Go Back

Dynamic, hard and strategic questions:

Using optimization to answer a marketing resource allocation question*

Alan K. Graham and Carlos A. Ariza
Business Dynamics Practice, PA Consulting Group, Inc.

One Memorial Drive, Cambridge, Massachusetts 02142
Voice: 617-864-8880, Fax: 617-864-8884
Alan.Graham@PAConsulting.com
Carlos.Ariza@PAConsulting.com

Abstract

For historical reasons, optimization has traditionally been slightly outside the mainstream of System Dynamics. However, computer technology has both made quantitative data more abundant and optimization more feasible. At the same time, modelers are encountering real situations and clients with high-stakes questions that are nearly impossible to answer without optimization – the systems involved are not only dynamic, and not only highly interconnected, but also combinatorially daunting. In particular, corporate marketers currently make allocation decisions impacting billions of dollars of shareholder value on the basis of intuitive "anchor and adjust" strategies, which can be far from optimal.

This paper presents an anonymized case study of one such situation. The company is in a high-tech industry undergoing rapid change. The company needed to fashion a go-to-market strategy balancing traditional and unfamiliar markets, an important component of which was allocation of marketing resources. Optimization revealed a potential valuation increase of roughly 30% relative to executives' intuitive allocations. Scenario analysis revealed the basic policy direction (more advertising) to be robust, and in the process resolved several traditional conundrums in dealing with adverse events in the marketplace.

Keywords: Marketing, Advertising, Resource Allocation, Optimization, Optimal, Marketing strategy, Dynamic Model, Shareholder value

---

* The authors wish to thank the other members of the PA team who contributed to the modeling and client work reported here. Alphabetically they are: Ben Bryant, Sharon Els, Peter Genta, Jane Hemingway, Tameez Sunderji, and Mary Tolikas. We are also grateful to the visionary yet practical people at the client company and their advertising agency who regrettably should remain anonymous at this point. You know who you are, and we thank you.

1. Introduction: Times have changed

Optimization fits naturally into System Dynamics. Indeed, they are sister sciences, both born from the study of servomechanisms in the 1930s and 1940s.[1] In the first four years of the *System Dynamics Review* (1985-1989), five articles focused on optimization.[2] And work had already been going on for some time, particularly at the University of Bradford by Coyle, Wolstenholme and others, complete with software support through the DYSMAP optimization package. Keloharju and colleagues at Helsinki School of Economics also extensively investigated optimization of dynamic economic systems.

But optimization remained a poor relation in the System Dynamics world. In *Industrial Dynamics,* Forrester in effect suggests that optimization is a temptation that distracts modelers from creating good and useful models in the first place:

> Another evidence of the bias of much of today's [late 1950s] management science toward the mathematical rather than the managerial motivation is seen in a preoccupation with "optimum" solutions. For most of the great management problems, mathematical methods fall far short of being able to find the "best" solution. The misleading objective of trying only for an optimum solution often results in simplifying the problem until it is devoid of practical interest. (Forrester 1961, Section 1.2)

Indeed, Forrester even saw dangers in the use of objective functions, either as measures of model fit to data, or as being useful or appropriate for evaluating policy:

> The insecure system designer looks for "objective" criteria for the evaluation of the pertinence of a model. The compulsion to use such objective measures is sometimes so strong that refuge is taken in procedures that lack a sound foundation (Forrester 1961, Appendix O)

> Since most industrial systems seem to operate so far from a hypothetical ideal, it is reasonable to hope that system improvements can first be obtained without requiring any compromise. Improving one factor may not require paying a penalty elsewhere. That is the situation here. (Forrester 1961, Section 18.3)

And under the conditions, these were very reasonable warnings:

1. The management science optimizations of the time were often analytical exercises on one- and two-equation models, clearly far simpler than would be trustworthy or appropriate to answering major operational or strategic questions.

2. The balance between available qualitative data and quantitative data tilted much more strongly than today toward the qualitative -- "There is usually a paucity of dependable records that can be compared with a system model" (Forrester 1961, Section 18.1.6). Modelers needed to achieve great strength in qualitative understanding of the system and its dynamics.

3. The initial experience with System Dynamics was with problems of instability and oscillation, where the policy changes often would not have a truly direct cost, and so tradeoffs were not as dominant an issue as they would have been in either a resource allocation or growth dynamics setting.

And so after the first flurry of optimization studies in the *System Dynamics Review*, a full ten years went by before the next publications using optimization.[3]

But over the intervening years, both the operating realities of modeling, and the needs of end-users – corporate clients -- have changed. Optimization on today's PCs is now susceptible to brute force, even for realistically complex models, if judiciously applied. Optimization software is now much more widely available, simply as a feature of shrink-wrapped simulation software. And in this era of spreadsheets and Enterprise Resource Planning (ERP) systems, quantitative data are vastly more plentiful (even if not yet of perfect quality). So optimization is much less of a constraint and distraction, and generally much easier.

Moreover, the situations in which dynamic modeling is used have shifted mostly away from oscillations into questions that more directly call for trade-offs, such as effectiveness, market share, and growth. All of these require explicit examination of complex tradeoffs over activities, and over time.

Finally, the corporate community has become more sophisticated, perhaps even jaded, in the realm of strategic thinking, becoming aware of the limitations of more standard approaches to strategy. We have seen both the "fall and rise of strategic planning" (Mintzberg 1994). And we have seen explicit recognition that there are many situations for which standard strategy tools and processes are and always will be inadequate, not only because of communication and execution (Campbell and Alexander 1997), but of fundamental limitations of methodology. Three McKinsey consultants expressed this nicely in the *Harvard Business Review* (Courtney *et al.* 1997):

> The old one-size-fits-all approach is simply inadequate. Over time, companies in most industries will face strategy problems that have varying levels of residual uncertainty, and it is vitally important that the strategic analysis be tailored to the level of uncertainty. [pg. 73]
>
> Our experience suggests that at least half of all strategy problems fall into levels 2 or 3, while most of the rest are level 1 problems. [pg. 71] At level 1, a single forecast of the future…is precise enough for strategic development. The forecast will be sufficiently narrow to point to a single strategic direction. [pg. 69] In order to perform the kinds of analyses appropriate to high levels of uncertainty [levels 2-4], many companies will need to supplement their standard strategy tool kit. Systems [sic] dynamics and agent-based simulation models can help in understanding the complex interactions in the market. Real-options can help in correctly valuing investments in learning and flexibility. [pg. 78]

To this analysis we would only add that it is not only uncertainty but also the feedback complexity of problems that render traditional strategy analysis tools misleading or useless. There is good evidence from controlled experiments that human beings, even seasoned managers, are surprisingly poor at decision-making in an environment rich in feedback, delayed effects and nonlinearities (Morecroft 1985, Sterman 1987, 1989a, 1989b). This general result was also true of the marketing case described in the following section, as we shall see.

At the present time, then, there is both capability and need to do optimization for dynamic, hard, and strategic problems. We need no longer focus on algorithms and computation, but instead on documenting the practicalities and specific benefits of using optimization for such problems.

2. The Case

2.1. The company

The (anonymized) company whose case we examine here operates in a high-tech industry, which like biotech, communications and media was undergoing continual major transitions, in underlying technologies, in the mix of markets being addressed, and the competitive landscape. On one hand, the future was quite uncertain, but on the other hand, it was a company with strong engineering capabilities, with the ability to develop and bring to market revolutionary technologies.

PA was called in to develop a system dynamics capability that would enhance the company's strategic planning. Traditional tools for planning such as financial spreadsheet analysis were becoming increasingly ineffective in analyzing and understanding the future of an increasingly complex corporation. In addition, the planning of individual business units was steadily less and less informative for overall corporate planning, which had to deal with aggregate risks of the overall product portfolio, and limitations on the total "size of the bets" that were placed on each prospective product line. Moreover, operating realities were increasingly disconnected from the results of more traditional strategy exercises. (Graham 2001) describes the larger set of model uses in strategy formulation at the company.

2.2 The Marketing Muddle-Through

Corporate Marketing became one of the earliest internal clients for the simulation capability, to support resource allocation decisions. The newly-centralized marketing organization had received a fixed annual budget by the time we started working with them. A new corporate marketing officer had been appointed to guarantee that resources would be fairly distributed across categories, but it quickly became clear that the existing allocation process provided no easy, trustworthy answers.

The approach typically recommended in marketing textbooks is to benchmark the resourcing decisions of successful companies, in terms of Cost Per Revenue Dollar (CPRD), and adjust according to intuition and judgment. And so marketing managers had provided an approximate answer by adjusting average marketing cost per revenue dollar up for growth product lines and down for mature product lines. They also asked product managers to request levels of spending for each of their products. Of course this was a time-consuming process, but eventually (about two months into the fiscal year) consensus was reached about the best way to allocate resources.

Apart from the model-and-optimize method reported here, there really was no other way to arrive at decisions.[4] In terms of decision theory, this was a classic example of "anchor and adjust" (find some number, any number, and adjust it in an intuitively correct direction) and "satisficing" or in slightly more informal terms, "muddling through."

But for this strategic, high-stakes dynamic problem, there was a combinatorial explosion. With more than 6 product lines, more than 3 channels, and more than 5 types of marketing (advertising, types of promotions, and more), managers were suddenly allocating among more than 90 buckets.

Moreover, the outcomes for each of the more than 90 buckets impacted both other product lines, and extended out over time. Some cross-connections: advertising one product line would to some extent boost brand awareness of another, installed base in one would help gain share in another. And all of this was in an environment where competitors also acted, and even overall product demand depended on the actions taken.

Finally, the complexity interacted with the uncertainty. Some of the product lines were new, their precise characteristics were unknown and their future was uncertain.

The characteristics of the decision environment – multiple inputs and outputs, delayed effects, feedback – were precisely the characteristics that laboratory experiments suggest that managers could not do well intuitively. So it would have been virtually impossible for intuitive solutions to come close to even near-optimum performance. For that matter, suppose a modeler did indeed have a complete and valid dynamic model, but was searching for a better marketing resource allocation only by one-at-a-time experiments. Because of the large number of control "levers" the modeler would almost certainly not be able to find a solution close to optimum within a reasonable amount of time and effort. So the problem was not only "dynamic," and not only "strategic," but also "hard."

2.3. The model

Figure 1 shows the three main sectors of the model: Market Demand, Market Share and Companies (both our client and its competitors). Each of those sectors was subscripted to represent several product lines. In addition, the market share sector tracked share in each of several sales channels, again for each product line. For each of the sectors in turn, we will describe the phenomena being represented, and then comment specifically about interactions among product lines and data available for calibration.

2.3.1. Market Demand. The product markets represented in this sector range from mature product categories (for which several years of historical demand information was available) to new product categories (for which data was more anecdotal, and in which the company's own actions in part determined how fast market demand would develop).

In its mature product lines, the company had traditionally relied on one technology platform to drive its growth. Although fast price declines cause the overall revenue numbers to grow more slowly than unit volume, there was still significant demand ahead from both new installations and replacements of these products. Industry demand for these products was fairly stable, and therefore the model represented the growth of the installed base by relying on demographic trends and historical performance. This was further enriched with information about industry changes in pricing and information about technology changes (which could cause end-users to hold purchases until a new generation of products is released.)

For new product categories, the model represented how vendors influence the speed at which the category developed. Through the combined marketing and sales efforts of all vendors, potential users became aware of new products. Once aware, potential users eventually moved to consideration of purchase, during which marketing activities could improve end-users' perception of core product features. When purchases and installed base grew, word-of-mouth effects kicked in to generate additional awareness, consideration, and purchase.

There was interaction among demand for different product categories, since they all related to a common technology platform. There were two types of interactions, which can be called sales-driven and installed base-driven. Sales-driven demand is the "pull-through" created by solution sales, in which a (somewhat) fixed quantity of one product's sales is tied to the sales of another. (An economist might call these *complementary goods*.) By contrast, installed base-driven demand reflects the fact that owning one product makes adoption of other more likely – for example, no one would purchase a VCR without a TV, but if a TV is present, VCR adoption becomes more likely. In the long term, there was potential substitution among the various product lines as well, although this did not appear to be a significant effect over the time span of interest.

Information was surprisingly available on the size of potential markets both past and future, primarily from third-party sources. While the optimism or pessimism of any given information source was subject to question, there were a variety of cross-checks available. First, there were usually several sources available. Second, there was information about the "demographics" that drive a given product market. For example, if one were trying to understand the market for car stereos, one would also check projections against sales projections for new and used cars. Third, expert input was often available from within the company on how a given market could evolve.

2.3.2. Market Share. The company competed against other firms (represented as a single aggregate in each product category) to capture a fraction of total demand. There were a multiplicity of drivers for market share, including price, product performance, installed base, and a variety of forms of sales and marketing measures. The latter included size and experience of the direct sales force, customer awareness (which was in turn driven by relative advertising and public relations spending), end-user promotions, and sales channel support of various types. In determining the effect of each of these measures on market share, the measure was taken relative to the competing firms. These measures were weighted differently for each product and channel, so that a product that was successful in one channel could be a failure in another.

There were two interactions across product categories in determining market share. First, installed base in similar products impacted purchasing behavior, especially in new product categories, where no vendor had a track record to guide purchasers. A vendor's share of the installed base in a similar, more mature product could increase a purchaser's confidence when buying in a relatively immature product market. This effect became less important as the new category grew, since decision-makers had more information about the vendors' performance in the new market. A second interaction was an advertising "halo effect", where advertising in one product category could spill over to increase customer awareness and consideration in another, given that products had similar benefits or characteristics.

Data around market share was diverse and critical to the optimization process. We had historical market shares from external vendors in the various product markets, and competitive price information as well. We got interview data for share drivers such as relative product performance and the weighting or importance of each driver. There were some commissioned studies on brand awareness in various target markets, and some generic third-party research as well. Internal sources gave us historical data on marketing spend in the various buckets. In the process of calibrating the model, we achieved a rough level of confidence in the relationship of market share to marketing resourcing decisions. Remaining uncertainty, especially about the future, would be dealt with through scenario analysis.

We were also able to cross-check the model against the limited number of available external studies that quantified the effect of advertising on market share, sales and profitability. (Urban *et al.* 1986) offered a controlled cross-sectional study of consumer goods advertising on market share dynamics, as did (Schroer 1990) more anecdotally. (Jones 1990) characterized advertising needed for competitive equilibrium. The description of how new product categories evolve benefited from (Urban, *et al.* 1993, 1994), as well as traditional academic market models of awareness, consideration and purchase derived from the Bass model and others.

2.3.3. Company Sectors. At a high level, the model determined pricing and fully-burdened costs for each of the product lines and thereby computed a profit-and-loss (P&L) statement, for each product category and for the corporation as a whole. (In parallel, there was a P&L for product lines within the aggregate competing firms and an aggregate P&L.) Based on cash flow available and internal policies, the model made budget allocations to a variety of activities, including the marketing activities that were the subject of the analysis reported here.

The model also calculated a shareholder valuation proxy measure based on multiples of earnings and revenue to estimate the impact of the company's actions on shareholder value.[5] In effect, the model computed a weighting of valuations from price/earnings ("P/E") and price/revenue ("P/R"). Revenue growth impacted both the P/E and P/R ratios.[6] This valuation measure was central to the optimization, since it was the basis for the objective function that was maximized.

Within the company sector, product categories interacted primarily through the corporate resource allocation process. Mature product categories funded the technical and market development of new product categories, to an extent determined by internal policies. Potentially, there were additional cross-connections through economies of scale in shared technologies,[7] and motivation and/or turnover of employees being influenced by aggregate stock price (and therefore employee option performance).[8] But for this particular company's history, structure, and performance, these were very minor factors, and thus not represented in the model.

In addition to the marketing and share data mentioned above, we had P&Ls both at the corporate level, and by business unit. We mapped the latter approximately onto product categories, and allocated corporate expenditures as well, so that the model's P&L approximated a true Activity-Based Costing system, cross checked against both the corporate totals and internal expert opinion. For the valuation proxy, we had the company's stock price. We also had a selection of data for similar industries from the S&P 500 that we regressed to find appropriate coefficients in the valuation formula.

So we had developed and calibrated a simulation model that captured key cause-and-effect relationships for about 75% of the company's product categories by revenue. The model was well-covered by the available data, in the sense that no causal relationship was distant from a variable for which we had data. Iteratively, model behavior came to nicely replicate the overall behavior shown in the data. The model was particularly rich in structure around the marketing budget and its effects in the marketplaces for each of the product categories, both mature and new.

3. The optimization

3.1. The process: Constrained optimization in multiple scenarios

The optimization varied only the "handles" which were under the client's control. So the optimization reported here was only allocation amongst marketing buckets (product category and type of marketing activity). For the Corporate Marketing client, we did not address the wider (and more difficult to implement) question of what the appropriate budget for marketing overall should be.

The objective function used in the optimizations was the shareholder value proxy, averaged over the next two years. The company was under severe pressure to restore profitability and its current stock value, and this objective function seemed closer to what the management and stockholders were actually desiring, by comparison to a more traditional cash flow discounted out into the indefinite future.

The optimization used a simulated annealing algorithm, which randomly varied parameters and gradually allowed the better combinations of parameters to emerge. It was therefore relatively well-protected against getting stuck in local optima, and choice of initial conditions was not critical. The initial resource allocation was the marketing managers' proposed budget allocations that result from the traditional, heuristics-based approach. (This allowed us to estimate the value added by the optimization exercise.) Typically, the optimization got reasonably close to a stable value after a few thousand simulations.

Through discussions with the client, we identified key uncertainties that could potentially have a major impact on the nature of the results. They were:

- What if the new advertising campaign is less effective than expected?
- What if the products in the new categories are less appealing to consumers than expected?
- What if competitors spend substantially more on advertising?
- What if investors are *really* short-term – looking at value only over the next 9 months?

Optimizing under each scenario was a form of robustness testing. (Forrester 1969 Appendix B.3) makes the distinction between sensitivity of behavior and sensitivity of policy conclusions.)

For each of these potential realities an optimal budget allocation was calculated, and so could answer the question "is this the best you can do?"

We also performed a number of more conventional sensitivity tests, verifying, for example, that the computed optimal policy in fact still produced a substantial increase in value, even if the importances or weights given to various drivers of market share were quite different. (The weights were derived from interviews, and were therefore not as certain as other data.)

3.2. Optimization findings

The optimization gave allocations to the ninety-plus buckets over which marketing had some control. For obvious reasons, only selected aspects of the results will be discussed here. Figure 2 shows optimization results under several scenarios. The improvement from the base case to the baseline optimized allocations provided a boost of *over 30% in value*. Moreover, the more general conclusion was remarkably robust under a variety of assumptions: *the company must substantially increase its use of advertising.*

In some respects, the optimum solution was not different in spirit from that originally proposed: the optimum solution focused resources on new product categories, and lowered the amount of short-term, tactical spending which had been the norm for much of the company's history. An important difference, however, lay in the magnitudes of changes. Without analytical tools, the solution reached intuitively tends to be less of a departure from historical spending, reflecting uncertainty and risk aversion from the part of managers. The basic results were generally greeted with relief by both the Corporate Marketing group and (perhaps unsurprisingly) their ad agency. Typical response: "we had a strong feeling we should be doing a lot more advertising, but we couldn't justify it." But marketing executives closer to individual products or channels tended to be quite skeptical about pulling money away from their bailiwicks – their own major products, or their own selling activities. Very understandable, but in retrospect, such views had clearly biased the intuitive allocation process.

3.2.1. Local optima or ridges? Although the optimization gave very clear and robust results along some dimensions (like the allocation to advertising in Figure 2), the search was surprisingly slow to converge along other dimensions. In particular, some solutions threw resources toward one new product category, and other solutions threw resources toward another. In general, the solutions seemed to concentrate resources, rather than disperse them uniformly among product categories, and within broad ranges, it didn't much matter which among several categories were chosen as the focus. But the valuations for all of the near-optimum solutions were very close to each other.

Of course, the model was pervasively nonlinear, both in dynamics and in response of valuation to, e.g. growth rates, so multiple optima are quite possible in theory. In particular, there are economies of scale in marketing (Jones 1990) that imply that focusing resource on a few product categories would be more effective use of resources than spreading them evenly.

It should make little practical difference whether there are truly local optima which happen to have nearly the same valuation, or whether there are just "ridges" in parameter space, where a variety of policies make very little difference to the outcome. But it does confer a major practical benefit to know that one of the two has happened, since either will differentiate those issues that executives must stand firm on, versus those issues that have relatively minor influence such that executives can yield on them for political or implementation reasons.

3.2.2. Does short-term vs long-term make a difference? Especially brand-related advertising (as opposed to point-product or promotional advertising) is often positioned as an investment, with the implication that the company should make a short-term sacrifice for the longer-term good. But with the particular situation here, even if one optimized over the very short term (specifically, the next 9 months), the optimal allocation to advertising didn't shift much, and it was still substantially higher than the then-current allocation.

3.2.3. Weak campaigns and products: Bolster or abandon? In some industries, like motion pictures, the rule is that if products don't do well initially, advertising budget is vigorously slashed. By contrast, Figure 2 shows that the optimal ad allocation for products that don't do well was larger. This might be termed the "Microsoft strategy", which relied on aggressive marketing of existing products even while better products were being developed.

3.2.4. Competitive advertising challenge: Meet it or slipstream? Another conundrum in the marketing world is how to respond to competitor's ad campaigns. In consumer marketing, there is a well-known correlation between "share of voice" (a company's share of total advertising being done regarding a given product category) and share of market (Jones 1990). This fact in isolation would imply that the company should attempt to match competitor's increases in advertising.

But there is another school of thought that might be characterized by a bicycle- or car-racing analogy. Sometimes it's better to save energy and let a competitor go ahead just a bit, so you can travel easily behind them in their slipstream. If the product markets were immature enough that a major effect of advertising was to raise awareness of the product line, a competitor who increases their advertising was spending their money expanding the overall market, making it easier to capture share. In this particular situation, the optimal response to more competitor advertising was to "slipstream" the competition, by spending somewhat less on advertising.

3.2.5. Developable markets: How much more? For markets where the whole product category is new, marketing expenditures really have two functions: not only to convince customers to buy the company's product rather than a competitor, but also to convince customers that they want to buy this kind of product in the first place. In the original marketing budget, marketing directors had allocated somewhat more funding to such products relative to their revenues. But optimization suggested that their allocations, even in percentage of revenue terms, were still too low by a factor of over 3. This was an example of the marketing executives understanding the issues and considerations properly, but by using intuition alone, missing the right quantitative results by substantial margins.

3.2.6. Mature markets: How much less? On the flip side, mature markets deal in products where not only are customers aware of the vendors and products, but also what features are available now and shortly from which vendors, where to find third-party product reviews, and in many markets, Value-added resellers to talk to about the purchase decision. In short, information is already abundant and marketing can add little. So in the corporate portfolio, the optimal share of marketing spend for mature products would be less than the share of revenues or value they create. And marketing executives know this. To offer a slightly disguised example, one mature product line contributed 30 percent of the value of the company, and was allocated only 20 percent of the marketing budget. But the optimization revealed that, under all the scenarios tested, the optimum share of the marketing budget was around 5 percent. Again, executives knew to make the adjustment, but underestimated the magnitude of adjustment required, in this case by a factor of 4.

4. Discussion

On the practical side, we have demonstrated "in the field," just as in the laboratory (Sterman *et al.* 1995), that intuitive solutions to dynamically complex business strategy problems fall short of what was achievable by systematic analysis. Intuitive adjustments to allocations were directionally correct, but magnitudes were too small by factors of three and four. The degree of improvement over the intuitive solutions – roughly 30 percent in value terms -- was surprisingly close to that found in a telecommunications strategy study (Graham and Walker 1998) and studies of very large development and construction projects (Graham 2000).

The study reported here can extend straightforwardly to wider strategic inquiry. Beyond marketing allocations is the total size of the marketing budget. Beyond those are similar questions about engineering allocations and budget, with consequences for product performance and its influence on share and revenues. Finally, there is price position (cut rate, or premium brand?) and choice of distribution channels. So optimization squarely attacks some of the most fundamental strategic questions:

- What businesses should we be in?
- What should our product positioning be (price, branding, performance)?
- Through what channels should we distribute the products?
- How should we respond to our competitor's actions?

Tactically, we would like to go further in exploring alternate approaches. Depending on the current state of the business and its strategy deliberation, we would use different constraints and different "handles." Rather than experimenting with different time horizons over which to optimize value, we would likely constrain short-term profitability explicitly – any strategy that makes the company into takeover bait should be avoided.

One the methodological front, we would like to explore the matter of multiple local optima and / or ridges. A faster optimization algorithm like hill-climbing, along with explicit examination of first and second partial derivatives of the objective function, would clarify the situation. Because strategic situations abound with economies of scale, winner-take-all, first-mover advantage and nonlinearities in general, near-local optima seem likely to continue to be both a technical issue and an opportunity to add strategic insight and value. No doubt we will expand the optimization and analysis process to more effectively identify and characterize such situations.

On the matter of robust conclusions, rather than optimizing under a variety of scenarios, one could explicitly optimize for robust conclusions by maximizing allocations, pricing, etc. with respect to not the valuation on a single scenario, but the minimum over all of the scenarios under consideration. One would also want to experiment with weighting the extent to which some optimal policies changed from scenario to scenario—this would sort out robust policy results from opportunities for real options analysis.

Further along the robustness path would be giving probabilities to a variety of scenario conditions and optimizing on a variance-weighted expected value of Monte Carlo simulations, in effect evaluating a risk-adjusted return. Still further along that path would be to explicitly optimize over scenario-dependent strategy options, in effect solving the real options problem for the company's portfolio of product lines and their strategies. The company would find out the best way to spend their money, given both a probability that some of the product line "bets" (investments) won't pan out, and the future opportunity to halt further investment under defined conditions.

Along a separate methodological front, we can start to answer the question "how dependent are the strategic conclusions on the information used to calibrate the model?" Clearly, from a technical viewpoint, one can always pick model parameters that negate or even reverse the improvement obtained from a given recommended strategy. Equally clearly, few if any such model parameter changes will be realistic, either *a priori* (just from knowledge of the real world, we know they can't be right) or *a posteriori* (if we try to make the model fit observed historical behavior, we can't still get the model to reproduce known history). There is theory and the beginnings of practice to address this question of quantifying confidence in results of dynamic models.[9]

For the overall field of System Dynamics over the past decades, we can observe major increases in both the ability to easily conduct optimizations and the need to explicitly optimize (rather than conduct only hand-experimentation). This experience suggests to us that optimization should be moved to a more central position in the practice of System Dynamics, to be taught and used on a more consistent basis.

Notes

[1] Jay Forrester generally identifies three primary sources for system Dynamics: study of management decision-making, computer simulation, and servomechanisms/cybernetics (Forrester 1961, Chapter 1). For an analysis of historical roots see also (Richardson 1991)

[2] Coyle 1985, Mohapatra and Sharma 1985, Kivijarvi and Tuominen 1986, Wolstenholme and Al-Alusi 1987, and Macedo 1989.

[3] Dangerfield and Roberts 1999, Kleijner 1999, and Bailey *et al.* 2000.

[4] There is a body of portfolio optimization theory, but it is not helpful in the corporate marketing situation, both because the cross-connections violate the assumptions, and because as a rule they focus on buying and selling of "buckets" of assets, rather than the allocation of resources to existing buckets. Such methods primarily focus on assets like oil-producing properties or stock holdings. These optimizations, however, take a view of risk and performance that, by

comparison to the problems discussed here, is extremely static and simplistic. Risk and payoff are assumed as inputs, and there is no interaction among the elements of the portfolio (Jorion 2001, Section21.4.2).

There are also a variety of marketing optimization tools currently available. However, they are all (to the authors' knowledge) tactical rather than strategic, for example focusing on how to spend money in an advertising campaign among a variety of print and broadcast media options.

[5] One alternative to a multiples-based valuation measure is to use net present value of discounted cash flows (DCFs). Although DCFs are theoretically appealing, they rely on projections that go beyond reasonable time horizons for planning, on a measure of terminal value that is itself multiples-based, and an often-prohibitive amount of financial detail surrounding capital expenditures and the balance sheet. In the current technological environment, estimates beyond four years are usually ignored, since technological and market change is difficult to predict and feel comfortable about characterizing.

[6] The formula for the shareholder valuation proxy can be considered to modify the well-known "market convergence" variant of DCF (See Copeland *et al.* 1995, pp. 293-295) in two ways. First, the formulation is all terminal value (since the model lacked detail needed for explicit capital investments and balance sheet tracking). Second, the formula represents the expectation that current revenues will in due time turn into market rates of profitability, so that the formula contains not only a term proportional to a profitability measure but also a term proportional to revenues.

[7] One of the most respected frameworks for strategy analysis among academics is the *resource-based view of the firm,* which instead of looking at a corporation's portfolio of product lines, looks at the underlying resources and capabilities that drive the success (or failure) of those product lines in the marketplace. (Wernefelt 1984) is the seminal work. Because the product lines were separated not only organizationally, but also technically and often geographically, there was negligible resource-sharing apart from funding, which the model did capture.

[8] Sterman models a corporation where the interaction between stock price, employee performance, and employee recruiting and turnover were major influences on the dynamics. The model is the central element in a publicly-available management game, described at http://web.mit.edu/sdg/www/

[9] This issue has been analyzed extensively inside PA, under the rubrics of "data-constrained outcome sensitivity analysis (DCOSA)" or "the cone of confidence." (Graham, Mullen and Choi 2001) provides more discussion of this issue and a case analysis.

Figure One

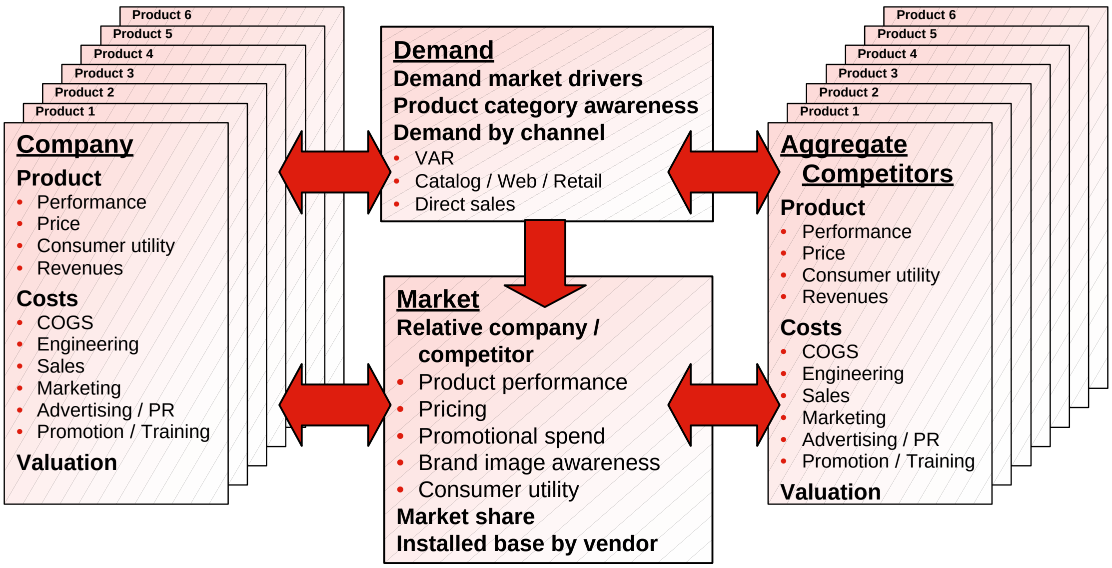

Figure 1. Model architecture, including Demand, Market share, and two Company sectors, competing across a number of product lines.

Figure Two

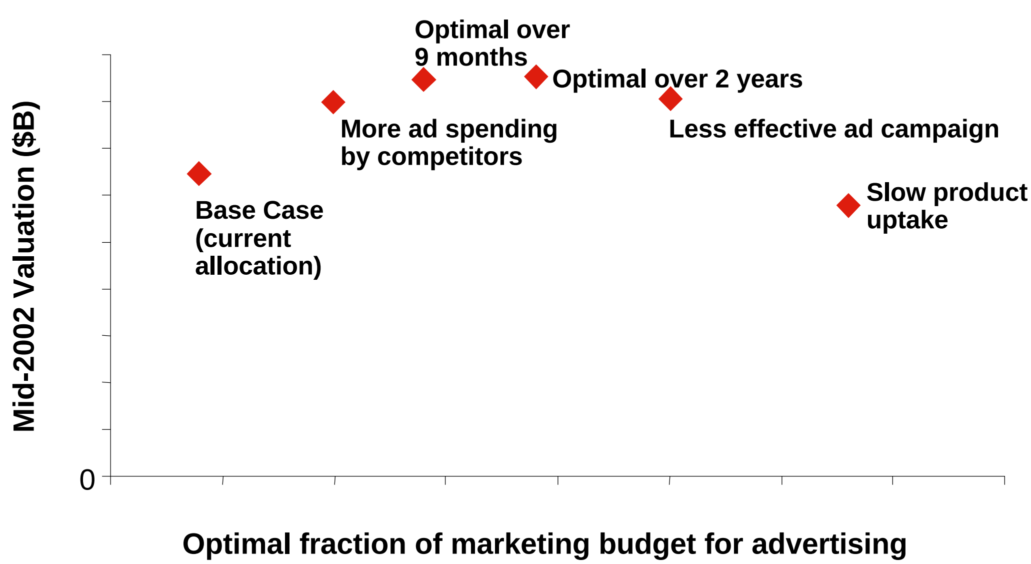

Figure 2. Optimal allocation to advertising, versus other marketing "buckets" for several scenarios, with current base case for comparison.

References

Bailey, Reid, Bert Bras and Janet Allen 2000. Using Response Surfaces to Improve the Search for Satisfactory Behavior in System Dynamics Models. *System Dynamics Review* 16(2) 75-90.

Campbell, Andrew and Marcus Alexander 1997. What's Wrong with Strategy? *Harvard Business Review*, Nov-Dec 1997, 42-51.

Copeland, Tom, Tim Koller and Jack Murrin 1995. *Valuation: Measuring and Managing the Value of Companies.* New York: John Wiley.

Courtney, Hugh, Jane Kirkland, and Patrick Viguerie 1997. Strategy under Uncertainty. *Harvard Business Review*, Nov-Dec 1997, 67-79.

Coyle, Geoffrey R. 1985. The Use of Optimization Methods for Policy Design in a System Dynamics Model. *System Dynamics Review* 1(1), 81-91.

Dangerfield, Brian and Carole Roberts 1999. Optimisation as a Statistical Estimation Tool: An Example in Estimating the AIDS Treatment-Free Incubation Period Distribution. *System Dynamics Review* 15(3), 273-291.

Forrester, Jay W. 1961. *Industrial Dynamics*. Waltham, Mass.: Pegasus Communications.

Forrester, Jay W. 1969. *Urban Dynamics.* Waltham Mass.: Pegasus Communications.

Graham, Alan K. 2000. Beyond PM101: Lessons for Managing Large Development Programs. *Project Management Journal* 31(4) 7-18.

Graham, Alan K. 2001. Business Dynamics at the Center Square in the Game of Strategy" (forthcoming)

Graham, Alan K., Thomas W. Mullen and Carol Y. Choi 2002. Using Fit-Constrained Monte Carlo Trials to Quantify Confidence in Simulation Model Outcomes. *Proceedings of the 2002 Hawaii Systems Conference* (forthcoming).

Graham, Alan K. and Robert J. Walker 1998. Strategy Modeling for Top Management: Going Beyond Modeling Orthodoxy at Bell Canada. *Proceedings of the 16th International Conference of the System Dynamics Society, Quebec, Canada, July 20-23.*

Jones, John Phillip 1990. Ad Spending: Maintaining Market Share. *Harvard Business Review*, Jan-Feb 1990, 38-42.

Jorion, Phillipe 2001. *Value at Risk: The New Benchmark for Managing Financial Risk.* New York: McGraw-Hill.

Kivijarvi, Hannu and Markku Tuominen 1986. Solving Economic Optimal Control Problems with System Dynamics. *System Dynamics Review* 2(2), 138-149.

Kleijner, Jack P. C. 1999. Sensitivity Analysis and Optimization of a System Dynamics Model: Regression Analysis and Statistical Design of Experiments. *System Dynamics Review* 15(4), 275-288.

Macedo, Julio 1989. A Reference Approach for Policy Optimization. *System Dynamics Review* 5(2), 148-175.

Mintzberg, Henry 1994. The Fall and Rise of Strategic Planning. *Harvard Business Review*, Jan-Feb 1994, 107-114.

Mohapatra, P. J. K. and S. K. Sharma 1985. Synthetic Design of Policy Decisions in System Dynamics Models: A Modal Control Theoretic Approach. *System Dynamics Review* 1(1), 63-80.

Morecroft, John D. W. 1985. Rationality in the Analysis of Behavior Simulation Models. *Management Science,* 31(7), 900-916.

Richardson, George P. 1991. *Feedback Thought in Social Science and Systems Theory*. University of Pennsylvania Press.

Schroer, James C. 1990. Ad Spending: Growing Market Share. *Harvard Business Review* Jan-Feb 1990, 44-47.

Sterman, John D. 1987. Testing Behavioral Simulation Models by Direct Experiment. *Management Science* 33(12), 1572-1592.

Sterman, John D. 1989a. Misperceptions of Feedback in Dynamic Decision Making. *Organizational Behavior and Human Decision Processes,* 43(3), 301-335.

Sterman, John D. 1989b. Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiment. *Management Science* 35 (3), 321-339.

Sterman, John D., Rebecca Henderson, Eric D. Beinhocker and Lee L. Newman 1995. A Behavioral Analysis of Learning Curve Strategy. Cambridge, Mass.: MIT System Dynamics Group Working Paper D-4354.

Urban, Glen *et al.* 1986. Market Share Rewards to Pioneering Brands: An Empirical Analysis and Strategic Implications. *Management Science* 32(6).

Urban, Glen L., John S. Hulland, and Bruce D. Weinberg 1993. Premarket Forecasting for New Consumer Durable Goods: Modeling Categorization, Elimination, and Consideration Phenomena. *Journal of Marketing* 57(2) 47-63.

Urban Glen L., Bruce Weinberg and John R. Hauser 1994. Premarket Forecasting of Really New Products. Cambridge, Mass.: MIT Sloan School of Management Working Paper WP-3689.

Wernefelt, Birger 1984. A Resource-Based View of the Firm. *Strategic Management Journal* 5, 171-180.

Wolstenholme, Eric F., Abdul-Satter Al-Alusi 1987. System Dynamics and Heuristic Optimization in Defence Analysis. *System Dynamics Review* 3(2), 63-80.