

7-1-2009

An application of value-of-information to decision process reengineering

Jeffrey Keisler

University of Massachusetts Boston, jeff.keisler@umb.edu

Mark Brodfuhrer

General Motors Corp.

Follow this and additional works at: http://scholarworks.umb.edu/msis_faculty_pubs

Recommended Citation

Keisler, Jeffrey and Brodfuhrer, Mark, "An application of value-of-information to decision process reengineering" (2009). *Management Science and Information Systems Faculty Publication Series*. Paper 22.

http://scholarworks.umb.edu/msis_faculty_pubs/22

This Article is brought to you for free and open access by the Management Science and Information Systems at ScholarWorks at UMass Boston. It has been accepted for inclusion in Management Science and Information Systems Faculty Publication Series by an authorized administrator of ScholarWorks at UMass Boston. For more information, please contact library.uasc@umb.edu.

JEFFREY M. KEISLER
University of Massachusetts-Boston

MARK BRODFUEHRER
General Motors Corp., Detroit, MI

An application of value-of-information to decision process reengineering

Abstract: Value of information (VOI) methods were used to guide changes to recurrent organizational decision processes, under a reengineering effort at a major automobile manufacturer to reduce supply and demand imbalances involving capacity for parts and products. We modeled representative decisions assuming as they would be made with and without the benefit of improved information flows, and calculated the resulting increase in expected value. By factoring in the entire range of decisions affected by a process change, we scaled the value of each organizational change to a life-cycle value. The results quantified the impact of organizational changes in order to refine and prioritize a portfolio of change projects.

1. Introduction

Value of change as value of information

This paper applies the decision analytic idea of value of information (VOI) to plan improvements to an organization's decision processes and information flows. The theoretical idea of applying VOI to value information systems is well-known: for each possible information system, model the decision maker's information for each state of nature information and calculate the expected value for relevant decisions. This is difficult in practice, as often there are too many possible information systems or too

many decisions to model. We successfully adapted techniques from decision analysis (DA) in concert with a business process reengineering (BPR) effort to guide improvements at General Motors.

Planning changes typically requires prospective performance measures (Sarkis et al, 1997) and although information can be more nebulous than, say, materials handling, there is a literature on analytic methods for evaluating prospective information systems. The conceptual notion of using VOI (Raiffa, 1968) for this purpose is well-known (e.g., Vazsonyi 1976) and appears in standard management information systems textbooks (e.g., Turban and Aronson 1997, pp. 567-568). Both Hilton (1981) and Demski (1972) noted with disappointment the lack of practical applications of this concept. Such applications are still hard to find, although there have been recent efforts to evaluate (Kumar 1997, Khouja & Kumar 2002) decision support systems with regard to speed and flexibility, and to consider their optimal design (Herrmann & Schmidt 2002). These approaches are suitable when assessments for highly specific decisions and conditions can be obtained and timing of specific decisions can be controlled, e.g., within automated systems.

For simplicity, we shall use the terms P_1 and P_2 to denote the *as-is* process (the process before changes) and the *to-be* process (with desirable improvements).

Under the as-is process, the decision maker assigns probability distributions (in some cases deterministic assumptions) $X|P_1$ for a set of parameters $\{X\}$ such as the mean and standard deviation of the forecast a specific product's demand. P_1 does not necessarily represent the best information available within the company or make the best use of available information. Let $V(\delta, X)$ denote the expected value of decision δ given

the distribution X , and let $\delta^*(X)$ denote the decision that maximizes $V(\delta, X)$. Under an improved process, the decision maker assigns different probability distributions $X|P_2$. Since the purpose of changing a process is to remove sources of error found in the pre-existing process, we assume P_2 will be more reliable than P_1 . The question of interest is not whether P_2 is better than P_1 , but rather how much value this change adds and at what cost. To measure this, we calculate the expected value for each alternative under P_2 and compare the value of the best alternative to the expected value (still calculated assuming P_2) of worse alternatives that might be selected under P_1 .¹

The change in the probability distribution faced by the decision maker could arise in various ways: different people (who have different information) could make the decision; the same decision makers could receive different sources of information; the decision could be made at a different point in time; or using different assumptions. We call the value added in going from P_1 to P_2 the *value of change* (VOC). We define:

$$V_1 = V(\delta^*(X|P_1), X|P_2),$$

$$V_2 = V(\delta^*(X|P_2), X|P_2), \text{ and}$$

$$\text{VOC} = V_2 - V_1$$

Now, an important step is to relate VOC for real processes to the theoretical value of information so that we can model both the P_1 and P_2 states. We can think of changes

¹ To the extent that any technique embeds its own biases, it over-estimates the value of its perceived optimal choice compared to other choices (Smith & Winkler 2006). In calculating VOC, we made the implicit assumption that the new process would have removed biases that might be present in the current process, but we recognized that results should be interpreted with some caution.

along the lines suggested by Keisler (1992), where a signal flowing through an advisor to a decision maker may be improved in various ways as it is collected, processed, or communicated. Such changes result in different situations for the decision maker. For example,

- 1) P_2 simply adds information to that from P_1 , thereby reducing uncertainty. In this case, VOC is standard expected value of information. Where possible, we make simplifying assumptions (e.g., normal distributions) so that we can calculate this value only using summary statistics such as the mean and standard deviation of forecasts under P_1 and P_2 .
- 2) P_2 eliminates some bias. In this case, ex-post VOI for a given decision is easily calculated (i.e., reduction in expected loss), but using expected value of information is less practical here as it would require assessing a distribution on the amount of bias.
- 3) P_2 provides a more accurate depiction of the level of uncertainty. Here, VOC is similar to Henrion & Morgan's (1992) "expected value of including uncertainty," or alternatively, we might think of the ex-post value of improved information about a variable's as standard deviation or some other parameter.

Meeting challenges of this approach

In our application, the biggest challenge we faced was the sheer volume of decisions and information involved. There could be much to model, many potential interactions and many system alternatives. Therefore, we needed to speed up the modeling process, either by making fewer models, creating the models more quickly or by making the models easier to populate. Equally important, just having a definition of VOC and a list of decisions is not enough to calculate anything. We need ways to go

from managers' verbal descriptions of their problems to something calculable. Furthermore, if subjective judgments about what happens under P_1 and P_2 were needed, it was unclear how to obtain them. Certain characteristics of VOI problems and certain aspects of the information systems made these challenges manageable.

It would take impractically high analytic resources to create computational models at the level of fidelity appropriate for making the underlying decisions we wished to improve. One way around this problem is to aim for a lower level of fidelity. VOI estimates for fixed problem structures are often only moderately sensitive to parameter values near the optimum. For example, where two alternatives have almost the same expected value, a small perturbation to one of them could drastically change the allocation of resources, but value of information would remain nearly constant (see Keisler, 2004a). This same notion applies again at the level of decisions about the information system. The actual decision may be more sensitive to the precision of estimates than is expected value. Thus, order-of-magnitude type facilitative models (Soderquist, 2003) may be reasonable and useful if the purpose is to produce only value of information and related results. In situations where we cannot know the actual conditions of future decisions, this sort of sampling of the future decision space keeps the modeling workload manageable.

If we were to obtain actual results, we needed to impose structure at several points. The DA work and the BPR work were closely coordinated to support each other. For both structuring models and estimating input values, we exploited GM's base of experience, research, computer models and written reports. With these resources, some judgments required little more than a survey of existing data. To avoid difficult assessments, we

structured VOC models to use such simple judgments, while further impacts on the decision maker's state of information were calculated.

2. Application

GM's problem

In the mid and late 1990's, General Motors (GM) suffered periodic imbalances between production and demand for parts that it used in its family of products. For example, a surge in demand for a new sporty version of a popular existing family car could in turn lead to a surge in demand for rear spoilers. If, however, the company had contracted for too few spoilers then it would have to sell more basic vehicles than sporty vehicles. Some customers would be without their most desired product and substantial discounting could be required to sell what was actually produced. Conversely, lower demand for a product line or a whole class of products could lead to disproportionately large overcapacity problems for certain parts. This would result in losses due to contractual obligations, unnecessarily high upfront costs, and product discounting.

There was no easy solution – it would have been impractical to increase a decision-maker's workload by prescribing that they actively produce more precise information and analysis for thousands of separate decisions. Instead, the goal was to improve the match between the needs dictated by the overall decision stream and information flow generated by the organization, i.e., the content, timing, definitions and rules used for creating and distributing reports.

Using DA with BPR

GM initiated a BPR effort with the management consulting firm Strategic Decisions Group in order to align the related planning processes across the organization. We discuss this BPR as background in order to set the broader context for our particular DA work and because it seems to be a good starting point for this type of effort. Several planning processes were embedded in other functions of the huge company, and hence the charge for this effort was for major improvements of existing processes rather than radical change that would eliminate many processes.

Our goal was to improve ongoing processes to support series of recurring similar decisions. This meant assuring that decision makers would have available information of sufficient quality and would make the best use of that information, i.e., that the company would have a better information system (broadly conceived) for capacity planning. Key to this was ensuring that once decisions, forecasts and comprehension of uncertainty were available to one group, these would cascade consistently and quickly through the organization. The timing of decisions could likewise be managed in order to exploit late-arriving information by postponing commitments where reasonable. The fundamental benefit of successful BPR would be a drastic reduction in financial losses due to insufficient or excess production.

We needed to prioritize the portfolio of possible changes because there were many of them and relatively few people to implement them. The company would benefit if the changes with the highest payoff were completed first. We used DA techniques to value changes in order to prioritize them. From the start, we viewed parts capacity planning as a family of decisions and we collected qualitative data on it using BPR methods. These data would frame the DA modeling. BPR was to generate and then

implement specifications for individual projects to move the company from its current (as-is) state to a desired (to-be) state (Manganelli and Hagen 2003).

Mapping the system

The first step of the BPR² was *as-is* mapping. Because the engagement was focused on improving decisions, mapping consisted of identifying all players representing all related parts of the process and asking what decisions do you make, what information do you produce, and what information do you use? If we had not asked about each of these all along, the DA modeling that followed would have been much more difficult. From these data, we physically mapped the pertinent existing information flows.

We posted the map and refined it as GM experts and stakeholders viewed and responded to it. The map contained 29 nodes at which information was manipulated and 56 links between them. We reviewed the map with GM staff to identify gaps between the information that was needed and what was provided. These gaps suggested potential changes or change requirements (CRs). We were left with a complex but detailed picture of the system. A greatly simplified version of a portion of this map, still giving a hint of complexity, is shown in Figure 1. Out of 67 changes we ultimately considered and 20 that we modeled, this figure highlights four of them. In the paper, we explain these four changes in detail to illustrate the process we used to analyze many of them.

² The larger BPR effort involved a *process facilitation team* of five people (two internal GM managers, three external consultants) working closely with a *core team* of over twenty GM information creators, users and managers, in turn reporting to a senior *board of key leaders*.

These changes generally consisted of connecting poorly linked points. “Poorly linked” often meant that information relevant to a decision was available somewhere in the organization, but the information was not readily available to the decision maker in correct and usable form and in time to make the decision. Perhaps the decision was made too soon, or a relevant report never reached the decision maker, or two decisions were made independently (so that one decision maker just made a guess about what the other would do) instead of in coordinated fashion. There was no one centralized organizational unit that collated all information and made all decisions; rather, different decentralized organizational units made decisions based on information they had available. There was no appetite for a radical restructuring of the organization, so we mostly focused on strengthening the existing weak links.

3. Working with change requirements

Initial definitions

We initially defined change requirements in terms that were natural for the users of information. We review a few of these requirements and then discuss how we moved toward modeling their value. We consider now the four CRs highlighted in Figure 1:

CR1 (industry and segment volume forecast) involved the formulation of vehicle demand forecasts. These forecasts were prepared by GM’s Business Decision Support Center (BDSC) and were then sent to GM’s North American Operations Portfolio Planning and Capacity Planning office. This information was used for a variety of production capacity decisions. The concern was that the forecasts erred not so much due to lack of precise detail in forecasting, but due to lack of coordination and consistency

between different subgroups. An obvious symptom of this problem was that different groups actually used different forecasts. The definition of this CR was clear enough that the implementation team, knowing the context, understood what problems the CR would solve. But it was not at all obvious how to measure the impact of this requirement to “ensure the right information is provided by the right functional area at the right time, especially in areas [x, y and z].”

CR2 (production cost information) addressed the terms for contracting capacity GM Purchasing reported to the vehicle level executive teams. These teams treated capacity terms as fixed, and therefore did not consider how costs would vary in non-linear fashion with volume, and built plans around a point-estimate forecast. GM Purchasing actually had the ability to change contract terms. Ideally, Purchasing could understand and communicate to vehicle level executive teams relevant threshold levels that would cause suppliers to incur incremental tooling investment costs. Then these teams could make plans that reflect the impact of demand uncertainty on cost.

CR3 (target volumes) stated that “BDSC needs to provide timely split information to North American Operations Portfolio Planning and Capacity Planning on items that have long term implications on (capacity decisions for) the Body Shop.” This included split information on trucks (the percentage of different bed types, drive types, 3 vs. 4 door vehicles, and lift gates vs. rear door designs) and vans. In the as-is process, planners combined vehicle forecasts with generic split percentages (i.e., assuming the same percentage of leather seats on vans as on trucks) rather than vehicle specific split percentages because this information wasn’t available until later.

CR4 (capacity analysis) addressed the need for a brand and body-style specific distribution in capacity planning analysis. The as-is process used a common standard deviation for all products, even though the company had discovered that forecasts for new products had greater uncertainty. The risk assessment should include the distinctions that represent the major sources of uncertainty.”

Most other CRs similarly modified the way that some information or rule applies in some decisions, e.g. “Engineering Groups should not modify Vehicle Group plans or build in contingency (to capacity levels) but they should be strictly compliant with stated plans.” The CRs were thus easy to understand, but difficult to value.

Choosing which CRs to model

At core team meetings, we discussed which CRs to model in detail. This was an iterative process. Out of 67 CRs, several were technical quick hits that did not require modeling, others would interact with corporate operations well beyond the scope of our effort, some involved organizational issues that would contribute to more efficient operations but did not directly affect any easily identified decisions or parameters. About half the CRs involved specific enough decisions and information that it was plausible to obtain data and develop quantitative models. We listed pertinent questions that would benefit from answers, and then identified twenty decisions tied to specific CRs where we saw potential for quantitative models to provide answers. With unlimited resources, we would have modeled each of them extensively. Instead, we considered which CRs were hardest to prioritize and which ones could be improved with a better understanding of their value drivers. For ten such CRs, we developed full quantitative VOC models using spreadsheets and explicit judgments of experts outside of the process facilitation team.

We analyzed ten other CRs less formally, still using spreadsheets to gain insight but not seeking expert judgments or additional documentation. For the other CRs where it was meaningful but less critical, we discussed but did not model how improved information flow could lead to increased value in downstream decisions.

Common model structure for the capacity decision

We constructed a number of “quick and dirty” influence diagrams, using as experts mostly members of the core team or other GM analysts. We refined influence diagrams and decision tree models for several of the decision points. It helped that GM had a strong history with DA (Kusnic & Owen, 1992). We could draw on archival influence diagrams as well as a corporate forecasting template that used accounting identities and pro-forma business case parameters such as market share and market size (as in DA textbook examples, e.g., McNamee and Celona, 2001).

We realized that, because they related to supply/demand imbalances, many of the decisions shared the common structure of the influence diagram shown in Figure 2. Although we believe it is likely that in many organizations, sets of decision flows would have their own common structures, it was still fortunate that we discovered one here. It simplified the task of modeling numerous changes and made the models easy to compare.

Specifically, the decisions involved setting capacity or quantity supplied for some product or part in the face of uncertainty about quantity demanded. Quantity demanded (e.g., the “net option demand” node in Figure 2) was driven by different parameters for different parts or options. It could often be derived hierarchically from either demand for a given family of GM vehicles, demand for a version of that vehicle (e.g., 2-door), or demand penetration for an option among buyers of that version of the vehicle. Profit

maximizing capacity levels would take into account contribution margins (both from parts sold as options and from vehicle sales enabled by the availability of parts), as well as fixed costs (e.g., tooling) associated with capacity.

We calculated the optimal decision (typically capacity, although we could easily incorporate different decision variables into the same basic influence diagram) under the given information for each process using critical fractiles or similar simple rules, and we assumed normal or similarly tractable distributions. We then calculated the expected corresponding value received for each case.

In some instances, we considered more than two possible information states, i.e., if there were multiple ways in which the CR might be implemented we estimated value of information about more than one variable. Although we felt that most of the changes could be modeled as independent, it was simple to modify multiple variables. This could help identify interactions that make it worthwhile to cluster certain changes.

We model this capacity decision as a “newsvendor” problem (Arrow, et al 1951, Hillier & Lieberman 2005) in that we view the relevant investments as creating a perishable asset. Typically, GM spends money to configure durable assets from its existing base (whose costs treat as sunk), e.g., buildings and general use machines, to produce a specific parts and designs. The dedicated configuration (and related spending on equipment such as dies) has little or no salvage value after the next major vehicle update (or after the dies wear out). Alternatively, GM’s capacity decision can simply result in a contract with a supplier that commits to make such capacity available at a given cost.

In this case, rather than newspapers, capacity is what is purchased and used. The expected additional contribution from the marginal unit of capacity is the product of its contribution if used and the expected number of times it is used. This varies by vehicle type, e.g., for demand a sporty car's demand may drop sharply after the first year, while a family sedan's demand may be relatively stable for a number of years. Capacity is added up to the point at which the marginal unit's expected contribution is less than or equal to its cost. This point is the critical fractile of the demand distribution, where the probability of demand exceeding capacity is equal to the ratio of cost to contribution margin. If, as we typically assumed, demand follows a normal distribution, the critical fractile is simply calculated and easily modified for different CRs. The critical fractile rule was implemented in a spreadsheet template (Figure 3) and as described below.

Demand (D) is normally distributed with mean $E(D)$ and standard deviation $\sigma(D)$; f_D denotes the probability density function for D , and Φ_D denotes the cumulative probability function for D .

Letting F denote fixed cost, K denote the capacity decision variable, and C denote annualized cost per unit of capacity, we note that the total cost is equal to $F + KC$.

We let M denote the contribution margin (marginal profit for a unit that was demanded, then produced and sold). When demand exceeds capacity (which happens with probability $1 - \Phi_D$) production is fixed at the capacity level. When demand is below capacity, production varies with (and is equal to) demand, and contribution is integrated as in equation 1 below. We add together the expected contribution from both cases and subtract costs. Thus, the expected value received calculated over the assumed distribution of demand levels is equal to

1) ~~CONFIDENTIAL~~

For the remaining calculations, we use the following notation:

Target percentile (Row 5 in Figure 3, for calculating critical fractile) $T = 1 - C/M$;

Optimal capacity = $\Phi_D^{-1}(T)$, where Φ_D^{-1} denotes the inverse cumulative normal distribution over D. (Row 8 in Figure 3)

Q = quantity produced to meet demand = $\min(D, K)$;

The annual profit π for given demand and capacity, where $\pi = MQ - F - KC$;

Since there is a $1-T$ chance that demand will exceed capacity, the expected annual production for the optimal capacity (Row 21 in Figure 3) is given by:

$$2) \quad E(Q) = (1-T)K + \int_{-\infty}^K x f_D(x) dx$$

This is calculated straightforwardly in the spreadsheet (rows 10-18 in Figure 3 calculate the linear loss type of integral in Equation 2).

Then we can write the expected profit as:

$$3) \quad E(\pi) = ME(Q) - F - KC.$$

If the demand distributions used are unbiased, expected loss compared to the ex-post optimal capacity depends only on the standard deviation of the demand forecast. So if P_1 has a normal distribution on D with standard deviation σ_1 and P_2 has the same mean as P_1 but has standard deviation reduced to σ_2 , the expected value of the contemplated change would be $E[\pi | \text{standard deviation} = \sigma_2] - E[\pi | \text{standard deviation} = \sigma_1]$.

We assumed normal distributions around variables and entered means and standard deviations to the spreadsheet. We were able to structure variations with additive or multiplicative relationships between variables so that, as uncertainty about vehicle

demand propagated to uncertainty about part demand, D would continue to follow a normal distribution³. The spreadsheet template proved flexible.

Input data for the spreadsheet models

In obtaining numerical inputs for our models, i.e., assessment, we were interested in not only future costs, but uncertainty levels for a stream of products whose details were not yet defined. Under these conditions, we could not rely on standard DA assessment questions that involve subjective probability judgments about specific events. Keeping in mind that we are aiming for order-of-magnitude type VOC results, we hoped it would suffice to sample the future decision space by detailing typical situations that might occur and running calculations for them. As earlier, it helped that we had access to the forecasts and results for past product decisions. In most cases, we had no reason to think that conditions had substantially changed and we assumed that a sampling of past conditions would be a reasonable proxy with which to simulate future conditions. To obtain base case assumptions about contribution margins relative to cost, we used as a representative vehicle (or part) one whose parameters fell near the median or mean of the range observed for the family of decisions toward which the change was targeted.

Often, the most critical parameter in our spreadsheet model was the standard deviation on the volume forecast. We gave this parameter the most detailed attention and

³ Value of information results should be robust to the exact choice of distribution as long as the same percentile is used, since for given mean and standard deviation the value of information calculated with the linear loss integral is bounded (e.g., by Chebyshev's inequality). If a distribution were clearly characterized by rare but major events, it would be preferable to build that into the model, as we did in one case.

our assessments drew on various empirical data. There were prior GM research studies on forecast error, including one which compared forecast error for major vs. minor vehicle changes. In some cases, we looked at old planning forecasts and uncertainty ranges given, the assumptions used in those forecasts, and how they compared with what actually occurred. There was not enough for formal distribution fitting, but there was enough to draw some rough conclusions. We conferred with GM internal staff who conducted product decision analyses, and when necessary with the planners and decision makers who provided their inputs.

We assigned probability distributions for hypothetical situations, e.g., if we had a major change to the C-platform sports car, and what would be the uncertainty in the forecast, what would be split between the lower and higher demand body styles. In most cases, this worked something like it did for CR1, described shortly. For simplicity, we often assigned normal distributions (suitable for forecast errors) with standard deviations approximated to multiples of 10% of the mean, e.g., 0% for no uncertainty, 10% of mean for low uncertainty, 50% for high uncertainty. These numbers were assigned by trained decision analysts (in the role of experts), but were not assessed using special DA techniques, and we based these estimates on data from similar previous situations.

The effect of CRs was typically to incorporate information about forecasts, splits, and even levels of uncertainty that ought to already be available. For this reason, we were not faced with difficult to assess subjective distributions. Instead, the post-change states were characterized simply as good practice, e.g., facts which should be known were assumed to be known, and forecast errors were assumed to be at the level that the

company should achieve when the process worked correctly. The pre-change states were the post-change states compromised by the various errors we described.

In the product plans we reviewed, cost of capacity and contribution margins were provided primarily by accounting and finance, demand forecasts were provided primarily by product managers. Uncertainty in those forecasts was usually articulated in the form of 10th, 50th and 90th percentiles that GM analysts obtained from product managers.

Mathematical models of the CRs

As we shall see in the four illustrative CRs, a few basic variations could be applied in combination to adapt the template to our VOC models. The numbers are disguised. Fixed costs were the same under P_1 and P_2 , and so are not included in the calculations of VOC. In the newsvendor-based VOC models for the specific CRs described here, capacity cost parameters and contribution margins were treated as deterministic.

CR1: This change requirement is directed at improved estimates on vehicle demand for new programs in particular, as these have the most risk. The focus on new programs affects the estimates of parameters involved and is relevant to determining the frequency with which decisions affected by the change occur. To model its impact, we utilized the spreadsheet template, and incorporated the assumption that the actual demand for a given vehicle will be equal to the forecast demand plus a normally distributed error term. The company must decide how much capacity to acquire. The cost terms here are at the vehicle level (fixed cost for capacity, unit cost for capacity, contribution margin, etc). We looked at several representative vehicles to estimate these numbers, including one where an inaccurate forecast had actually led to problems.

We first estimated the standard deviation of the error without further improvement – where the decision maker fails to use the best forecast that combines views from around the company. For this, we estimated a baseline error for standard vehicles (where assumptions are better shared, and where there was some expertise within the company even about the level of uncertainty in the forecasts). We then qualitatively estimated the size of potential errors for other types of vehicles, by looking at empirical data showing the variation among parties’ forecasts and between these forecasts and the eventual sales that actually occurred for historical examples.

We assumed that with the improved estimation process, the decision maker rightly uses a standard error more like that for most other vehicles. The decision maker’s distribution is centered on the correct mean under the new process, but not necessarily under the original process, and we compare the expected value obtained under the assumption that the initial process had an unbiased estimate for the mean (P_{1a}), or where it was assumed too low (P_{1b}) or too high (P_{1c}). For P_{1a} , P_{1b} and P_{1c} , the assumptions for C, M, E(D) and $\sigma(D)$ were entered in cells C3,C4, C6 and C7 of the template shown in Figure 3, while K and E(π) are calculated in cells C8 and C22; assumptions and calculations for P2 are in the same rows of column D, and VOC is calculated in cell C23. Thus, P_{1a} has the correct mean and the wrong standard deviation on demand, and P_{1b} and P_{1c} have not only the wrong standard deviation but also the wrong mean, while P_2 has the correct mean and standard deviation.

CR1	C	M	E(D)	$\sigma(D)$	K	E[$\pi(K) P2$]	VOC
P_{1a}	\$1000	\$5000	100,000	50,000	142,000	\$355,378,000	\$9,467,000
P_{1b}	\$1000	\$5000	80,000	50,000	122,000	\$364,775,000	\$30,000

P_{1c}	\$1000	\$5000	120,000	50,000	162,000	\$337,455,000	\$29,360,000
P_2	\$1000	\$5000	100,000	25,000	121,000	\$364,805,000	–

Given the large standard deviation and high margin, this situation would typically result in a lot of extra capacity. In case P_{1b} , where the original forecast is below expected demand, the capacity at the critical fractile happens to be near the correct capacity under P_2 (because the high standard deviation leads to inflated capacity on high margin products), so then the ex-post value of the change is low.

CR2: The original CR referred to the fact that targets for vehicle production were being treated as requirements for parts production. To model this situation, we interpret the change requirement as meaning that a new decision rule should be specified. We assume that without the change, the decision rule takes the mean of the true input distribution on the quantity of a part that will be required. It treats this mean as a point estimate without uncertainty, and sets an optimal capacity for this point estimate. We calculate the expected value of this production level and compare it to the expected value for the proposed new decision rule: treat the point estimate as having the uncertainty it actually does, and set the optimal capacity for that distribution. The change will not reduce standard deviation in forecast distributions, but will change how the distribution is utilized. The value added here relates to value of information, but is not exactly the same –value is added by using information correctly, not merely by acquiring information.

We selected a representative part, e.g., a high-markup option such as sunroofs, and estimated F and C for the production of that part. The contribution margin for the part is the difference between the premium we charge for a vehicle with that part

(alternatively, the amount by which we would have to discount a vehicle if the part were not available) and the variable cost of producing that part. The as-is decision rule is modeled as: $K = E(D)$. The to-be decision rule finds the optimal capacity given that standard deviation on demand for the part is σ , using our template.

CR2	C	M	$E(D)$	$\sigma(D)$	K	$E[\pi(k) P_2]$	VOC
P_1	\$600	\$2400	175,000	0	175,000	\$264,733,000	\$10,227,000
P_2	\$600	\$2400	175,000	52,500	210,411	\$274,960,000	–

We assumed $F = 0$ (because capacity is outsourced). Under P_2 , K targets the 75th percentile of demand. VOC is positive even though the expected utilization of the plant is lower for the new K derived under P_2 than the K derived under P_1 , i.e., 79.4% vs. 88%. The VOC here is for a decision about a high-priced option. Most parts are cheaper. We expected decisions on approximately five such high-priced options per year.

CR3: This CR mandates that information about the relative demand for different body-styles be timely and complete. To model this CR, we assumed that under the current system, North American Operations capacity planning ultimately formulates its own estimates about how many units of each body-style are needed. This estimate is based on the mean and standard deviation (15%) from BDSC’s forecast over the demand summed over both body styles. We considered a specific example of 4-door (style a) vs. 3-door body (style b) minivans.

Under P_1 , manufacturing makes an uninformed guess that there is a 50%–50% split among body-styles a and b, and makes a capacity decision for style. Under P_2 , the estimated split is 60% for style a, and 40% for style b.

Again, we targeted the 80th percentile. Note, identical values of C , M and F were used for CRs 1, 3 and 4, because for the relevant vehicles, these were approximately correct and we saw no reason to change the sample case more than necessary. For other CRs that were most relevant to vehicles with relatively low margins, different parameter values were used.

From the viewpoint of P_2 , P_1 gets the critical fractile wrong and yields a high probability of shortfall on vehicle type a , and low expected utilization on vehicle type b .

CR3	C	M	$E(D)$	$\sigma(D)$	K	$E[\pi(K) P_2]$	VOC for vehicle type
P_{1a}	\$1000	\$5000	120,000	15%	150,298	\$298,362,000	\$17,166,000
P_{1b}	\$1000	\$5000	120,000	15%	150,298	\$128,048,000	\$15,637,000
P_{2a}	\$1000	\$5000	144,000	15%	180,358	\$315,528,000	<i>TOTAL a + b</i> \$32,803,000
P_{2b}	\$1000	\$5000	96,000	15%	120,239	\$143,685,000	

CR4: The issue here is that the more radical the change in a vehicle from its predecessor, the more uncertainty there is about its demand. Some vehicles represent updated versions of older products, while other vehicles are entirely new. Instead of assuming that the standard deviation on the sales forecast is always the historical 30%, an improved approach would allow decision makers to customize forecast distributions to the type of vehicles. For example, vehicles that are essentially unchanged have lower standard deviation (perhaps 10%), and radically new vehicles have higher standard deviation (perhaps 50%). We calculated P_1 , assuming that the standard deviation σ is 30% of mean demand. Holding expected sales constant, we calculated P_{2a} assuming σ is 10% and P_{2b} assuming it is 50%.

Here we calculated capacities as before, and computed profit levels for both types of vehicles.

CR4	C	M	$E(D)$	$\sigma(D)$	K	$E[\pi(K) P_2]$	VOC for each case
P_1	\$1000	\$5000	150,000	30%	187,873	a-\$361,988,000	\$17,015,000
						b-\$288,527,000	\$6,487,000
P_{2a}	\$1000	\$5000	150,000	10%	162,624	\$379,003,000	-
P_{2b}	\$1000	\$5000	150,000	50%	213,122	\$295,014,000	

Several others versions of the model developed along the same lines, and illustrate how a wide variety of changes might translate to template models. For example, one major part came in two styles, one of which appeared in about 60% of vehicles and the other in 40% of vehicles. Faulty communication occasionally led to larger capacity being committed for the wrong version. This was equivalent to having a 5% chance of using the wrong assumption for mean volume for the two parts (P_1) and the rest of the time using the right assumption (P_2), and VOC was thus 5% of the decreased profit for an instance when the mistake occurred.

Another CR was to consider uncertainty about the mix of demand at the brand level conditionally for each body type (P_2), as opposed to just calculating body type split and brand mix independently and multiplying the percentages to get the demand for each specific vehicle (P_1). The reasoning is that different option packages would be required for each combination (e.g., sunroof with the exciting two-door Pontiac, no sunroof with the four-door Pontiac or with any Chevrolet). Thus, there is more uncertainty about demand under P_2 than P_1 , and for a high markup item, more capacity is appropriate.

Other CRs could allow improved information about contribution margins, or other variables that we treat as fixed in our basic model.

Lifecycle value of change

To scale up from VOC on one decision to VOC across all affected decisions, an enumerative approach would have been to list all such decisions and redo the model for each of them. In our situation – and we suspect in most real ones – even if it had been possible to name each future decision, there would not have been time to model them all. Instead, we analyzed a single decision in-depth for which the change is hoped to produce its benefit, e.g., setting production capacity for sports car spoilers. We estimated the VOC for an instance of that decision and then surveyed the classes of decisions that would go through the same process and that we expected to have similar benefits as a proportion of spending involved (this is key to defining the relevant set of affected decisions). This approach matched how GM organized its production scheduling systems, so obtaining these numbers was straightforward.

For parts decisions, we factored in the number of parts per vehicle model that would be affected by a change. From the type of program (e.g., sports car production) we determined the set of affected vehicles. We asked how many times per year new programs of this type occur. We asked how many years the change would benefit a new program, e.g., the current process could be self-correcting for the relevant programs, so that an imbalance only would apply for one year (which was common), or an imbalance once created might persist throughout the life of the program. Finally, we recognized that decisions for different products involving the same type of decision may have costs and quantities that are lower or higher by some amount than the specific case we analyzed.

With simple arithmetic, we combined all these into a factor representing the number of equivalent decisions by which to scale up the single decision's VOC to get a lifecycle VOC.

4. Results

Impact

We summarized the results for all the VOC models (not just CR1-CR4) in a portfolio-analysis type dashboard as in Table 1. We could easily think about the cost of changes in terms of person days. We informally weighed costs against value. Based on their payoffs, we designated the changes as having high, medium or low priority. Senior management accepted these recommendations, and implementation teams were assigned with guidelines that fleshed out the formal CR definitions.

The VOC models had benefits beyond prioritization. In some cases, the requirements were further refined as a result of insights that were directly due to these models. In CR1, for instance, we added guidance about when the value of improved information is high and hence when that information should be required. For CR4, we added to the original requirement the direction to ensure that forecasts make explicit which sources of uncertainty have and have not been incorporated, and that large downsides or upsides be noted separately. Many other CRs were also enhanced this way. In some cases, the decision rules embedded in VOC calculations were not just sources to inform the definition of CRs, but served as prototypes to potential solutions for those CRs. For example, the company could make better use of critical fractile methods that comprehend uncertainty in setting the amount of flexible capacity needed, and in setting

various other capacity levels. In other cases, variables with high value information themselves (and thus topics for future corporate research) were identified, e.g., a need to understand how the price elasticity of demand for various options differs by body style.

We can assess the ultimate impact of this effort within the context of the BPR effort as well as within the broader corporate context. The modeling phase of the project ended at the point that the implementation of changes began. As such, the modeling effort played a crucial role in selecting which changes to pursue. Beyond driving recommendations, the detailed models aided in understanding and articulating to the various stakeholders the reasons why one change or another would pay off more, thus facilitating consensus about how to move forward.

Quantitative benefits: The direct benefit of this project was in the form of improved estimates of the annual value that changes would provide, and the use of these estimates to prioritize changes. It was difficult to prioritize these changes without the quantitative model results. Thus, as a rough approximation, without this work all the changes from this set were equally likely to have been in the high-priority group as the low-priority group (which were not implemented until two or more years later). Because of the improved estimates, GM would benefit from the high valued changes instead of the average changes over two years between the implementation of the first and second round of changes. By this reasoning (Keisler 2004b), the value added by our prioritization was on the order of \$100M. This figure assumes that our value estimates were reasonably accurate (as the board of key leaders seemed to agree). We performed at least some analysis on one third of the CRs. About half the entire set of CRs had effects

on decisions that were easily enough articulated that this approach could plausibly have been used.

Qualitative benefits: As we look back on the project several years later, the changes that were thought to be important appear to have really been so. The large-scale BPR effort was a major success. The insights and implications of VOC modeling clearly contributed to elements of that success.

Key leaders Pat Jansen, Senior Manager, Capacity Planning and Richard Willson Director of Manufacturing Planning at GM described the benefits realized from the improvements to the decision process:

“System improvements have been identified and implemented to correctly translate and communicate program intent into the detailed level required by Purchasing. This active management of supplier tooling rates has resulted in fewer constraints at new product launch due to early attention to flexibility requirements and product mix. ... Savings in overall tooling investment [had been] documented and the potential for over-tooling due to poor communication has been minimized.”

Brian Hagen (at that time the lead partner from Strategic Decisions Group, Inc., on the BPR initiative) confirms the criticality of VOC modeling to the overall effort:

“The "Value of Change" approach provided a significant breakthrough on the project as it helped our project team simplify our characterization of the link between strategic change and corporate value creation. Ultimately, the approach allowed the project team to pinpoint the sources of value resulting from change,

more easily prioritize recommendations, and explain – in simple terms – the impact of recommended changes to a senior level executive team.”

This effort provided a language for improving the definitions and describing the effects of changes, and for deciding implementation priorities. It generated significant direct value in the form of a focused implementation plan, and it generated indirect value by making the reasons for the recommended changes more transparent to the rest of the organization.

Lessons learned

Some of our tactics were improvised to the specific situation we faced, with both constraints and resources that practitioners might not always have. Specifically, we benefited from having organizational memory about DA. Certainly, if practitioners have a similar resource, they should use it. On the other hand, we undertook a large analytic project without realizing how much modeling work would have to be done in a short time, so the focus on rapid model development might not be as great in other cases.

We can also draw some general lessons. First, we now know that it is feasible to apply value of information to the design of organizational information systems and decision processes. It remains for others to demonstrate that this can be done for more subtle system design questions that would require more finely tuned analysis. In a realistic setting, the dimensionality of the problem may be too great to create comprehensive models.

Several shortcuts, concepts and insights seem generally helpful to push our approach through to a successful conclusion in any setting. It was productive to integrate BPR and DA methods, as the former are well-suited to mapping processes and defining

changes. Classic elements of concurrent engineering facilitate this, i.e., team members had shared, multiple and requisite skills, and planned interactions structured to ensure feedback and feed-forward. To support DA, process mapping should explicitly address decision points and information that flows into them. DA cannot dominate this step. Instead, managers must define the changes and cluster them around certain decision points. The analyst then translates these changes into DA models. Representing different changes with separate models for separate decisions keeps the task simpler than trying to create a unified simulation of the entire organization. If the BPR is focused on a set of related decisions, some common structure – such as the newsvendor problem – will probably facilitate rapid creation of many similar models.

To translate from verbal descriptions to quantitative models of changes, we formulate from the problem description a typical DA model that a decision maker might use. We track the flow of information from its origin to the decision maker's final use of it. Specifically, we focus on characterizing the amount of uncertainty in key parameters such as demand forecasts and on how much of uncertainty can be reduced merely by removing rather specific sources of noise or confusion in the organization so that the best available information will be used at the most opportune time. It is important to identify carefully the categories of future decisions that are affected by the change, to then tally the benefits of the change over that stream. We know that the results from this approach can at best be only approximate, but that VOI may often be robustly estimated and applied – and produce insight – even if the core decision model is inexact.

The concept of VOI (and VOC) turns out to be a useful tool for improving decision processes, not just individual decisions. By understanding when information is

valuable, we can help organizations sort through the glut of information they possess. Many in-house decision analysts or richly embedded consulting teams could adapt this approach to resolving systemic organizational problems they observe. By quantifying the impact of these organizational changes, we gain ability to plan and implement them.

Acknowledgments:

We thank Dr. Brian Hagen of Strategic Decisions Group for his strong personal and intellectual support during this effort as well as his comments on this paper. Kim McGill of General Motors and Suzanne Wurster of Strategic Decisions Group were members of our process facilitation team. They helped to define and inform the modeling effort described here and we thank them. The entire change effort, and the decision analytic piece in particular both owe all their success to the engagement, cooperation and support of many others at General Motors, including extended team members and subject matter experts. We also thank the anonymous reviewers and numerous helpful readers.

References:

- Arrow, K., Harris, T. and Marschak, J. (1951) Optimal Inventory Policy. *Econometrica*, 19, 250-272.
- Davenport, T. (1993) *Process Innovation: Reengineering Business Processes through Information Technology*. Harvard Business School Press, Cambridge, MA.
- Demski, J. (1972) *Information Analysis*. Addison Wesley, Reading, MA.

- Henrion, M. and Morgan, M. G. (1992) *Uncertainty: A Guide to Dealing with Uncertainty and Quantitative Risk in Policy Analysis*. Cambridge University Press, Cambridge, UK.
- Herrmann, J. and Schmidt, L. (2002) Viewing product development as a decision production system. *Proceedings of 2002 ASME Design Engineering and Technical Conference*. DETC2002/DTM-34030.
- Hillier, F.S. and Lieberman, G.L. (2005) *Introduction to Operations Research*, 8th edition, McGraw Hill, New York, NY.
- Hilton, R. (1981) The Determinants of Information Value: Synthesizing some General Results. *Management Science*, 27, 1, 57-65.
- Keisler, J. (1992) *A Framework for Organizational Decision Analysis*. Doctoral Dissertation. Harvard University, Cambridge, MA.
- Keisler, J. (2004a) Technical Note: Comparative Static Analysis of Information Value in a Canonical Decision Problem. *Engineering Economist*, 49, 4, 339-349.
- Keisler, J. (2004b) Value of Information in Portfolio Decision Analysis. *Decision Analysis*, 1, 177-189.
- Kusnic, M. and Owen, D. (1992) The Unifying Vision Process: Value beyond Traditional Decision Analysis in Multiple-Decision-Maker Environments. *Interfaces*, 22, 6, 84-93.
- Kumar, R. (1997) Understanding the Value of Information Technology Enabled Responsiveness. *The Electronic Journal of Information Systems Evaluation*, 1. (<http://is.twi.tudelft.nl/ejise/>).

Manganelli, R. and Hagen, B. (2003) *Solving the Corporate Value Enigma*. AMACOM, New York.

McNamee, P. and Celona, J. (2001) *Decision Analysis For The Professional* (3rd edition). SmartOrg, Inc., Menlo Park, CA.

Raiffa, H. (1968) *Decision Analysis*. Addison-Wesley, Reading, MA.

Sarkis, J., Presley, A., and Liles, D. (1997) The Strategic Evaluation of Candidate Business Process Reengineering Projects. *International Journal of Production Economics*, 50, 261-274.

Smith, J., and Winkler, R. (2006) The Optimizer's Curse: Skepticism and Post-Decision Surprise in Decision Analysis. *Management Science*, 52, 311-322.

Soderquist, C. (2003) Facilitative modeling: Using small models to generate big insights. *Systems Thinker*, 14, 9, 2-6.

Turban, E., J. Aronson. (1997) *Information Technology for Management*. Wiley, New York.

JEFFREY M. KEISLER is an Associate Professor of Management Science and Information Systems at the University of Massachusetts-Boston. Previously he had over ten years experience as a decision analyst, primarily with Strategic Decisions Group, Argonne National Laboratory, and General Motors. His B.S. is in Mathematics and Computer Science from the University of Wisconsin-Madison, his M.B.A. is from the University of Chicago, his S.M in Engineering Science and his Ph.D. in Decision Sciences are from Harvard University. In addition to The Engineering Economist, Dr. Keisler's work has been published in such journals as Risk Analysis, Interfaces, Decision Analysis, and Energy Economics. His research focuses on value of information, various aspects of decision process design, and portfolio decision analysis.

MARK BRODFUEHRER is a Team Manager with General Motors Strategic Initiatives group. He has thirty years of automotive experience including manufacturing, engineering, and research and development. Over the last ten years, he has guided business decision support projects, applying analytic methods such as decision and risk analysis, system dynamics, and game theory to a wide variety of issues of importance to senior leadership. These issues have included global vehicle programs, software and connected vehicles, fuel cells and battery technology, brand and channel growth, crisis preparedness, order-to-delivery, option packaging, manufacturing flexibility and capacity planning. He has presented and contributed to INFORMS. Mark graduated from Lehigh University with a BS in Mechanical Engineering.

Figure 1: Map of information flows (simplified)

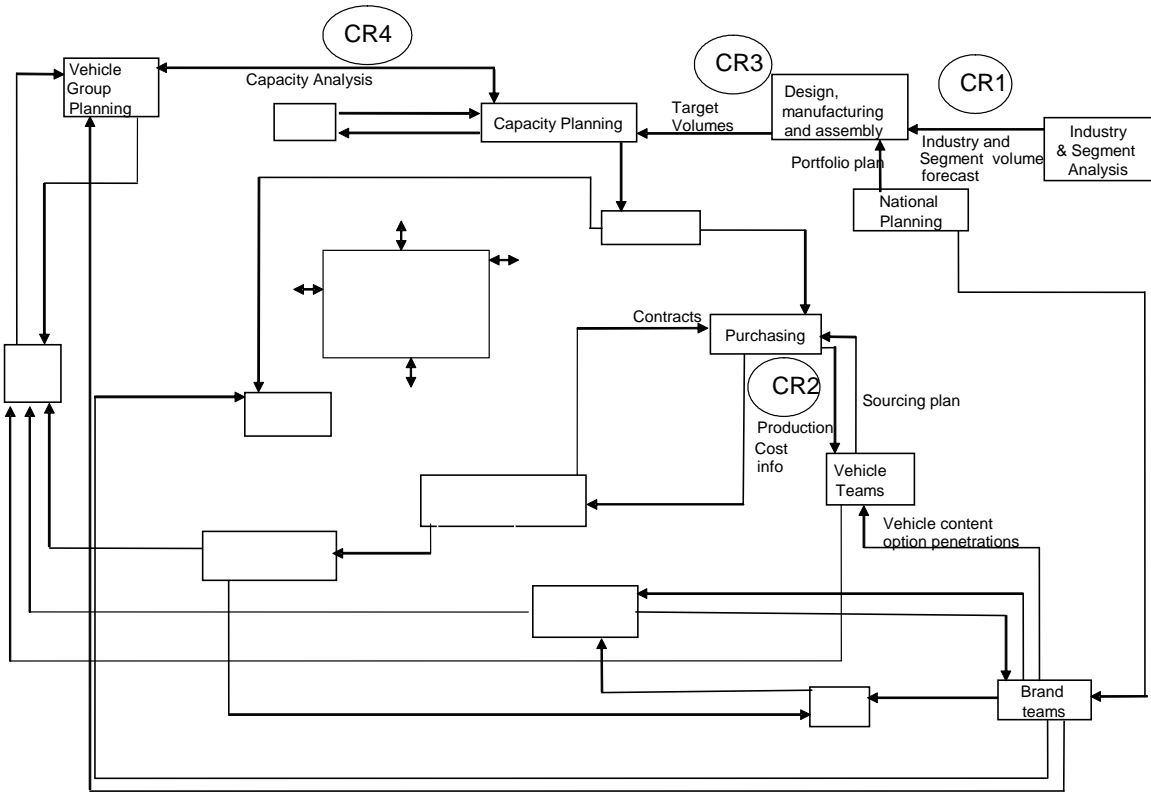


Figure 2 Generic influence diagram for newsvendor type parts capacity decisions




Figure 3: Value-of-change spreadsheet template (illustrative values).

	A	B	C	D	E
1		Calculation of P2 View	P2 View	P1 View	
2	Fixed cost (F)	10000000	\$ 10,000,000	\$ 10,000,000	Calculations
3	Unit cost of capacity (C)	250	\$ 250	\$ 250	using capacity
4	Unit contribution (M)	1000	\$ 1,000	\$ 1,000	from P1,
5	Target percentile (T)	=1-C3/C4	75%	75%	demand
6	Demand mean [E(D)]	130000	130,000	150,000	distribution
7	Demand Std Dev	39000	39,000	67,500	from P2
8	Optimal capacity (K)	=NORMINV(C5,C6,C7)	156,305	195,528	
9	Nominal profit	=C4*C8-C3*C8-C2	\$ 107,228,825	\$ 136,646,044	
10	Nominal extra capacity	=C8-C6	26,305	45,528	65,528
11	Standardized extra capacity	=C10/C7	0.674	0.674	1.680
12	Stdized prob density of D at K	=NORMDIST(C11,0,1,0)	0.318	0.318	0.097
13	Cumulative probabily of D at K	=NORMDIST(C11,0,1,1)	0.750	0.750	0.954
14	Right tail std hazard of D at K	=C12-C11*(1-C13)	0.149	0.149	0.019
15	Left tail std hazard of D at K	=C12+C11*C13	0.824	0.824	1.699
16	E(D) D < K	=C8-C7*C15/C13	113,476	121,400	126,023
17	E(D) D > K	=C8+C14*C7/(1-C13)	179,573	235,800	211,636
18	Pr(D < K)	=C13	75%	75%	95%
19	Expected unused capacity	=C8-(C8*(1-C18)+C18*C16)	32,122	55,596	66,276
20	Expected utilization	=(C8-C19)/C8	79.4%	71.6%	66.1%
21	Expected quantity	=C20*C8	124,183	139,932	129,252
22	Expected profit	=C4*C21-C3*C8-C2	\$ 75,106,714	\$ 81,050,081	\$ 70,369,633
23	VOC	=C22-E22	\$ 4,737,081		

Table 1 Results for portfolio of possible changes (illustrative values)

Change requirement	Typical beneficiary (*disguised)	Annual value of change per application	Frequency	Annual value
1 BDSC communication	Sportarama*	\$10M-\$50M (key assumption is that higher quality forecasts could be used and differ from brand forecasts up to 20%)	All new programs (3-5 per year)	\$100M High priority
2 Clear direction to purchasing regarding options	Most vehicles	Savings on tooling of \$10M per program per year, if it is true that purchasing currently uses expected demand as capacity, options are 40% of tooling cost.	Recover after 2 years. Applies to most programs (10 per yr)	\$40M Medium priority
3 More detailed truck/van split info	Mini-vans 3 rd vs. 4 th door	\$30M (key assumption is that plug used for assembly decision is off by about 10%.	Maybe important for one program per year, can recover within 1-2 years at some cost.	\$15M Low priority
4 Brand / body style specific distributions	Chevy 999* carryover has less uncertainty than new segment	\$10M	Every program, but value of change will vary – lower if program is typical or lower contribution	\$50M Medium priority
5 through 20 would be shown on later pages				