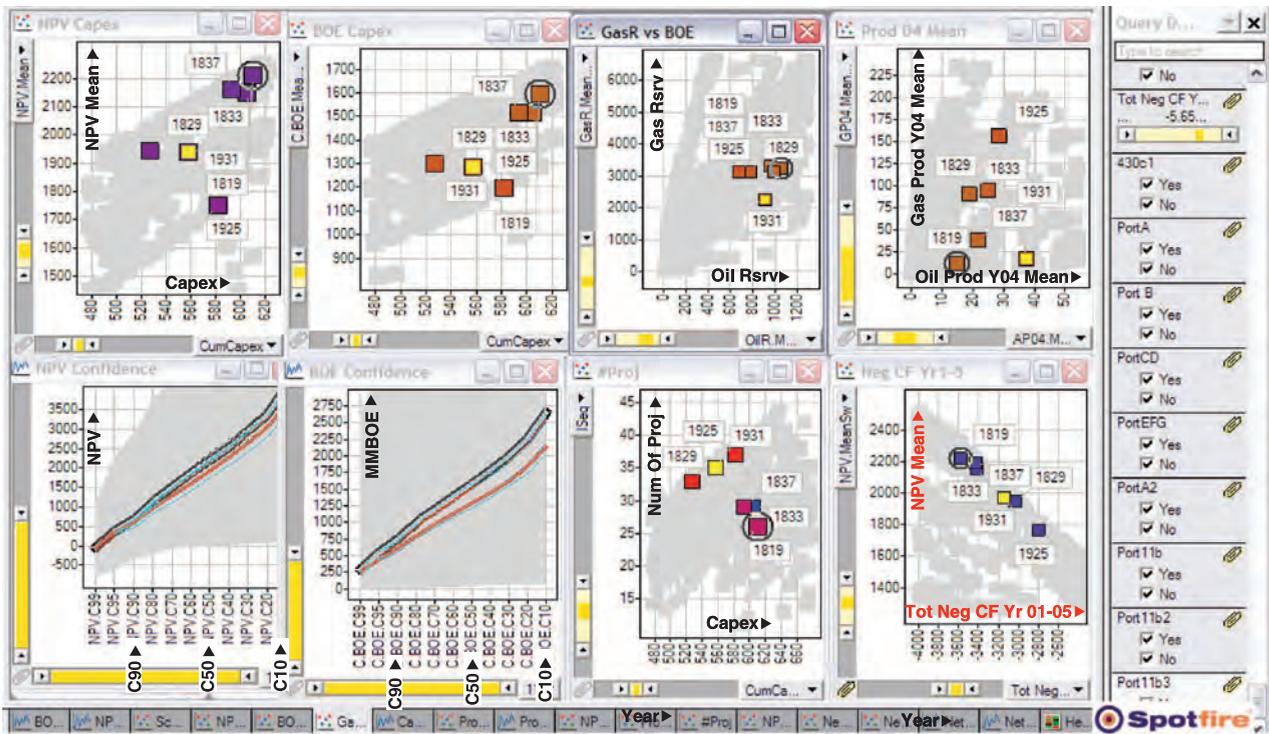


Figure 19. Eight key charts zoomed into a Capex range of 480–620. Each of six portfolios selected as finalists are shown with their strategy numbers. All portfolios visible. Lower right chart is a mean expected NPV (net present value) versus total negative cash flow (CF) for the first five years.

that projects funded in every marked portfolio are more important than projects funded in only one or two. Portfolio managers can quickly segregate the opportunities into groups:

- Group 1. Projects picked in every desired portfolio.
- Group 2. Projects picked in most portfolios (and which strategies picked them)
- Group 3. Projects picked in some marked portfolios, which we infer are substitution candidates, but are valuable only at the margin of the portfolio.
- Group 4. Projects never picked in desired portfolios, which thus look less attractive.

Knowing which projects fall into which groups greatly raises the confidence of decision makers. Group 1 should have the highest priority for funding. Group 2 projects would probably be worth funding but could be exchanged with a Group 3 projects for operational reasons. The projects in Group 4 might be targets for improvement, postponement, or possible divestment.

In the census table in Figure 20, 18 projects in group 1 are present, those funded in all six portfolios (Appling to Fingers). Group 1 prospects should not need much discussion because they are funded in all of the finalist portfolios. Perhaps the most important discussion is to double check that all the prospects in group 1 can be practically implemented in the same portfolio. If not, the operational strategies need to be changed and the portfolios reoptimized.

Group 2, coincidentally, also has 18 prospects; six prospects (Jackson to Morgan) are funded in 5 of the 6 portfolios. Jackson is not funded in the 1819 strategy, but is funded relatively early, as shown by the cell values of the table, in all the others strategies that put some value on the year 4 production. Cobb, Speaker, Clemente, Kaufax, and Morgan are all missing from strategy 1925 but funded in the other five. Strategy 1925 put a premium on year 4 gas and oil production and de-emphasized NPV. As a side note, Kaufax in strategy 1829 has a value of 558 in gray because the selected portfolio in strategy 1829 only uses 527 in

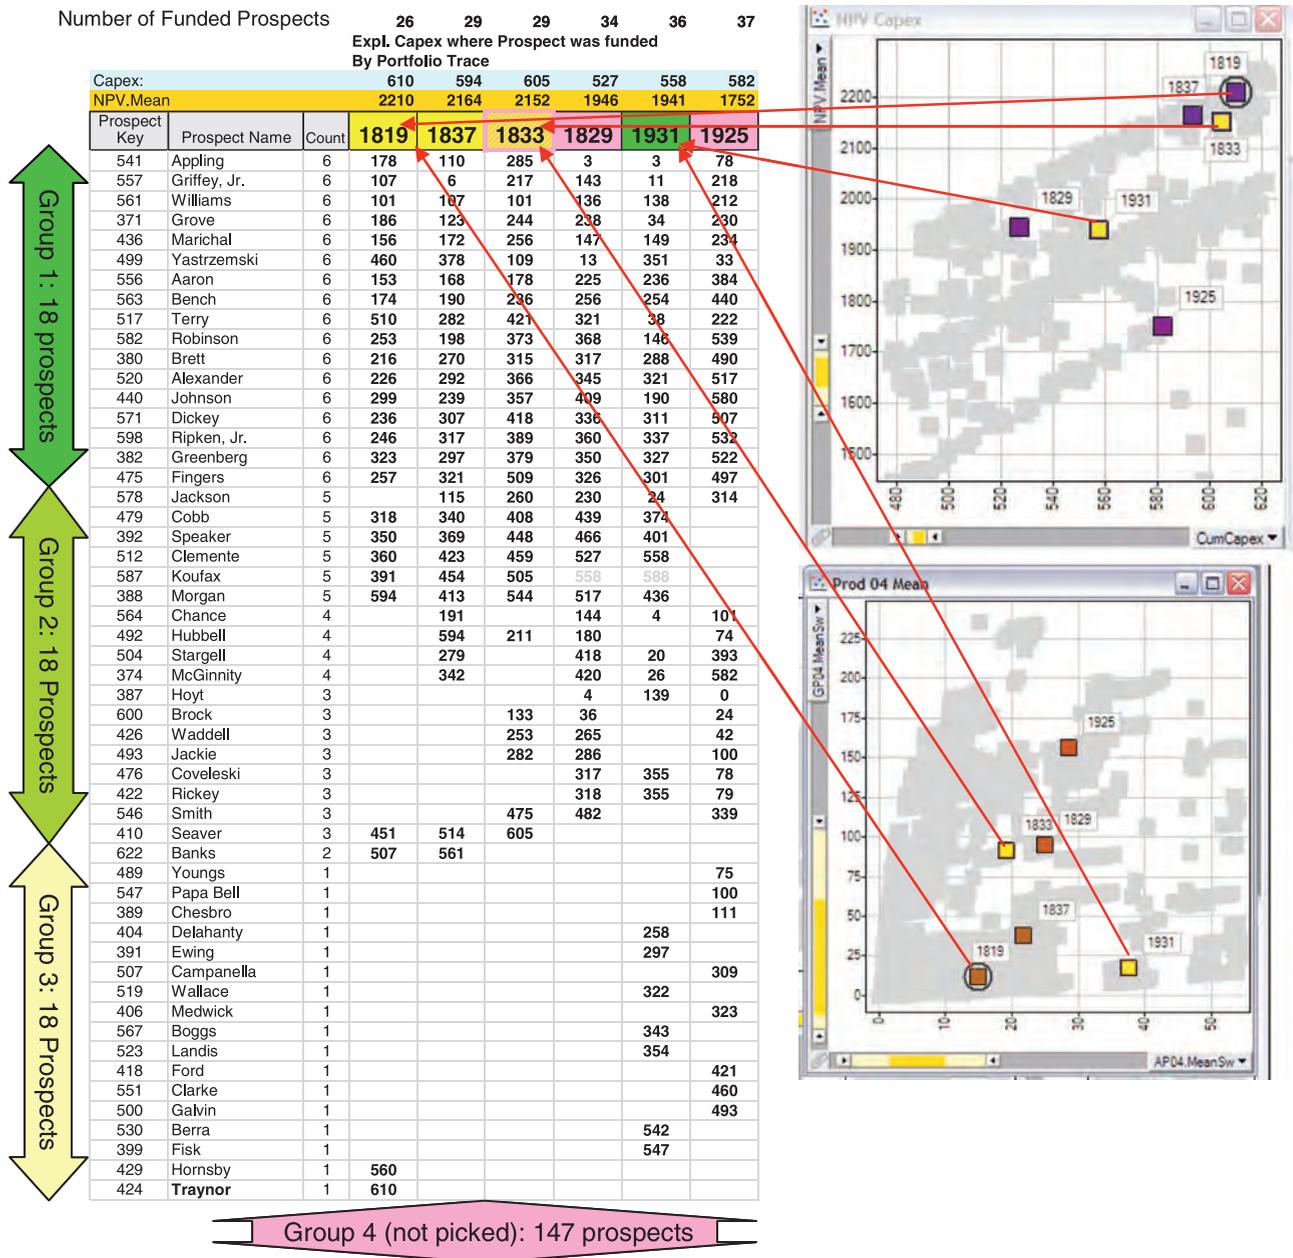


Figure 20. A project census table showing which opportunities (rows) are funded in which of the marked portfolios (columns), sorted with the most often picked opportunities at the top. The value in each cell is the cumulative Capex value at the point the project was funded in the portfolio; lower numbers indicate projects picked before higher numbers in the same portfolio.

Capex. However, because that portfolio is being compared to portfolios that use as much as 610, the table shows that according to strategy 1829, Kaufax would have been funded in that strategy for portfolios where Capex is greater than or equal to 558.

In the lower part of group 2 are the prospects Chance, Hubble, Stargell, and McGinnity that are funded in four out of the six portfolios, which

emphasized high year 4 production with some emphasis on NPV. Prospects Brock, Waddell, Jackie, and Smith are only picked in the gas production portfolios. Hoyt, Covelski, and Rickey are picked in portfolios that emphasize oil and gas production but in none of the high NPV portfolios. Seaver is picked in all three of the high NPV portfolios (and relatively late in the Capex allocation) but in none of those that deliver the higher year 4 production.

Group 3 also has 18 prospects, but these, with the exception of Banks, are each only funded in one of the six portfolios. Prospects young, Papa Bell, and Chesbro are all picked early in the strategy 1925 portfolio, which delivers the most year 4 production but the least NPV. Therefore, if near-term gas production is really important, these prospects would be important to fund. All other prospects in this group were funded in the latter half of the Capex allocation, indicating that these prospects add value to their portfolios at the margin. Delahanty, Ewing, Wallace, Boggs, Landis, Berra, and Fisk were funded in strategy 1931, a low NPV, high year 4 oil production portfolio.

What is left is group 4, the 147 of the original 200 prospects that did not make it into any of the six portfolios. Group 4 contains prospects that are unattractive for some reason, but it also contains some good prospects that may have not been funded in group 1, 2, or 3 because of operational strategy rules; for instance, too many wells in one play in one budget cycle. The next step in the analysis is to look at 3 or 4 yr worth of portfolios to see which of this group 4 are worthy of funding in later years. Those prospects not funded after 4 yr worth of portfolios may be considered for divestment.

CONCLUSION: USING VISUALIZATION AND GRADIENT SEARCH FOR PORTFOLIO OPTIMIZATION AND MANAGEMENT

A visualization approach to portfolio management is appropriate in a variety of real-world situations. visualization helps portfolio managers who want to understand the trade-offs between goals to better negotiate budgets and performance goals given their current opportunity inventory.

The combined visualization and gradient search approach is very useful in an environment where budget is subject to change. The approach strategy driven and always searches for the next best project to do with the next amount of Capex. Being strategy driven, the approach is superior to testing strategies than are goal-driven LP models.

Finally, this process is well suited to working with situations where many acceptable portfolios

are present and for looking at the funding decisions they have in common. Knowing which opportunities are funded in most of the desirable portfolios and which opportunities are funded, only sometimes give the decision maker a higher level of confidence in deciding which opportunities to fund, analyze, and debate further, which to shelve, and which opportunities to monetize.

APPENDIX 1: OPTIMUM PORTFOLIOS FROM THE GREEDY ALGORITHM AND COMPARISONS TO LP OPTIMUM PORTFOLIOS

With an LP optimization model, we specify the goals as constraints, and let the dual variables vary. Diligent investigators will interrogate the output values of the dual variables to determine which performance exchange rates are acceptable to management and whether any are unacceptable. If any are unacceptable, a trial and error process is used to change one or more of the goal or resource constraints and rerun the model to find a new optimum portfolio and to find a set of dual variables, where all have values acceptable to management.

In the gradient search technique, we specify a range of values for performance exchange rates (i.e., LP dual variables) and run combinations of those dual variable values as part of our portfolio strategies. For each strategy, we use the greedy algorithm to sequentially select the next best project to add to each preceding portfolio generating portfolios with unknown performance and resource usage levels, but all using acceptable levels of performance exchange rates. Through the visualization process, we interactively and collaboratively in real time filter out portfolios that have unacceptable performance or resource usage levels. When finished, we are left with portfolios that have both acceptable absolute performance results and acceptable performance exchange rates. No recycle is necessary unless the all combinations of the ranges of performance exchange rates approved by management are incapable of creating portfolios with overall performance levels acceptable to management.

Are the portfolios created by two different optimal processes the same? If we were to take one portfolio generated by the gradient approach and plug its performance results into the constraints in an equivalent LP model, the theorem of complementary slackness (Woolsey 1975; Anstee, 2006) insures that we get the same optimal portfolio from the LP model.

APPENDIX 2: WHAT HAPPENED TO THE MARKOWITZ CONVEX EFFICIENT FRONTIER?

The risk-reward plot in the lower right of Figure 18 is similar to the Markowitz risk-reward plot but differs from the classical Markowitz efficient frontier in that it is concave upward instead of convex upward. The Markowitz models assume

that you can diversify and reduce overall risk by funding more projects at reduced working interest at will and at no cost.

The Markowitz efficient frontier gets its convex upward shape by doing more projects for the same budget on the assumption that doing more projects at smaller levels of investment mathematically reduces risk. More projects, if independent of each other, will reduce the statistical standard deviation of the uncertain result and reduce the downside risk of the portfolio. However, even if projects are independent of each other, it means doing more projects at lower working interests. That forces you to consume your opportunity inventory at a greatly accelerated rate. Oil and gas opportunities are not like stocks and shares: Oil and gas properties have an investment life—they get used up. Shares can be sold today and bought back tomorrow. Selling down from 25% to 15% working interest in an oil and gas property is harder, and buying back working interest or replacing it with new opportunities is harder still.

The problem with this assumption in the oil and gas business is we can neither change working interest of our projects quickly, at will, nor at minimal cost; changing working interest is time consuming and costly. The gradient search-visualization model described in this article assumes you can fund at the fixed working interest levels offered, delay, or not fund at all, but you do not have the freedom to fund an arbitrary level of the working interest.

Finally, the Markowitz process assumes that you can do just as good a job managing 100 projects at 20% working interest as 40 projects at 50% working interest or 20 projects at 100% working interest. In reality, however, when

the number of projects funded exceeds your ability to manage them properly, your portfolio risk must rise. For example, one expert said: “We believe that a policy of portfolio concentration may well *decrease* risk if it raises, as it should, both the intensity with which an investor thinks about a business and the comfort level he must feel with its economic characteristics before buying into it.” (emphasis in original) (Warren Buffett, cited in Cunningham, 2002, p. 82).

REFERENCES CITED

- Anstee, R., 2006, Math 340: Linear programming theorems: University of British Columbia, January 2006, <http://www.math.ubc.ca/~anstee/math340/3401ptheorems.pdf> (accessed April 15, 2007).
- Bjorndal, M. H., 1995, Some thoughts on combinatorial optimization: *European Journal of Operational Research* 83, p. 253–270.
- Cunningham, L. A., 2002, *The essays of Warren Buffett: Lessons for investors and managers*, revised ed.: New York, John Wiley & Sons, p. 82.
- Markowitz, H. M., 1952, Portfolio selection: *Journal of Finance*, v. 7, no. 1, p. 77–91.
- Rasey, S. M., 2005, The BlitzPort process: <http://wiserways.com/blitzport/process2.html> (accessed April 20, 2007).
- Woolsey, R. E. D., 1975, *Operations research for immediate applications*: New York, Harper and Row, p. 65–71.