Go Back

Product Development Strategy for the High-tech Startup: A System Dynamics Post-Mortem

Corey Lofdahl
SAIC
Burlington, MA USA

Clyde Lofdahl
LCG
Littleton, CO USA

Prepared for presentation at
The 19th International Conference of the System Dynamics Society
Atlanta, GA – July 2001

Abstract

This study reviews the case of a failed high-technology product development effort, the Multi-Chip Module (MCM) project. Specifically it reviews the strategic decisions that led to the project's demise, and using simulation it constructs an alternate strategy that would have proven more successful. The difference between the two strategies centers on the relative timing between product development and marketing efforts. The first strategy, which was actually pursued, stresses acquiring customers early and learning about the production process while filling actual orders. The second, hypothetical strategy centers instead on learning about the production process first and then, only after production proficiency has been achieved, shifting attention and resources to the task of customer acquisition. Simulation-based scenario analysis demonstrates the second strategy's superior performance as measured by the MCM project's path to profitability.

Corey Lofdahl holds degrees in electrical engineering, computer science, and international relations from the University of Colorado, Brown University, and MIT. He is a branch manager with SAIC's Simulation and Information Technology Operation.

Clyde Lofdahl holds degrees in mechanical engineering and business administration from the University of Idaho and Harvard University. He is a management consultant and investor.

Introduction

Many businesses have shut their doors since the dot-com bubble burst in April 2000, so many in fact that one popular web site – the dot-com deadpool – exists simply to mark (and sometimes gloat over) their passing. While some failed companies have undoubtedly been the unwitting victims of collapsing capital markets, others were fatally flawed from the beginning and should never have been conceived, funded, or staffed. Thus nascent companies come in three flavors: those that are (1) perfectly conceived, (2) flawed and fixable, and (3) flawed and unfixable. For the purposes of this paper, let us assume that the majority of new companies are of type (2) or (3) – that is, they are flawed – so the question then turns to whether or not they are fixable. This is usually determined by funding and staffing a company and then letting it compete in the marketplace where the staff "fixes" the company in real-time. If the company makes a profit, then it was fixable; if it doesn't, then it wasn't.

Of course the staff, committed as they are to the company's success, always believe their situation is salvageable given enough time and money. This is why pulling the life-support on dying companies is so difficult. This paper explores the application of simulation generally, and system dynamics specifically, to the crafting of strategies that increase the likelihood of success for new commercial ventures. Simulation is offered for this task because fixing key policy problems with a model on a laptop in an afternoon is much less costly and painful than fixing them with real employees and real money in real-time. In considering the problem of crafting corporate strategy, it is easy to focus on the internal workings of the firm. However, it should be recognized that firms exists in markets where they compete with other firms. In this manner, the crafting of internal policy blurs into the formulation of competitive strategy. Whether simulation is used to determine a company's internal workings or its strategy across a range of competitive and market conditions, the key is to avoid costly problems and to maximize the company's chances of profitability and survival. If solutions cannot be found for likely problems, then the company should be considered unfixable and should go unfunded.

Working in the realm of high-technology complicates matters by increasing the unfriendly factors of complexity, risk, and uncertainty. High-tech product development, by its very nature, implies a diminished history on which decisions might be based. Every time, it seems, feels like the first time. And while the decisions are difficult, the stakes are high: the careers of the project team and millions of investment dollars are on the line. Given the paucity of data and the high stakes associated with high-tech product development, it makes sense to spend extra time on those critical decisions, policies, and strategies that will ultimately determine the fate of the firm.

Focusing of critical decisions is however not as easy as it sounds though. How does one know that a particular decision is critical? In a high-tech research and development (R&D) environment, the importance of seemingly simple decisions is often revealed only later, and seemingly important questions are later revealed to be of only minor significance. Recognizing this, it is similarly impossible to give high priority to all decisions as "analysis paralysis" inevitably results. Even after acknowledging these observations as true, it still makes sense to spend extra time on certain key questions and problems. This study presents a case study of a high-tech, product development effort that pursued a strategy that ultimately led to its demise. It does so using system dynamics simulation to provide an example of what a crucial, initial, and incorrect strategy looks like. An alternative strategy is then developed that shows how a different set of decisions might have yielded a superior result. This study is organized into three parts: first, a background of the motivations and facts underlying the product development example is provided; second, a model is developed that captures the relevant relationships of the product development process; third, a dynamic analysis is performed with an eye towards developing a better guiding strategy that leads to a better final outcome. Concluding thoughts are presented in the final section.

Part 1 – Product Background

The project took place at a Silicon Valley firm from 1984 to 1992 and centered on the development of Multi-Chip Modules (MCM), products that concentrate computational capability in a smaller area than had been possible until that time. The idea was to put a whole computer on a single silicon wafer, but the parent company had little previous experience in this area. In fact, the parent company's primary expertise was in wire, cable, and electrical interconnection products, and it had watched the integrated circuit (IC) and computer industries grow up around it. It had never participated in either of these important and lucrative markets; the MCM project was to provide its entry.

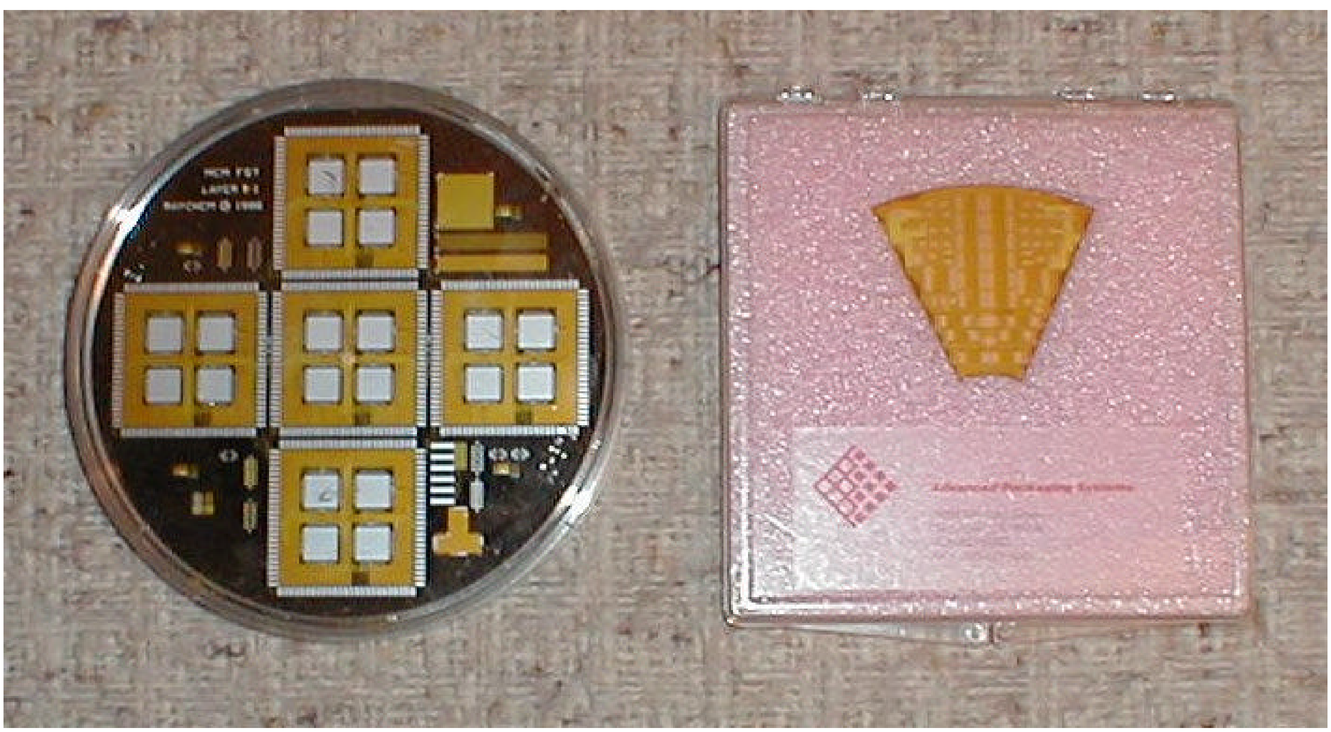

Figure 1. Example Multi-Chip Module (MCM) Products

The development strategy was to use established IC manufacturing equipment and methods to produce MCMs, which were in effect a printed circuit boards shrunk down to 1/10 scale. The MCMs were used to connect ICs and associated discrete electrical components -- i.e., capacitors, resistors, and other devices. At the time, printed circuit board standard conductors were 10 mils, or 10 thousandths of an inch wide. Conductors in the MCM were intended to be 1 mil or one tenth as wide, a significant technical leap. However, the project was launched with an initial development team that did not include IC expertise. The possibility of hiring an expert was discussed and several were interviewed, but their salary demands were substantially higher than the company's salary structure allowed. Thus the development team was formed entirely with in-house personnel.

The product development philosophy of the company for many years had been, "Take a big order and learn to make it on the run!" This strategy had proven successful on many occasions and was a matter of pride within the company. Time spent in R&D by a small team was viewed with suspicion. Management's view was, "There is nothing like a big customer yelling for his order to focus effort and speed development!"

The first MCMs produced were test pieces that proved manufacturing capability, but within 90 days a senior company executive booked an order with a large computer manufacturer to, "give the MCM team something to chew on." It took over 18 months to deliver this order. During this time, additional orders were booked with major computer manufacturers and large military electronic contractors in the US and Europe. Orders were not the problem; deliveries were.

The foremost delivery problem was always manufacturing yield, and the essence of the yield problem centers on geometry. If a 5-inch diameter wafer is used in IC manufacturing, it will product fifty 0.5 inch square ICs. If there are 5 defects (i.e., electrical opens or shorts in the circuitry) in the wafer, there can be up to 5 defective ICs. 45 good ICs out of 50 produced means a 90% yield, which is quite acceptable. If a 5-inch diameter wafer is used in MCM manufacturing, it will produce five 1.5 in square MCMs. If there are the same 5 defects in this wafer, there can be up to 5 defective MCMs, for a yield of zero. So when wafer (IC-like) manufacturing methods are used, the larger the final product, the smaller the yield.

The following list outlines the number of MCM products per wafer for customer orders during the MCM project:

1 per wafer	13% of orders
2 per wafer	41% of orders
3 per wafer	22% of orders
4 per wafer	14% of orders
5 or more per wafer	10% of orders

The “book a big order and learn to make it on the run” philosophy was ruinous for the MCM project. The pressure from customer orders with firm delivery dates meant that there was never sufficient manufacturing capacity to produce development orders from which the group could learn to improve yields. During the first year manufacturing yields were about 5%. Eight years later they were in the 20-25% range. To break even financially required yields of 50%, and this was never achieved.

The high outflow of cash or “burn rate” from the large-scale production facility and its attendant multi-shift crews meant that top management was always concerned. Any suggestion of scaling back until yields could be improved was quickly dismissed. This MCM facility was known throughout the IC and computer business worlds, and scaling back the effort would have proven embarrassing to the parent company. After eight years of heavy losses, the MCM facility was shut down in 1992.

Part 2 – The Model

Years afterwards, the project's facility manager was introduced to system dynamics, and the Product Development (PD) model resulted from the initial modeling effort. The PD model, shown below in Figure 2, was created to explore, in retrospect, how the project might have been better understood, funded, and managed.

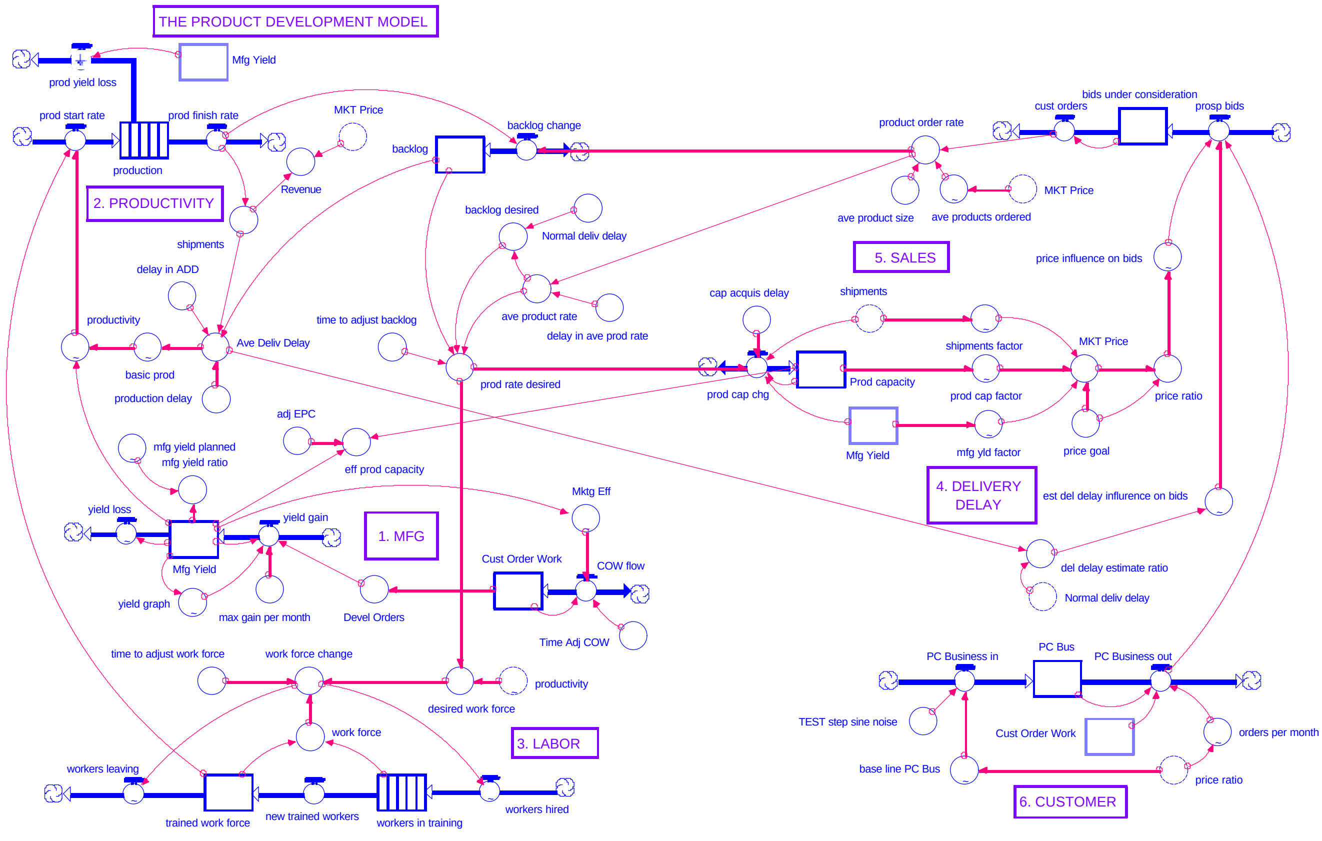

Figure 2. The Product Development (PD) Model

In keeping with the traditional strengths of system dynamics modeling, the actual physical processes and behaviors of the MCM manufacturing effort are modeled rather than the abstract financials. The model is comprised of six sections: (1) manufacturing, (2) productivity, (3) labor, (4) delivery delay, (5) sales, and (6) customer, each of which are discussed below.

Section 1. Manufacturing

This section contains the key relationship in the model – that between the *manufacturing yield* (Mfg_Yield) and *customer order work* (Cust_Order_Work) stocks. The relationship between the two is simple: the more customer order work, the slower the increase of manufacturing yield. This relationship requires some explanation. To begin, recognize that the MCM manufacturing process is complex, which means there are many process parameters that can be varied, and the yields change from batch to batch even when the parameters are ostensibly kept the same. Learning takes place slowly in such an environment as knowledge and experience build gradually. Experience accumulates more quickly if those responsible for manufacturing are given the chance to `play' with the process – that is, if they are given the chance to vary the parameters and experiment with unusual configurations in a low-pressure, consequence free environment. Conversely, experience accumulates more slowly if pressure is placed upon the manufacturing team per the theory of *bounded rationality*[*]. In times of stress people tend to rely on few and certain data. Here, if the pressure is on to deliver an order, people naturally tend to concentrate on what worked before rather than risk an even lower yield. So the more customer orders, the more stress, the slower the learning process, and the lower the yield. Upon recognizing the importance of manufacturing yield, the strategy should have been not to, "book big orders", but rather

[*] For more on this topic, see John D.W. Morecroft's 1983 article, "System Dynamics: portraying bounded rationality," in the **International Journal of Management Science 11**(2), pp. 131—42.

to keep the project small and minimize market contact until production capability (i.e., manufacturing yield, product design, and other skills) had been learned and demonstrated.

Section 2. Productivity

If the previous section concerns production “know how,” this section concerns its application to actual production. Thus manufacturing yield is used to generate both usable and flawed product. Usable or finished product is shipped, which when multiplied by price generates revenue. Costs are then subtracted from revenue to determine profit.

Section 3. Labor

This section covers the normal hiring and training of workers in accordance with the production demands as represented by order backlog. As demand increases, it also drivess new investment in production capacity.

Section 4. Delivery Delay

A basic axiom of business is that low price leads to more customers and high price to fewer customers. Likewise, timely deliveries have a positive influence on customers and late deliveries a negative influence. This section tracks delivery delay and relates it to the sales process as a function of backlog – the larger the backlog, the longer the delivery delay.

Sections 5 and 6. Sales and Customer

These sections concentrate on the process of acquiring customers and generating business from them. Competition was never a serious factor for MCM – the main problem was delivering product in an efficient manner. These sections include several influences including price and delivery delay on customers. Product price was a major issue in the MCM project, which is captured by the key price factors of shipment size, production capacity, and manufacturing yield.

Profit and Loss

As noted previously, revenue is generated in the productivity section when product is shipped and billed. A separate section, not shown in Figure 2, calculates profit and loss essentially by counting the costs of production including materials, labor, marketing, technical staff, overhead, and intellectual property. Profit is calculated by subtracting these costs from revenue.

The resulting model, while not a perfect representation of reality, provides a reasonable measure of correspondence to the MCM project as it was actually run. In the next section, the model is used to explore and improve the operational strategy that eventually led to the demise of the MCM project. Additionally, several alternative policies are proposed, implemented, and evaluated in an effort to craft a more workable strategy that might have led to a more positive outcome.

Part 3 – Crafting Alternative Strategy

The key variable for this discussion is *marketing effect*, a graphical function that relates the effort and expense necessary to identify and acquire new customers with the technical state of the production process as measured by *manufacturing yield*.

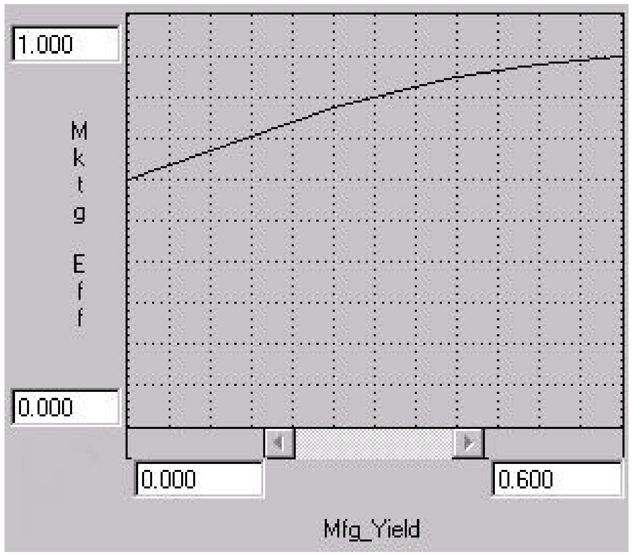

Figure 3. Actual MCM Policy – Acquire Customers Early

Figure 3 captures the product development strategy as it was actually practiced. The horizontal or X-axis measures the *manufacturing yield* in terms of percentage of usable product. The scale's maximum range is 60%, with postulated breakeven being about 50% and the best achieved yield being about 25 to 30%. The vertical or Y-axis shows the output of *marketing effect*, the effort expended to acquire customers using a dimensionless measure ranging from 0 to 1 with 0 being the minimum effort and 1 the maximum. Note that the graph depicts fairly consistent and high effort to acquire customers regardless of the value of *manufacturing yield*, starting off at a value of 0.6 and increasing to a value of 0.9. In this fashion, the graph captures management's policy of booking orders early and learning on the fly.

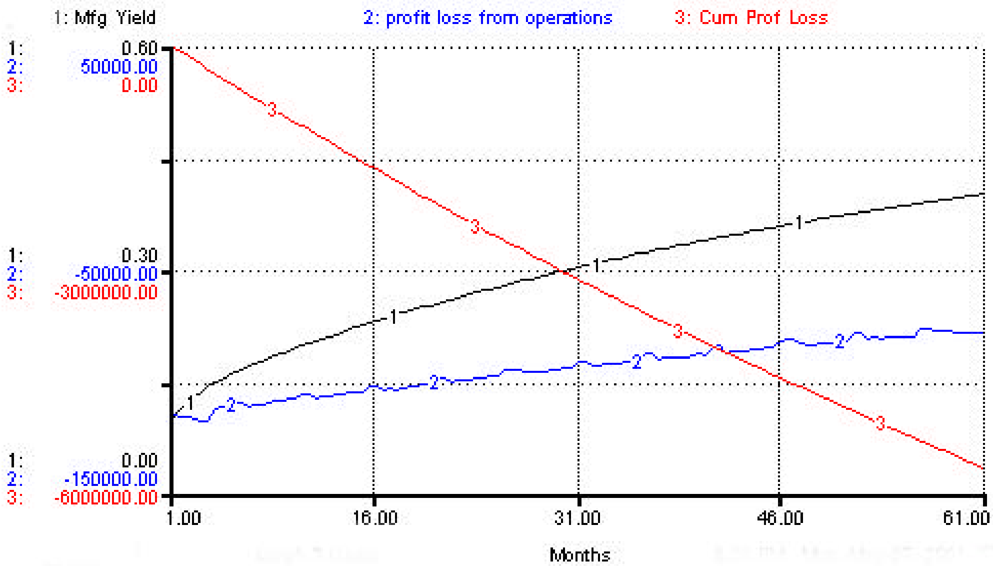

Figure 4. Actual MCM Performance

Figure 3 demonstrates the consequences of the pursued strategy, consistent losses on the order of $100,000 per month. Driving this consistent loss of money is the presence of customer orders that effectively prohibit experimentation and learning about the production process – such learning would have allowed the manufacturing yield to improve. In this scenario, while *manufacturing yield* increases, it never reaches the 50% mark deemed necessary for profitability.

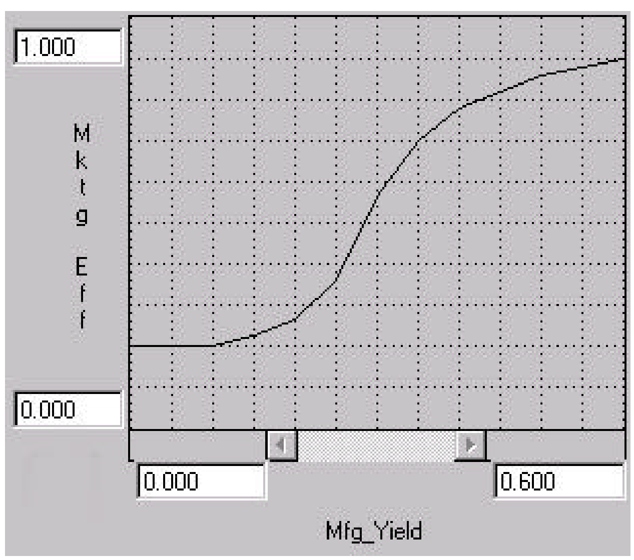

Figure 5. Acquire customers after production understood

Figure 5 graphically describes the strategy of expending serious energy acquiring customers only after production yields approach a close to profitable level, in this case about 35%. Below that, only token customer acquisition efforts are undertaken.

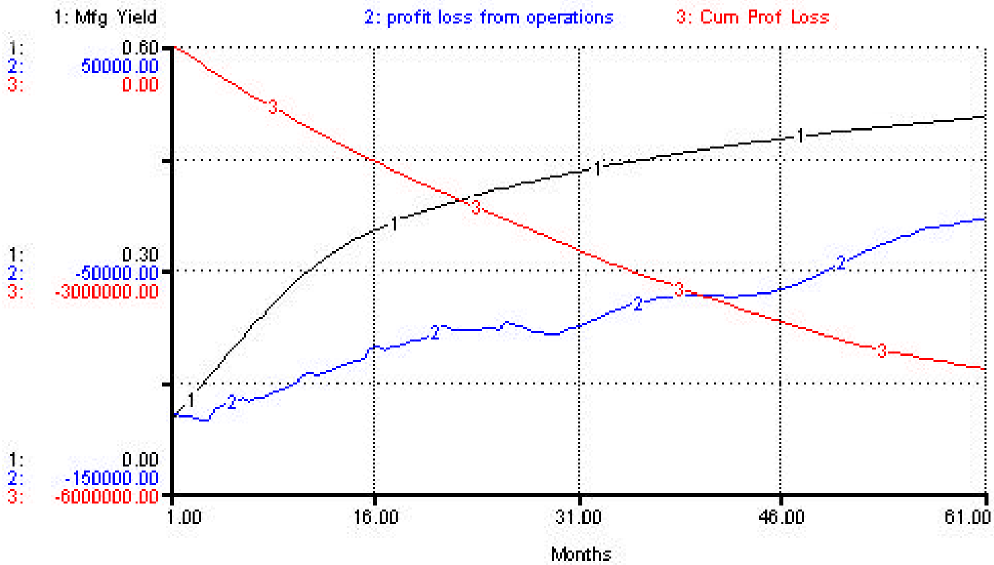

Figure 6. Performance of delayed customer acquisition

Figure 6 shows the results obtained by implementing the graphical policy presented in Figure 5. Note specifically an improved rise of *manufacturing yield*, which leads to improved financial performance in the form of lower losses. So while profitability remains elusive in this scenario, the MGM project's financials are at least headed in the right direction.

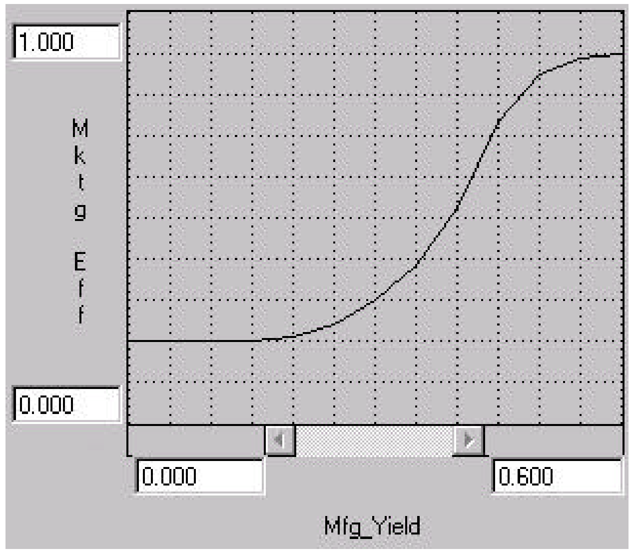

Figure 7. Acquire customers after production well understood

If waiting a while before acquiring customers is good, how about waiting even longer? Figure 7 graphically describes the strategy of expending serious energy to acquire customers only after production yields approach an even more profitable level, in this case about 45%. Once again, only token customer acquisition efforts are undertaken below that.

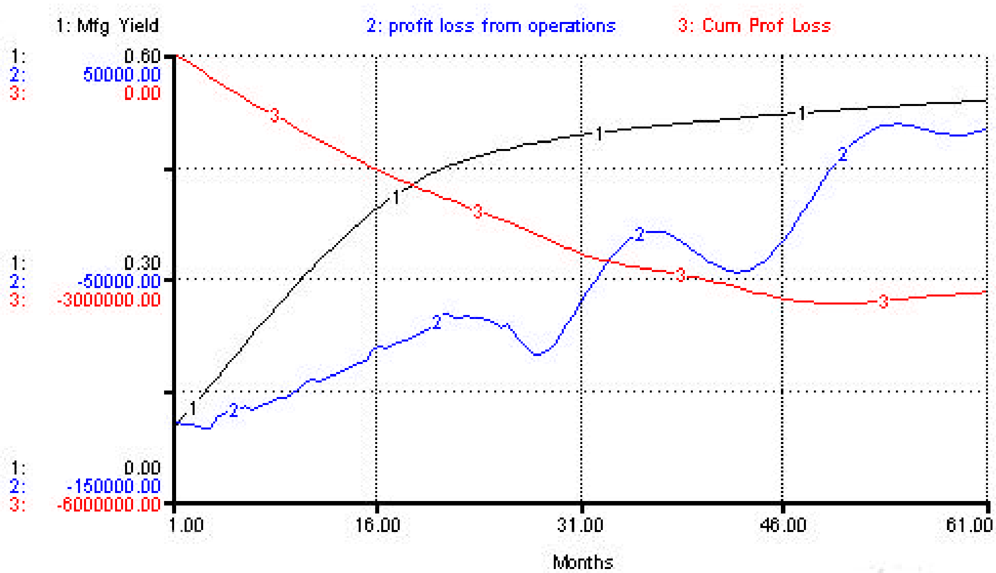

Figure 8. Performance of customer acquisition delayed still further

Figure 8 shows the results obtained by implementing the graphical policy presented in Figure 7. Note an even sharper rise of *manufacturing yield*, which once again leads to improved financial performance in the form of profitability around month 50. Had such performance actually been achieved, the MGM product development project would have been deemed a success. Note also the oscillations of the profit and loss curve. These stem from costs on the operational side, most notably the salaries of the technical staff. Staff levels oscillate due to the time delays associated with acquiring and training staff in response to product backlogs.

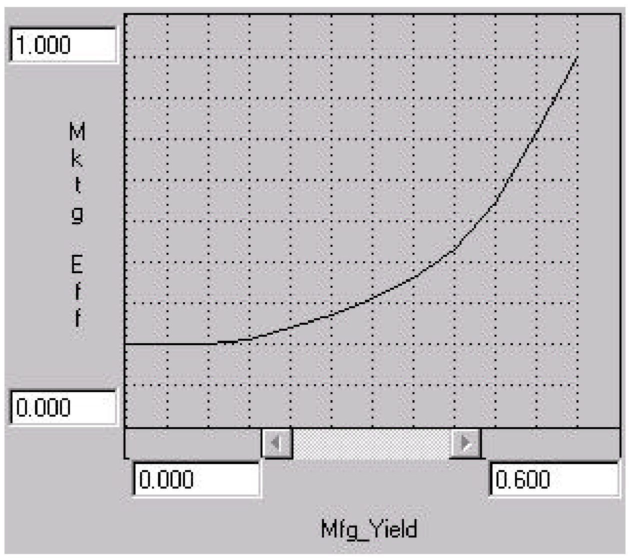

Figure 9. Gradual acquisition of customers as production is understood

One aspect of the policies presented in Figures 5 and 7 is their troubling binary nature. That is, there is a tendency to work at a reduced level until *manufacturing yield* reaches a previously defined target. Often however, such targets are ill specified and difficult to define. Figure 9 shows a policy curve with a different shape denoting the gradual increase of marketing effort as manufacturing yield improves. This seems a more plausible and sustainable policy.

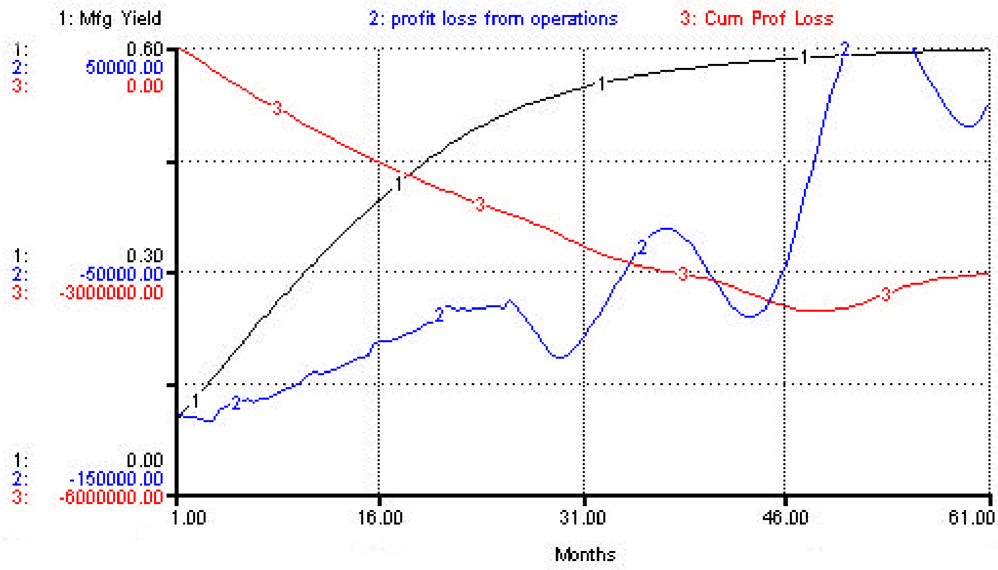

Figure 10. Performance of gradual acquisition of customers

Figure 10 shows the results obtained by implementing the graphical policy presented in Figure 9. Even though *manufacturing yield*, climbs at about the same rate as Figure 8, profitability is much higher and the response oscillates similarly. The key point is that even though the Figure 9 policy curve is shaped differently than Figures 5 and 7, it pushes off marketing effort until the production process is better understood, which is the key to profitability in the MCM domain.

Conclusion

In reviewing the results of the product development model, the question of whether or not success was "dialed in" is addressed. That is, the conclusion of doing development first and then marketing as opposed to development and marketing simultaneously is, in retrospect, almost self-evident. Moreover, the algorithm used to generate increases in manufacturing yield are oversimplified, primarily because capturing the complexity of the MCM production process is impossible and would not have been worth the effort even if it were possible.

The benefits of PD model are more subtle. First, the model shows in an operational fashion the types of losses that can be incurred should the manufacturing problems go unsolved for a long period, as actually occurred. Similar models could have been run using a spreadsheet model, though they would have lacked the operational flavor of this model as well as its complex feedback relationships. Moreover, the model shows the complex, causal connections that link marketing effort with the production effort. Without such explicit connections, one might assume that the two efforts are independent – that decisions made in marketing are unrelated to production or vice versa. One might also assume, as did the senior executives of the MCM project's parent company, that marketing effort would spur the development effort. This proved not to be the case as the drive to deliver product to the customer effectively stifled learning about the manufacturing process. This in turn kept manufacturing yields at low and unprofitable levels.

The product development model shows that a strategy focusing on customer orders at the expense of research and development lead to high burn rates, dissatisfied customers, and a drawn out learning curve. In light of this, it seems that a strategy predicated on learning first and selling second would have lowered the burn rate, retained the goodwill of customers, and accelerated learning. In making this distinction, recall that the project continued for eight years, spent millions of dollars, and consumed the time of many bright people. While this lesson might seem obvious or unique, such stories are all too common in business. The problem was not that the people were insufficiently dedicated, informed, or intelligent, but that there was no framework to integrate and make sense of the disparate and diffuse data. The system dynamics model developed and referenced herein provides a framework that delivers value not because it uncovers new information but because it orders and interprets available information in such a way as to provide new insights.