

Assessing the Benefits of Integrated Vehicle Life Cycle Planning

Wendy Lu Xu
Barry L. Nelson
Northwestern University

Wallace J. Hopp
University of Michigan

Jonathan H. Owen
General Motors R&D

February 15, 2009

Abstract

Many decisions are involved in managing the vehicle life cycle. These include product portfolio planning (which models to launch at which times), plant assignment (which plants to use to produce each vehicle) and production allocation (how much of each vehicle to produce in each plant). Since these decisions are complex and made at different points in time they are typically decoupled in practice. But they clearly impact one another, since portfolio decisions constrain assignment decisions, which in turn constrain the allocation decisions. We present a Markov Decision Process model, which maximizes expected long-run average profit, to assess the value of integrating portfolio planning and plant assignment. Our analytical and numerical results suggest that decoupling these decisions can lead to substantial loss of profit, particularly when plant capacity utilization is moderate and flexible tooling makes it significantly cheaper to introduce new models into plants.

Keywords: Product Portfolio Planning, Product-to-Plant Assignment, Markov Decision Processes.

1 Introduction

Introduction of new and updated vehicles is one of the most vital activities of an automotive company. Without a steady stream of offerings that are perceived by the buying public as “new” and “fresh” within a segment, a firm cannot remain competitive. No amount of manufacturing efficiency or quality control can offset the disadvantage of having “stale” products. As a result, all automotive companies devote a tremendous amount of managerial attention at all levels, including in the executive suite, to the problem of planning the evolution of their vehicle portfolio over time. A central issue in this planning process is when to refresh or replace models.

On the surface, this looks like a marketing problem. The firm must be sensitive to customer tastes (e.g., “freshness” may be more important to buyers of a stylish sports car than to buyers of a utilitarian pickup). It must also take into account the pace of innovation by the competition. With these, the firm can predict the decline in demand for a model as it ages and balance this with the cost of design and tooling needed to launch a refreshed model to determine an appropriate refresh cycle.

But this problem goes well beyond marketing. Decisions about when and how often to refresh/replace models impact the workload on engineering resources, as well as the options for where the vehicles can be produced. Decisions about the assignment of vehicles to plants in turn affect subsequent decisions about how many vehicles of each type to produce in each plant.

To cope with the scale and complexity of this multidimensional planning problem, firms generally decouple the portfolio problem from the downstream decisions it impacts. That is, the portfolio problem is solved primarily as a marketing problem, with some checks on the resource implications of the decisions. For instance, to ensure that engineering resources are not overtaxed, the firm may restrict the number of major and minor refreshes in a given year to not exceed a specified number as a rough approximation.

But accounting for plant assignment considerations is more subtle because the life cycle cost of refreshing can depend on where the vehicle is produced, as well as where other vehicles are produced (i.e., because plant capacities may be shared by multiple vehicles). It would be simple to assume that a refreshed model will be produced in the same plant(s) as the old model it replaces. But this will not always be the best solution due to plant closings to consolidate capacity, new plant construction to add capacity or technology, improvements in flexibility that make it feasible to produce more models in a single plant and other business trends. Consequently, portfolio planners do not always know where a model will be produced at the time a refresh is planned. Therefore, they typically assume average tooling and production costs to evaluate the economics of replacement decisions. The actual costs will not be revealed until after the assignments are made and production is scheduled.

Because portfolio decisions influence plant assignment decisions, which in turn influence produc-

tion allocation decisions, decoupling may lead to suboptimal decisions. This presents the following questions:

1. How large can the shortfall in expected profit be from decoupling portfolio planning from plant assignment and production allocation?
2. Under what conditions is decoupling the decisions likely to lead to large profit shortfall, and when is such decoupling likely to result in good performance?
3. How can planners incorporate implications for the downstream plant assignment and production allocation decisions into the portfolio planning process?

In this paper, we address the first two questions in some detail and provide the basis for a method to address the third. To do this, we develop a Markov Decision Process (MDP) model that integrates portfolio planning, vehicle assignment and production allocation.

There is a considerable body of literature on new product development/R&D project portfolio planning which addresses the problem of allocating limited resources to products/projects in a way that maximizes the return, provides balance, and aligns with the strategy of the enterprise. Heidenberger and Stummer (1999) gave a comprehensive review of approaches to quantitative modeling for project selection and resource allocation, which include benefit measurement methods, mathematical programming, and other techniques. Other methods include graphical and charting techniques. Cooper et al. (1999) reported the portfolio management practices and performances of 205 U.S. companies. Dickinson et al. (2001), Loch and Kavadias (2002) and Stummer and Heidenberger (2003) addressed the portfolio problem over multiple periods. But none of these computed the return for new products/projects by taking into account downstream manufacturing decisions.

The way we model the portfolio decision process has some similarity to the work on machine maintenance problems. Surveys of the maintenance/replacement literature include Barlow et al. (1965), MacCall (1965), Pierskalla and Voelker (1976), Sherif and Smith (1981), Valdez-Flores and Feldman (1989) and Wang (2002). Decreasing demand over time in our model is analogous to deteriorating system performance in maintenance models. Vehicle model refresh, which is aimed

at increasing demand, is similar to maintenance, which is intended to increase performance. According to the classification of maintenance models in Pierskalla and Voelker (1976) and Wang (2002), our portfolio decision problem is similar to a discrete-time maintenance model with a single maintenance action, complete information and no inspection, which can be formulated as a Markov Decision Process with a failure limit policy. Among the machine maintenance studies, Sloan and Shanthikumar (2000) is most similar to our problem in that it models maintenance and production schedules simultaneously. The difference between our model and their work is that they decide either to perform maintenance or to produce one of several alternative products, while in our model we decide whether to refresh a vehicle and where to assign it at the same time, with production planning as a second-stage decision. Also our measure of performance includes revenue, which is determined by production quantity and is not a decision variable in their model.

Plant assignment and production planning is a complicated problem that must take into account timing of product introductions, limited tooling of each product, limited capacity of plants, and other issues. Inman and Gonsalvez (2001) described a tool for optimizing the allocation of vehicle products to plants considering the capacity constraints of individual lines, body shops and plants, with the objective of minimizing lost sales, equalizing utilization among plants, and maximizing the amount of inter-plant chaining. Alden, Costy and Inman (2002) presented a decision support tool to assess the cost of plant modifications needed to build a new vehicle product. In our model, plant assignment is a decision in the MDP associated with tooling costs, while production planning is formulated as a linear optimization problem.

The remainder of the paper is organized as follows: Section 2 formulates two MDP models for a two-plant, two-product system, one for the case where the portfolio planning decision is integrated with the plant assignment and production allocation decisions, and one where they are decoupled. This section also gives basic structural results for the optimal policies. Section 3 presents a numerical analysis that investigates the size and causes of large shortfall between the results of the integrated and decoupled models. Section 4 summarizes our conclusions, which are:

1. The shortfall in expected profit from failing to integrate plant assignment and product allo-

cation considerations into the portfolio planning process can be large (e.g., more than 50%). However, such a shortfall is not assured, since it is always possible to “get lucky” by having the decoupled solution match the integrated solution. Instead, because we cannot predict when this will occur, our results suggest that there are conditions under which the *risk of a bad decision* from using a decoupled model is large.

2. The risk of a large profit shortfall from using a decoupled model to address the portfolio planning problem is greatest when: (a) plant capacity utilization is moderate, and (b) it is significantly cheaper to introduce new models in plants equipped with flexible tooling than in plants equipped with dedicated tooling.
3. Although it is valuable to integrate the portfolio planning problem with the plant assignment and production allocation problems, doing so with an exact optimization model, such as the MDP we formulate here, is only feasible for very small systems. It might be possible to disaggregate the full problem, which for large auto makers involves dozens of products and 10 or more plants, into separate segments consisting of only a few products and plants within each, in order to make our MDP model viable. However, to incorporate effects that cut across segments (e.g., constraints on engineering and/or marketing resources), it may be necessary to embed these in a simulation-based optimization model. The results we presented here can be viewed as a first step toward such a large-scale model.

2 Model Formulation

We begin by considering a set of vehicles that are currently being produced in a number of plants. The objective is to maximize the expected profit in the future. We focus on decisions at three stages of the vehicle life cycle that affect future profits: portfolio planning, plant assignment and production allocation. At the portfolio planning stage, decisions are made whether or not to refresh an existing vehicle model in a given year. When a model is refreshed, it is assigned to one or more plants for production. Assigning a model to a plant incurs a retooling cost, whose magnitude depends on what is currently being produced in the plant. In general, it costs less to retool a

plant to produce a new model of the same vehicle than to retool it to produce an entirely different model. After refreshing, we assume that demand for the model increases in a stochastic manner, but then decreases at a deterministic rate until the next refresh. We further assume that demand is known before production quantities at each plant must be determined. These are constrained by plant capacity, which includes the option of overtime. Once the refresh, assignment and production allocation decisions have been made for a model, profit is determined by revenue minus production, overtime and tooling costs. Revenue is computed by assuming that sales equals production. That is, we are not considering inventory or shortage costs.

Decisions at portfolio planning and plant assignment stages can be made in either an integrated or decoupled manner. In either case, however, we assume that plant assignment and production allocation decisions are made in an integrated fashion, since it is a comparatively simple matter for manufacturing engineers considering plant assignments to analyze production plans for various demand scenarios. For the integrated case, future plant assignment (and production allocation) decisions are taken into consideration when making portfolio planning (refresh) decisions. For the decoupled case, plant assignment (and hence production allocation) information is not available to portfolio decision makers. Hence, portfolio decisions are made using historical data about tooling and overtime costs. In either case, once plant assignments are known, we assume that production allocations are optimized via linear programming.

To capture these decision mechanics in a model we focus on a two-product two-plant system, later extending to a three-product, two-plant model. To facilitate our presentation, we begin by describing the model of the integrated case and then adapt this model to represent the decoupled case.

2.1 MDP Formulation for a 2-Product 2-Plant System: Integrated Decisions

We label the two vehicles A and B, and the two plants 1 and 2. We formulate the integrated decision problem for a two-product two-plant system as a Markov Decision Process (MDP) with an objective to maximize long-run average profit. To do this, we define the following:

- *State space* \mathcal{S} with states (d_A, d_B, y_A, y_B) ; $y_A, y_B \in \{1, 2, 1\&2\}$ are the current plant assign-

ments for product A and B, indicating the product is assigned to plant 1, 2 or both 1 and 2; $d_A, d_B \in \{1, 2, \dots, M\}$ are the demand levels for product A and B in the current period. We use a function $f(\cdot)$ to map a demand state into a demand level and assume $f(\cdot)$ is increasing in this state.

- *Decision epochs* are the beginning of every year $t = 1, 2, \dots$
- *Action space* \mathcal{A} with action (a_A, a_B) ; $a_A, a_B \in \{0, 1, 2, 1\&2\}$ are the refresh and assignment actions for products A and B, respectively, where 0 represents “no refresh”, 1 denotes “refresh to plant 1”, 2 denotes “refresh to plant 2” and 1&2 indicates “refresh to both plants 1 and 2”.
- *Reward* equals net revenue minus tooling costs. Assigning a newly refreshed vehicle model to a plant incurs a tooling cost $c_{tool}(y_A, y_B, a_A, a_B)$, which depends on the current product-plant assignment (y_A, y_B) , and assignment action (a_A, a_B) . Define x_{A1} as the production quantity of product A in plant 1, and define x_{A2}, x_{B1} and x_{B2} analogously. Then $\mathbf{x} = (x_{A1}, x_{A2}, x_{B1}, x_{B2})$ represents a production plan. We can find the production plan that achieves the maximum net revenue, $\pi(d_A, d_B, y_A, y_B)$, by solving the following linear programming model \mathbf{P}_π :

$$\pi(d_A, d_B, y_A, y_B) = \max_{\mathbf{x}} \left\{ m_A(x_{A1} + x_{A2}) + m_B(x_{B1} + x_{B2}) - c_{OT,1}(x_{A1} + x_{B1} - K_{reg,1})^+ - c_{OT,2}(x_{A2} + x_{B2} - K_{reg,2})^+ \right\} \quad (1)$$

subject to:

$$x_{A1} + x_{A2} \leq f(d_A), \quad x_{B1} + x_{B2} \leq f(d_B), \quad (2)$$

$$x_{A1} + x_{B1} \leq 1.5K_{reg,1}, \quad x_{A2} + x_{B2} \leq 1.5K_{reg,2}, \quad (3)$$

$$x_{i1} = 0 \quad \text{if} \quad y_i = 2, \quad i = A, B,$$

$$x_{i2} = 0 \quad \text{if} \quad y_i = 1, \quad i = A, B, \quad (4)$$

where m_A and m_B are contribution margins for products A and B; $c_{OT,1}$ and $c_{OT,2}$ are unit overtime costs for plants 1 and 2; $K_{reg,1}$ and $K_{reg,2}$ are regular capacities of plants 1 and 2, $1.5K_{reg,1}$ and $1.5K_{reg,2}$ are total (regular plus overtime) capacity of plants 1 and 2, with the assumption that overtime capacity is half the regular capacity (because automobile factories

usually run two shifts for regular production and an extra overtime shift when necessary). Constraints (2) avoid producing more than needed. Constraints (3) guarantee that the total quantity of production of a plant does not exceed its capacity. Constraints (4) define the product-plant assignment. Finally, we note that we do not consider inventory or shortage costs, since we assume demand is known when we make production allocation decisions.

To model the change in demand, we assume that immediately after a model is refreshed, demand follows a binomial distribution $\text{bin}(M, p)$. A large p value implies that after a refresh, demand is very likely to increase to a high level. We use $P(\cdot)$ to denote the probability function. For the years in which a model is not refreshed, we assume demand decreases by one level, i.e., from d_A (d_B) to $d_A - 1$ ($d_B - 1$). We define $V(d_A, d_B, y_A, y_B)$ as the value function in state (d_A, d_B, y_A, y_B) , and g as the optimal average profit per year, and write the optimality equation of the MDP model as

$$\begin{aligned} & V(d_A, d_B, y_A, y_B) + g \tag{5} \\ = & \max_{(a_A, a_B) \in \mathcal{A}} \left\{ \pi(d_A, d_B, y_A, y_B) - c_{\text{tool}}(y_A, y_B, a_A, a_B) + \right. \\ & \left. \sum_{j_A=1}^M \sum_{j_B=1}^M \tilde{P}(j_A, a_A) \tilde{P}(j_B, a_B) V(D(d_A, j_A, a_A), D(d_B, j_B, a_B), Y(y_A, a_A), Y(y_B, a_B)) \right\} \end{aligned}$$

where

$$\begin{aligned} \tilde{P}(j, a) &= \begin{cases} P(j) & a = 1, 2, 1\&2, \\ 1 & a = 0, \end{cases} \\ D(d, j, a) &= \begin{cases} j & a = 1, 2, 1\&2, \\ \max(d - 1, 1) & a = 0, \end{cases} \\ Y(y, a) &= \begin{cases} a & a = 1, 2, 1\&2, \\ y & a = 0. \end{cases} \end{aligned}$$

We can now characterize the structure of the optimal policy for the integrated-decision scenario. Omitted proofs can be found in Online Appendix B. We begin by justifying the MDP solution.

Theorem 1. *There exists a stationary optimal policy for the MDP defined by (5) that has a constant gain rate. Moreover, the corresponding value iteration algorithm converges.*

To characterize the structure of the optimal policy, we first require the following technical result.

Proposition 1. *The value function $V(d_A, d_B, y_A, y_B)$ is non-decreasing in d_A (respectively d_B).*

Now we present the main result of the subsection.

Theorem 2. *If the action “no refresh for product A (respectively B)” is optimal for state (d_A, d_B, y_A, y_B) , then it is optimal for state $(d_A + 1, d_B, y_A, y_B)$ (respectively $(d_A, d_B + 1, y_A, y_B)$) for all $1 \leq d_A \leq M - 1$ (respectively $1 \leq d_B \leq M - 1$).*

Theorem 2 says that the optimal policy follows a one-dimensional threshold structure in demand level. As we show in Appendix B, that Theorem 2 holds for an N -product K -plant system and general probability function $P(\cdot)$, not just the 2-product 2-plant system and binomial distribution.

2.2 MDP Formulation for a 2-Product 2-Plant System: Decoupled Decisions

When portfolio and plant assignment decisions are made separately, portfolio decision makers may look at historical data to get an idea of how much it will cost to retool a factory for a new model, even though they do not know the specific factory that will be retooled. We model the decoupled decision making case as a three-stage process. First, average net revenue and tooling costs are estimated assuming a fixed refresh cycle; this imitates the real process of using historical data to obtain averages. To be conservative in our comparison of integrated and decoupled decision making we make the process as favorable as possible for the decoupled case by getting good average numbers for net revenue and tooling costs. In practice, such numbers would be estimated from historical averages. But here we solve an off-line MDP with fixed refresh cycle length for optimal plant assignment and production allocation, and compute average tooling costs under the optimal policy. Then, the portfolio decisions are made using the average net revenue and costs from this computation. Finally, we solve the plant assignment problem as an MDP with a restricted action space determined by the portfolio decisions. We describe this three-stage process in detail in Online Appendix A. We can demonstrate similar results to those of Theorem 1 for the MDP of each stage, which we state as Theorem 3 in Online Appendix A.

3 Numerical Results

In this section we assess the value of integrating the portfolio planning and plant assignment decisions by means of a numerical study. We use value iteration to find the optimal policy for both the

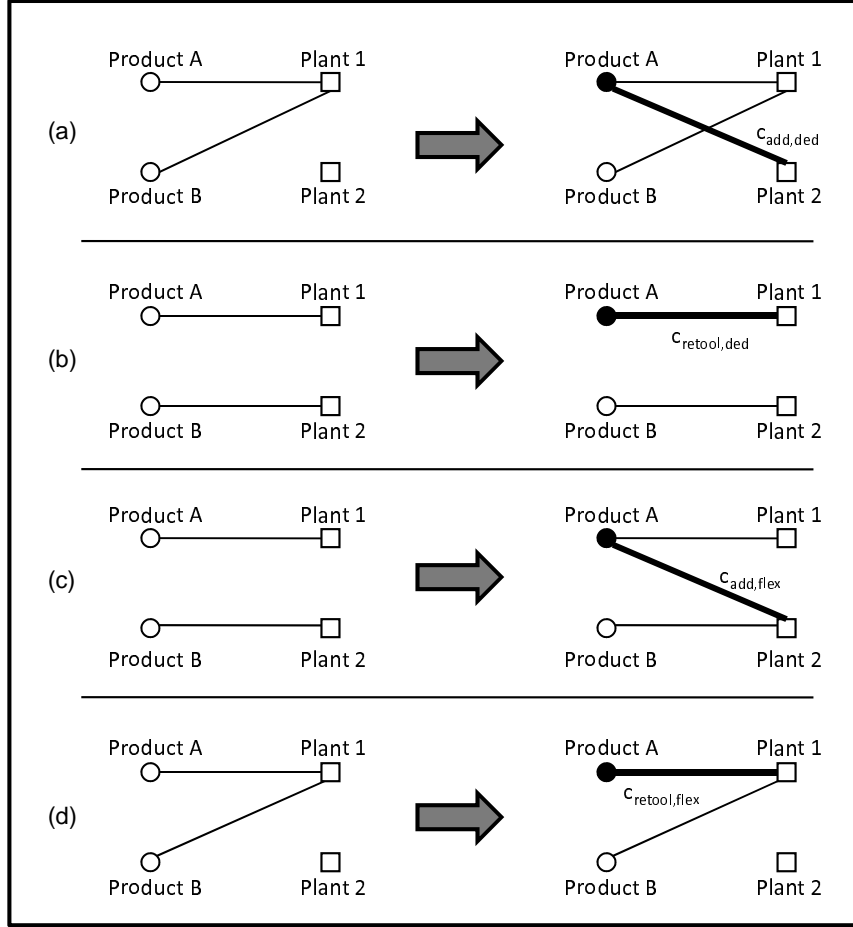


Figure 1: Tooling Costs (Refreshed product represented by filled circle)

integrated and decoupled decision models, as stated in Section 2, and compute the percentage difference in long-run average profit between the two models.

First we define four types of tooling costs, as shown in Figure 1, which represent costs to tool and retool both dedicated (single-product) and flexible (multi-product) plants. We denote the cost of tooling a dedicated plant that is not currently open by $c_{add,ded}$ (Figure 1 (a)), and the cost of retooling a plant currently producing a single vehicle by $c_{retool,ded}$ (Figure 1 (b)). We denote the cost of adding a model to a plant that is currently dedicated to another model by $c_{add,flex}$ (Figure 1 (c)), and the cost of retooling a plant that is already flexible and capable of producing two models by $c_{retool,flex}$ (Figure 1 (d)). Since installing tools in a plant for a newly added model is more expensive than updating tooling for a model that is currently produced in the plant, we assume that $c_{retool,flex} < c_{add,flex}$ and $c_{retool,ded} < c_{add,ded}$.

To limit our consideration to realistic cost scenarios, note that the cost of retooling consists of model-level and plant-level costs. Plant-level costs (e.g., power supply, ventilation, material handling) serve all models that are produced in the plant. Model-level costs are specific to the models produced (e.g., tooling for pressing sheet metal). When a model is retooled, we update both the model and plant, and hence incur both costs. When a plant produces multiple models, the plant-level costs are apportioned to the various models. Therefore, the total cost of retooling ALL models in a plant producing multiple models will be more expensive than retooling one model in a dedicated plant (because of the additional model-specific cost). This implies that $2c_{retool,flex} > c_{retool,ded}$. However, retooling ONE model in a multiple-model plant is less expensive than retooling one model in a single-model dedicated plant (because we are only paying part of the plant-level costs). This means that $c_{retool,flex} < c_{retool,ded}$.

Of course, the cost of opening an empty plant to produce one model incurs both model-level and plant-level costs; while installing tools in a plant that currently produces only one model so that it can also produce a second model incurs primarily model-level costs (because plant-level tooling is already there and can be shared). This implies that $c_{add,flex} < c_{add,ded}$.

Finally, to reduce the number of parameters we need to vary, and to simplify the analysis, we assume that $c_{add,ded}/c_{retool,ded} = c_{add,flex}/c_{retool,flex}$, and therefore we have $c_{add,ded}/c_{add,flex} = c_{retool,ded}/c_{retool,flex}$. Our solution methodology does not require these conditions, however. We let $c_{total} = c_{add,ded} + c_{add,flex} + c_{retool,ded} + c_{flex,ded}$.

We can now state the following proposition:

Proposition 2. *If a 2-product 2-plant system is symmetric, i.e., $m_A = m_B \triangleq m$, $c_{OT,1} = c_{OT,2} \triangleq c_{OT}$, and $K_{reg,1} = K_{reg,2} \triangleq K_{reg}$, then for a given demand value function $f(\cdot)$ the system can be fully described by the following five parameters, in the sense that the optimal policy for integrated and decoupled decisions, as well as the percentage difference in long-run average profit between the integrated and decoupled decision models, is solely determined by these five parameters:*

1. *Dedicated-to-flexible cost ratio* = $c_{add,ded}/c_{add,flex} = c_{retool,ded}/c_{retool,flex}$;
2. *Tool-to-retool cost ratio* = $c_{add,ded}/c_{retool,ded} = c_{add,flex}/c_{retool,flex}$;

3. *Margin-normalized overtime cost* $= c_{OT}/m$;

4. *Utilization* $= f(M)/K_{reg}$;

5. *Tooling-cost-to-revenue factor* $= c_{total}/(m * K_{reg})$.

Proposition 2 allows us to normalize all the system parameters, which generalizes our analysis beyond the specific cases considered in our numerical study.

To keep computation reasonable, we set the number of demand levels to $M = 5$. We use demand value function $f(d) = 0.2d$ to reflect a 20% annual decrease in demand. We use $p = 0.9$ for the binomial distribution of first-year demand, so that after a model refresh, demand rises to the maximum level with about 0.65 probability and reaches 80% of the maximum with 0.95 probability.

We vary the five parameters to cover the region of practical interest as follows:

1. Dedicated-to-flexible cost ratio $1 \leq c_{add,ded}/c_{add,flex} = c_{retool,ded}/c_{retool,flex} \leq 2$;

2. Tool-to-retool cost ratio $1 \leq c_{add,ded}/c_{retool,ded} = c_{add,flex}/c_{retool,flex} \leq 2$;

3. Margin-normalized overtime cost $0.05 \leq c_{OT}/m \leq 0.5$;

4. Utilization $0.5 \leq f(M)/K_{reg} \leq 5$;

5. Tooling-cost-to-revenue factor $2 \leq c_{total}/(m * K_{reg}) \leq 20$.

By examining the results of these numerical cases, we made the following observations:

OBSERVATION 1. Integrated decision making typically outperforms decoupled decision making, although there are certain parameter sets under which the decoupled decision achieves the same performance as the integrated decision (this occurs when the estimates of tooling costs and net revenue happen to be correct). Decoupled decision making cannot outperform integrated decision making, so the smallest percentage difference is zero.

OBSERVATION 2. The parameters “tool-to-retool cost ratio” and “margin-normalized overtime cost” do not have a significant effect on relative performance of the two systems. Therefore, from here on we fix the tool-to-retool cost ratio at 1.8 and the margin-normalized overtime cost at 0.2 to simplify analysis.

The intuition behind observation 2 is that, while “tool-to-retool cost ratio” and “margin-normalized overtime cost” do affect the expected profit, their impact does not depend on whether or not the portfolio planning and plant assignment are integrated. In the former case, this is because “tool-to-retool cost ratio” affects the optimal pace of refreshment similarly in both the integrated and decoupled settings. In the latter case, the “margin-normalized overtime cost” affects the cost of production, and hence the optimal loading of plants, which affects decision making in the integrated and decoupled cases similarly.

We now examine the effect of utilization and dedicated-to-flexible cost ratio on the value of integrating decisions. To do this, we introduce a “statistical” approach for analyzing data. The reason this is needed is that, as we noted earlier, the decoupled decision can often result in the same actions, and hence the same profit, as the integrated decision. So looking at the average profit across a range of cases does not fully characterize the difference between the two approaches. To better reveal this difference, we report multiple percentiles, as opposed to only the mean.

To illustrate this, we consider the impact of utilization on expected profit. With the tool-to-retool cost ratio and margin-normalized overtime cost fixed, we vary the dedicated-to-flexible cost ratio and tooling-cost-to-revenue factor over the ranges given above and calculate the percentage difference between the long-run average profits under optimal integrated and optimal decoupled decision making. Over all the cases, those for which which integrated decision making outperforms decoupled decision making by the largest amounts are of particular interest. So, to see these, we sort the percentage differences, and use the y quantile to measure the performance of the top $100(1 - y)$ percent of the cases. If the y quantile for the utilization value u_1 is larger than that for the utilization value u_2 , then we conclude that the performance advantage of integrated decision making over decoupled decision making is larger when utilization is u_1 than when it is u_2 for the top $100(1 - y)$ percent of cases. By looking at multiple percentiles, we can get a sense of how likely and how large the increase in profits will be by using integrated decision making in place of decoupled decision making.

With this definition of a performance improvement in mind, we can now state our observations about the impact of utilization and dedicated-to-flexible cost ratio on the value of integration.

OBSERVATION 3. The value of integrating portfolio and plant assignment decisions is largest when the system is not over-capacitated or under-capacitated, as shown in Figure 2, where we vary utilization from 0.5 to 5 with step size 0.1. For each utilization value, we also vary the dedicated-to-flexible cost ratio from 1.1 to 1.7 with step size 0.1, and tooling-cost-to-revenue factor from 2 to 20 with step size 0.2. For each utilization this gives about 650 cases whose percentage differences are sorted. We depict the y quantile of percentage difference with y values 0.5, 0.7, 0.9 and 1.

We can see from Figure 2 that the value of integrating portfolio and plant assignment decisions is maximized when utilization is around 1.5. Since utilization is defined as maximum demand over regular capacity, a utilization of 1.5 means total maximum demand equals total capacity in the system. When the system is over-capacitated, say with utilization smaller than 0.8, in 90% out of 650 cases the integrated decision outperforms the decoupled decision by less than 2%, so the value of integrating portfolio planning and plant allocation is not likely to be large. In contrast, when the system is under-capacitated, say with utilization larger than 4.5, in 90% out of 650 cases the integrated decision outperforms the decoupled decision by less than 3%, and in 70% out of 650 cases the integrated and decoupled decision systems perform about the same, so the value of integrating the two decisions is not large in this case either.

This observation is potentially of great practical significance, since it implies that integrating portfolio planning and plant assignment decisions is most important precisely when capacity is fully utilized. Since firms are likely to make capacity decisions over the long term that push utilization near this point, this suggests that integrating these two decisions could substantially improve profits in practice.

OBSERVATION 4. As the dedicated-to-flexible cost ratio increases, the value of integrating portfolio and plant assignment decisions increases, as shown in Figure 3, where we vary the dedicated-to-flexible cost ratio from 1.1 to 1.7 with step size 0.1. For each ratio, we also vary the utilization from 0.5 to 5 with step size 0.1, and tooling-cost-to-revenue factor from 2 to 20 with step size 0.2. Therefore for each dedicated-to-flexible cost ratio we have about 4,000 cases. We depict the y quantile curve of percentage difference for $y = 0.95, 0.99, 1.0$. Note that only the 1.0 quantile (i.e., the maximum value) is strictly increasing in the dedicated-to-flexible cost ratio. In fact, our data show




Figure 2: Effect of Utilization on the Value of Integrating Decisions

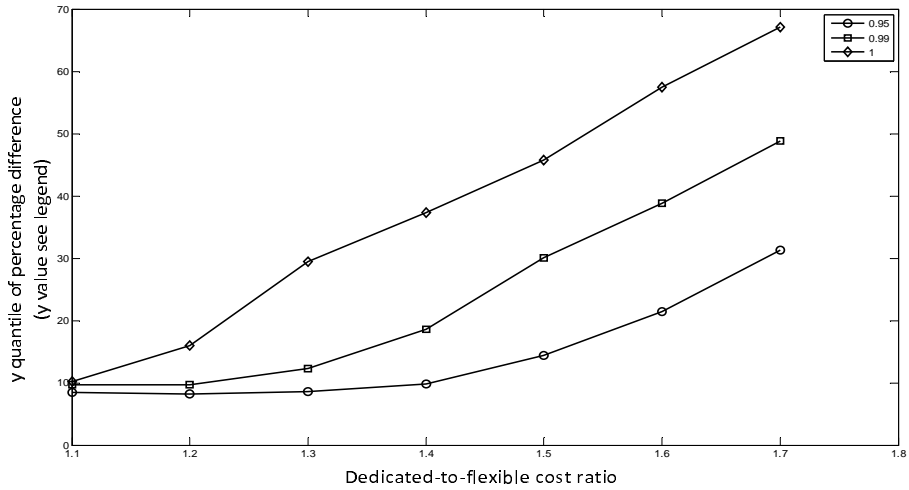


Figure 3: Effect of Dedicated-to-Flexible Cost Ratio

that any quantile above the 0.99 quantile is monotonically increasing. This means that if we focus on any percentage less than 1% of the 4,000 cases with the largest differences between integrated and decoupled decisions, the value of incorporating the impact of current portfolio decisions on later-stage planning decisions increases as the dedicated-to-flexible cost ratio increases. The 0.95 and 0.99 percentiles are not strictly increasing, but still show an overall increasing trend in the dedicated-to-flexible cost ratio.

We can explain the intuition behind this observation by examining the optimal policies more closely. To do this, we define the following two steady-state measures under optimal integrated and decoupled decision policies: (i) average number of plants that are in use, which we call *average*

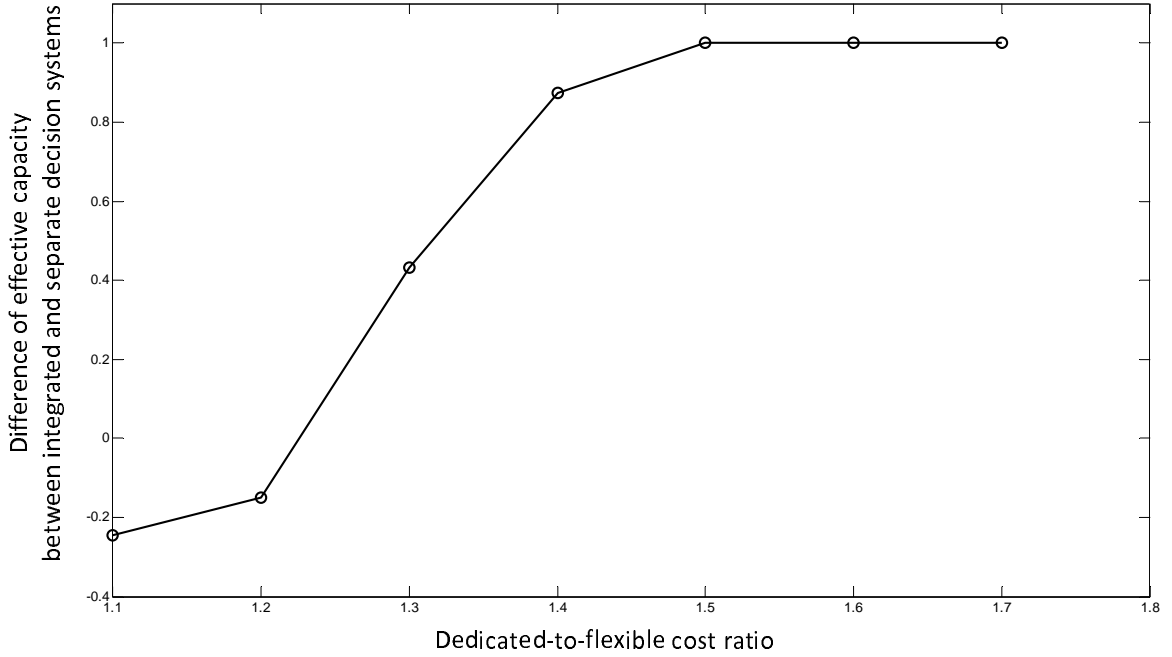


Figure 4: Average Effective Capacity

capacity, since the size of the plants is fixed; and (ii) average number of plants that produce more than one product, which we call *average flexibility*. The sum of (i) and (ii) can be interpreted as the average *effective capacity* in the system, since a larger number indicates that the sum of the maximum amounts of each product that can be produced is greater. Figure 4 depicts the difference in the average effective capacity under the integrated and decoupled decision policies for the 5% of cases with the largest difference. Notice that this difference increases as dedicated-to-flexible cost ratio increases. That is, compared with decoupled decisions, integrated decisions tend to use either more capacity or more flexibility when the dedicated-to-flexible cost ratio is larger. The intuition behind this result can be understood as follows. A larger dedicated-to-flexible cost ratio implies that flexibility (i.e., multi-model production) is more attractive. The optimal integrated policy is more likely to make use of flexibility when it is attractive to do so, which gives it an economic advantage, as depicted in Figure 4. This is consistent with intuition since, compared with the decoupled policy, the integrated policy has better knowledge about downstream stage decisions and uses a more realistic assessment of future tooling costs, and hence responds to change of system parameters better.

To gain insight into the influence of system size, we conducted a similar analysis to that described

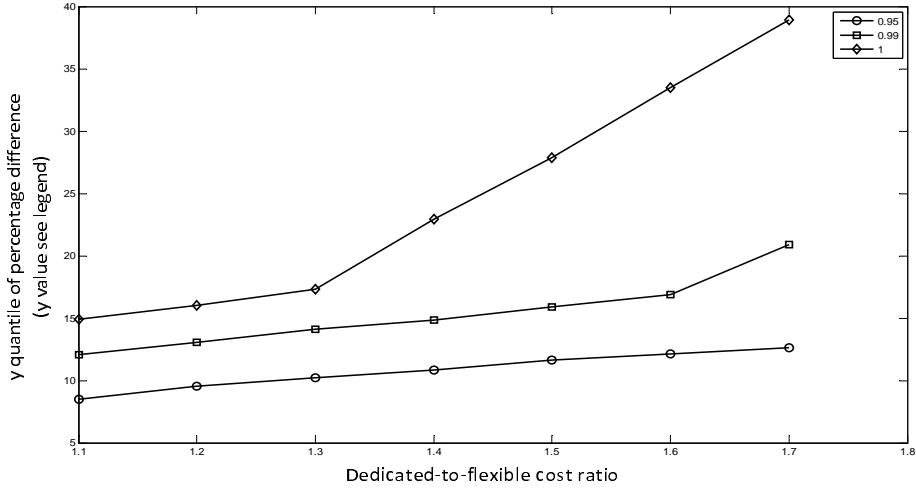


Figure 5: Effect of Dedicated-to-Flexible Cost Ratio for a 3-Product 2-Plant System

above on a 3-product, 2-plant system. Formulas in Section 2 can be extended easily to this case and the analytical results still hold for this larger system. These additional numerical studies yielded the following observation:

OBSERVATION 5. As the number of products gets larger, the impact of the dedicated-to-flexible cost ratio on the value of integrating portfolio and plant assignment decisions becomes stronger. Figure 5 depicts the 0.95, 0.99 and 1.0 quantile curves against the dedicated-to-flexible cost ratio. Note that all of these are monotonically increasing in the dedicated-to-flexible cost ratio, as compared to only the 1.0 quantile in Figure 4 for the 2-product system. Indeed, looking at the data from the numerical studies shows that all percentiles above the 0.85 quantile increase monotonically in the dedicated-to-flexible cost ratio. This suggests that the monotonicity trend is much stronger in a statistical sense. However, the percentage difference values shown in Figure 5 cannot be directly compared to those of Figure 4 as there is no easy way to calibrate the parameters so that the 3-product and 2-product systems are comparable on a case by case basis.

4 Conclusions

From this analysis we conclude that decoupling plant assignment and product allocation decisions from portfolio planning may lead to a significant reduction in expected profit (e.g., more than 50%). However, there are always cases when the decoupled decision happens to guess the right

downstream cost and match the integrated solution. But because we cannot predict when this will occur, our study focuses on the conditions under which the risk of a bad decision from using a decoupled model is large. Our numerical results show that this risk of a large profit penalty from using a decoupled model to address the portfolio planning problem is greatest when plant capacity utilization is moderate and flexible tooling makes it significantly cheaper to introduce new models into plants.

Using an exact optimization model to integrate portfolio planning and plant assignment decisions, as we have done in this paper, is only feasible for very small systems. For realistically sized problems it may be necessary to use a simulation-based optimization model to efficiently search over all possible decisions. However, it may be effective to decouple the problem by segments and use an analytic model like that developed here to optimize each segment for a given resource allocation (e.g., of engineering resources, marketing resources, or vehicle launches). Developing a modeling platform with which to search for the optimal resource allocation is the logical next step toward a practical tool for integrated product portfolio and plant allocation planning.

References

- Alden J. M., T. Costy and R. R. Inman (2002). Product-to-plant allocation assessment in the automotive industry. *Journal of Manufacturing Systems* 21, 1, 1-13.
- Barlow, R. E., F. Proschan and L. C. Hunter (1965). *Mathematical Theory of Reliability* John Wiley & Sons, New York.
- Cooper, R. G., S. J. Edgett and E. J. Kleinschmidt (1999). New product portfolio management: practice and performance. *J. Prod. Innov. Manag.* 16, 333-351.
- Dickinson, M. W., A. C. Thornton and S. Graves (2001). Technology Portfolio Management: Optimizing interdependent projects over multiple time periods. *IEEE Transaction on Engineering Management* Vol. 48, No. 4, 518-527.
- Heidenberger, K. and C. Stummer (1999). Research and development project selection and resource allocation: A review of quantitative modeling approaches. *Intl. J. Manag. Rev.* Vol. 1, 197-224.
- Inman, R. R. and D. J. A. Gonsalvez (2001). A mass production product-to-plant allocation problem. *Computers and Industrial Engineering* 39, 255-271.
- Loch, C. H. and S. Kavadias (2002). Dynamic portfolio selection of NPD programs using marginal returns. *Management Science* Vol. 48, No. 10, 1227-1241.
- McCall, J. J. (1965). Maintenance policies for stochastically failing equipment: a survey. *Management Science* 11, 5, 493-524.
- Pierskalla, W. P. and J. A. Voelker (1976). A survey of maintenance models: the control and surveillance of deteriorating systems. *Naval Research Logistics Quarterly* 23, 3, 353-388.

- Puterman, M. L. (1994). *Markov Decision Process*. Wiley, New York.
- Sherif, Y. S. and M. L. Smith (1981). Optimal maintenance models for systems subject to failure - a review. *Naval Research Logistics Quarterly* 28, 1, 47-74.
- Sloan, T. W. and J. G. Shanthikumar (2000). Combined production and maintenance scheduling for a multi-product single-machine production system. *Production and Operations Management* 9, 4, 379-399.
- Stummer, C. and K. Heidensberger (2003). Interactive R& D portfolio analysis with project interdependencies and time profiles of multiple objectives. *IEEE Transaction on Engineering Management* Vol. 50, No. 2, 175-183.
- Valdez-Flores, C. and R. M. Feldman (1989). A survey of preventive maintenance models for stochastically deteriorating single-unit systems. *Naval Research Logistics* 36, 419-446.
- Wang, H. (2002). A survey of maintenance policies of deteriorating systems. *European Journal of Operational Research* 139, 469-489.

ONLINE APPENDIX A

MDP Formulation for a 2-Product 2-Plant System: Decoupled Decisions

Estimating Tooling Costs:

In the decoupled case, we do not find the optimal plant assignment for every candidate portfolio plan in order to compute plant level costs. Instead, we estimate these costs by averaging the cost of all possible plant assignments. We do this by solving an MDP model that assumes a fixed T -year ($T = 5$) refresh cycle, where we define:

- *State space* \mathcal{S} with states $(t_A, t_B, d_A, d_B, y_A, y_B)$, where $t_A, t_B \in \{1, 2, \dots, T-1\}$ are age of product A and B, and d_A, d_B, y_A , and y_B are defined as in Section 2.1.
- *Decision epochs* are the beginning of every year $t = 1, 2, \dots$
- *Action space* is defined so that a model is refreshed when it reaches age T , which leads to a T -year refresh cycle:

$$\mathcal{A}(t_A, t_B, d_A, d_B, y_A, y_B) = \begin{cases} \{0\} \times \{0\} & \text{if } t_A \neq T \text{ and } t_B \neq T, \\ \{0\} \times \{1, 2, 1\&2\} & \text{if } t_A \neq T \text{ and } t_B = T, \\ \{1, 2, 1\&2\} \times \{0\} & \text{if } t_A = T \text{ and } t_B \neq T, \\ \{1, 2, 1\&2\} \times \{1, 2, 1\&2\} & \text{if } t_A = T \text{ and } t_B = T. \end{cases}$$

- *Reward* equals net revenue minus tooling costs, as defined in Section 2.1.

We make the same assumptions about demand change as in Section 2.1. We define $V(t_A, t_B, d_A, d_B, y_A, y_B)$ as the value function in state $(t_A, t_B, d_A, d_B, y_A, y_B)$, and g as the optimal average profit per year, and have the following optimality equation:

$$\begin{aligned} & V(t_A, t_B, d_A, d_B, y_A, y_B) + g & (6) \\ = & \max_{\substack{(a_A, a_B) \in \\ \mathcal{A}(t_A, t_B, d_A, d_B, y_A, y_B)}} \left\{ \pi(d_A, d_B, y_A, y_B) - c_{tool}(y_A, y_B, a_A, a_B) + \right. \\ & \left. \sum_{j_A} \sum_{j_B} \tilde{P}(j_A, a_A) \tilde{P}(j_B, a_B) \times \right. \\ & \left. V(T(t_A, a_A), T(t_B, a_B), D(d_A, j_A, a_A), D(d_B, j_B, a_B), Y(y_A, a_A), Y(y_B, a_B)) \right\} \end{aligned}$$

where

$$\begin{aligned} \tilde{P}(j, a) &= \begin{cases} P(j) & a = 1, 2, 1\&2, \\ 1 & a = 0, \end{cases} \\ D(d, j, a) &= \begin{cases} j & a = 1, 2, 1\&2, \\ \max(d-1, 1) & a = 0, \end{cases} \\ Y(y, a) &= \begin{cases} a & a = 1, 2, 1\&2, \\ y & a = 0, \end{cases} \\ T(t, a) &= \begin{cases} 1 & a = 1, 2, 1\&2, \\ t+1 & a = 0. \end{cases} \end{aligned}$$

The average tooling cost \bar{c}_{tool} is computed under the optimal policy as the long-run average tooling costs per refresh.

For the net revenue, we compute the average of all possible plan assignments for each demand level:

$$\bar{\pi}(d_A, d_B) = \frac{1}{9} \sum_{y_A} \sum_{y_B} \pi(d_A, d_B, y_A, y_B).$$

Portfolio Planning:

We can now formulate an MDP to find the best decoupled solution to the portfolio planning problem by defining:

- *State space* \mathcal{S} with states (d_A, d_B) , which is defined in Section 2.1.
- *Decision epochs* are the beginning of every year $t = 1, 2, \dots$
- *Action space* \mathcal{A} includes action (a_A, a_B) . $a_A, a_B \in \{0, r\}$ are actions for product A and B, respectively, where 0 represents no refresh, r refresh.
- *Reward* equals net revenue minus tooling costs. We assume assigning a newly refreshed model to a plant incurs a tooling cost \bar{c}_{tool} , which is calculated under the optimal policy of the MDP model (6), and net revenue per year is $\bar{\pi}(d_A, d_B)$ for demand level (d_A, d_B) .

We make the same assumptions about demand as in the integrated decision model. We define $V(d_A, d_B)$ as the value function in state (d_A, d_B) , and g as the optimal average profit per year, and write the optimality equation of the MDP model as

$$\begin{aligned} & V(d_A, d_B) + g & (7) \\ = \max & \begin{cases} \bar{\pi}(d_A, d_B) + V(\max(d_A - 1, 1), \max(d_B - 1, 1)) & (a_A, a_B) = (0, 0) \\ \bar{\pi}(d_A, d_B) - \bar{c}_{tool} + \sum_{j_B} P(j_B) V(\max(d_A - 1, 1), j_B) & (a_A, a_B) = (0, r) \\ \bar{\pi}(d_A, d_B) - \bar{c}_{tool} + \sum_{j_A} P(j_A) V(j_A, \max(d_B - 1, 1)) & (a_A, a_B) = (r, 0) \\ \bar{\pi}(d_A, d_B) - 2\bar{c}_{tool} + \sum_{j_A} \sum_{j_B} P(j_A) P(j_B) V(j_A, j_B) & (a_A, a_B) = (r, r) \end{cases} \end{aligned}$$

We call the optimal portfolio planning policy $pol(d_A, d_B)$, which decides whether or not to refresh model(s) for the given demand level.

Plant Assignment:

Finally, once a portfolio planning policy has been selected, we find the best plant assignment for it by solving an MDP with:

- *State space* \mathcal{S} includes state (d_A, d_B, y_A, y_B) , which is defined as in Section 2.1.
- *Decision epochs* are the beginning of every year $t = 1, 2, \dots$
- *Action space* $\mathcal{A}(d_A, d_B, y_A, y_B)$ is defined as follows

$$\mathcal{A}(d_A, d_B, y_A, y_B) = \begin{cases} \{0\} \times \{0\} & \text{if } pol(d_A, d_B) = (0, 0), \\ \{0\} \times \{1, 2, 1\&2\} & \text{if } pol(d_A, d_B) = (0, r), \\ \{1, 2, 1\&2\} \times \{0\} & \text{if } pol(d_A, d_B) = (r, 0), \\ \{1, 2, 1\&2\} \times \{1, 2, 1\&2\} & \text{if } pol(d_A, d_B) = (r, r), \end{cases}$$

which means that we only assign a product when it is refreshed according to the decisions at the previous (portfolio planning) stage.

- *Reward* equals net revenue minus tooling costs, which are defined similarly as in Section 2.1.

We make the same assumptions about demand as in the integrated decision model. We define $V(d_A, d_B, y_A, y_B)$ as the value function in state (d_A, d_B, y_A, y_B) , and g as the optimal average profit per year, and write the optimality equation of the MDP model as

$$\begin{aligned}
& V(d_A, d_B, y_A, y_B) + g \\
= & \max_{\substack{(a_A, a_B) \in \\ \mathcal{A}(t_A, t_B, d_A, d_B, y_A, y_B)}} \left\{ \pi(d_A, d_B, y_A, y_B) - c_{tool}(y_A, y_B, a_A, a_B) + \right. \\
& \left. \sum_{j_A} \sum_{j_B} \tilde{P}(j_A, a_A) \tilde{P}(j_B, a_B) V(D(d_A, j_A, a_A), D(d_B, j_B, a_B), Y(y_A, a_A), Y(y_B, a_B)) \right\}
\end{aligned} \tag{8}$$

where

$$\begin{aligned}
\tilde{P}(j, a) &= \begin{cases} P(j) & a = 1, 2, 1\&2, \\ 1 & a = 0, \end{cases} \\
D(d, j, a) &= \begin{cases} j & a = 1, 2, 1\&2, \\ \max(d-1, 1) & a = 0, \end{cases} \\
Y(y, a) &= \begin{cases} a & a = 1, 2, 1\&2, \\ y & a = 0. \end{cases}
\end{aligned}$$

The optimal policy gives the plant assignment for the given demand level and current assignment. We have the following theorem to justify the MDP solution.

Theorem 3. *The following hold:*

1. *For the cost estimate stage, the state space can be divided into T subsets, and for each subset, there exists a stationary optimal policy for the MDP defined in (6) that has a constant average gain, and the corresponding value iteration algorithm converges;*
2. *For the portfolio planning stage, there exists a stationary optimal policy for the MDP defined by (7) that has a constant average gain. Moreover, the corresponding value iteration algorithm converges;*
3. *For the plant assignment stage, if the portfolio planning decides that both products are refreshed when demand drops to the lowest level, i.e., $pol(1, 1) = (r, r)$, then there exists a stationary optimal policy for the MDP defined in (8) that has a constant average gain. Moreover, with an aperiodicity transformation, the corresponding value iteration algorithm converges.*

ONLINE APPENDIX B

Proofs of Analytical Results

PROOF OF THEOREM 1:

Since state space \mathcal{S} and action set \mathcal{A} are finite, by Theorem 9.1.8 of Puterman (1994), there exists a deterministic stationary optimal policy.

We now prove that the model is communicating. For an arbitrary pair of states $s = (d_A, d_B, y_A, y_B)$ and $s' = (d'_A, d'_B, y'_A, y'_B)$, under action $a = (y'_A, y'_B)$, state s reaches state s' with probability $P(d'_A)P(d'_B) > 0$ (from the property of binomial distribution probability function). Therefore, s' is accessible from s under a stationary deterministic policy d^∞ with $d(s) = (y'_A, y'_B)$. Since s and s' are arbitrarily chosen, this completes the proof. Hence, by Theorem 8.3.2 of Puterman (1994), there exists a stationary optimal policy with constant gain.

From Section 8.5.4 of Puterman (1994), we know that through a simple transformation, all policies can be made aperiodic. Then by Corollary 9.4.6 of Puterman (1994), the value iteration algorithm converges and the span can be used as a stopping criterion.

PROOF OF PROPOSITION 1:

We use induction and value iteration to prove the proposition.

It is clear that at the first iteration of the value iteration algorithm, $V_0(d_A, d_B, y_A, y_B) = \pi(d_A, d_B, y_A, y_B)$ is non-decreasing in d_A . We assume that value function is non-decreasing in d_A at iteration k , we will prove that it is non-decreasing in d_A at iteration $k + 1$.

$$\begin{aligned}
 & V_{k+1}(d_A, d_B, y_A, y_B) \\
 &= \max_{(a_A, a_B) \in \mathcal{A}} \left\{ \pi(d_A, d_B, y_A, y_B) - c_{tool}(y_A, y_B, a_A, a_B) + \right. \\
 & \quad \left. \sum_{j_A} \sum_{j_B} \tilde{P}(j_A, a_A) \tilde{P}(j_B, a_B) V_k(D(d_A, j_A, a_A), D(d_B, j_B, a_B), Y(y_A, a_A), Y(y_B, a_B)) \right\}
 \end{aligned} \tag{9}$$

where

$$\begin{aligned}
 \tilde{P}(j, a) &= \begin{cases} P(j) & a = 1, 2, 1\&2, \\ 1 & a = 0, \end{cases} \\
 D(d, j, a) &= \begin{cases} j & a = 1, 2, 1\&2, \\ \max(d - 1, 1) & a = 0, \end{cases} \\
 Y(y, a) &= \begin{cases} a & a = 1, 2, 1\&2, \\ y & a = 0. \end{cases}
 \end{aligned}$$

Since $\pi(d_A, d_B, y_A, y_B)$ is non-decreasing in d_A , $c_{tool}(y_A, y_B, a_A, a_B)$ does not depend on d_A , $D(d, j, a)$ is non-decreasing in d and $V_k(d_A, d_B, y_A, y_B)$ is non-decreasing in d_A , the expression within braces in (9) is non-decreasing in d_A . Therefore $V_{k+1}(d_A, d_B, y_A, y_B)$ is non-decreasing in d_A .

PROOF OF THEOREM 2:

Define

$$\begin{aligned}
 & R(d_A, d_B, y_A, y_B, a_A, a_B) \\
 &= \pi(d_A, d_B, y_A, y_B) - c_{tool}(y_A, y_B, a_A, a_B) + \\
 & \quad \sum_{j_A} \sum_{j_B} \tilde{P}(j_A, a_A) \tilde{P}(j_B, a_B) V_k(D(d_A, j_A, a_A), D(d_B, j_B, a_B), Y(y_A, a_A), Y(y_B, a_B)),
 \end{aligned}$$

where

$$\begin{aligned}\tilde{P}(j, a) &= \begin{cases} P(j) & a = 1, 2, 1\&2, \\ 1 & a = 0, \end{cases} \\ D(d, j, a) &= \begin{cases} j & a = 1, 2, 1\&2, \\ \max(d-1, 1) & a = 0, \end{cases} \\ Y(y, a) &= \begin{cases} a & a = 1, 2, 1\&2, \\ y & a = 0. \end{cases}\end{aligned}$$

Suppose action $(0, \beta)$ is optimal at state (d_A, d_B, y_A, y_B) for some $\beta \in \{0, 1, 2, 1\&2\}$, we have

$$R(d_A, d_B, y_A, y_B, 0, \beta) \geq R(d_A, d_B, y_A, y_B, a_A, a_B) \quad \forall a_A \in \{1, 2, 1\&2\}, \quad a_B \in \{0, 1, 2, 1\&2\}. \quad (10)$$

We will prove $(0, \beta)$ dominates all the “refresh for model A” actions at state $(d_A + 1, d_B, y_A, y_B)$, i.e.,

$$R(d_A + 1, d_B, y_A, y_B, 0, \beta) \geq R(d_A + 1, d_B, y_A, y_B, a_A, a_B) \quad \forall a_A \in \{1, 2, 1\&2\}, \quad a_B \in \{0, 1, 2, 1\&2\}. \quad (11)$$

From the optimality Equation (5) we have

$$\begin{aligned}R(d_A + 1, d_B, y_A, y_B, 0, 0) - R(d_A, d_B, y_A, y_B, 0, 0) \\ = \pi(d_A + 1, d_B, y_A, y_B) - \pi(d_A, d_B, y_A, y_B) \\ + V(\max(d_A, 1), \max(d_B - 1, 1), y_A, y_B) - V(\max(d_A - 1, 1), \max(d_B - 1, 1), y_A, y_B),\end{aligned}$$

and for all $a_B \in \{1, 2, 1\&2\}$,

$$\begin{aligned}R(d_A + 1, d_B, y_A, y_B, 0, a_B) - R(d_A, d_B, y_A, y_B, 0, a_B) \\ = \pi(d_A + 1, d_B, y_A, y_B) - \pi(d_A, d_B, y_A, y_B) \\ + \sum_{j_B} P(j_B) V(\max(d_A, 1), j_B, y_A, a_B) - \sum_{j_B} P(j_B) V(\max(d_A - 1, 1), j_B, y_A, a_B).\end{aligned}$$

From Proposition 1, $V(d_A, d_B, y_A, y_B)$ is non-decreasing in d_A . Also, $\max(d_A - 1, 1)$ is non-decreasing in d_A . Therefore for $a_B \in \{0, 1, 2, 1\&2\}$, we have

$$R(d_A + 1, d_B, y_A, y_B, 0, a_B) - R(d_A, d_B, y_A, y_B, 0, a_B) \geq \pi(d_A + 1, d_B, y_A, y_B) - \pi(d_A, d_B, y_A, y_B). \quad (12)$$

On the other hand, for $a_A \in \{1, 2, 1\&2\}$, since d_A appears in $R(d_A, d_B, y_A, y_B, a_A, a_B)$ only through $\pi(\cdot)$, we have

$$\begin{aligned}R(d_A + 1, d_B, y_A, y_B, a_A, a_B) - R(d_A, d_B, y_A, y_B, a_A, a_B) = \pi(d_A + 1, d_B, y_A, y_B) - \pi(d_A, d_B, y_A, y_B) \\ \forall a_A \in \{1, 2, 1\&2\}, \quad a_B \in \{0, 1, 2, 1\&2\}. \quad (13)\end{aligned}$$

Therefore,

$$\begin{aligned}R(d_A + 1, d_B, y_A, y_B, 0, \beta) - R(d_A, d_B, y_A, y_B, 0, \beta) \\ \geq \pi(d_A + 1, d_B, y_A, y_B) - \pi(d_A, d_B, y_A, y_B) \quad \text{from (12)} \\ = R(d_A + 1, d_B, y_A, y_B, a_A, a_B) - R(d_A, d_B, y_A, y_B, a_A, a_B) \quad \text{from (13)} \\ \forall a_A \in \{1, 2, 1\&2\}, \quad a_B \in \{0, 1, 2, 1\&2\},\end{aligned}$$

or equivalently,

$$\begin{aligned}R(d_A + 1, d_B, y_A, y_B, 0, \beta) - R(d_A + 1, d_B, y_A, y_B, a_A, a_B) \\ \geq R(d_A, d_B, y_A, y_B, 0, \beta) - R(d_A, d_B, y_A, y_B, a_A, a_B) \\ \geq 0 \quad \text{from (10)} \\ \forall a_A \in \{1, 2, 1\&2\}, \quad a_B \in \{0, 1, 2, 1\&2\},\end{aligned}$$

and (11) follows.

Note that the proof of Theorem 2 does not depend on the size of the system. It does require that $P(M) \neq 0$ so the MDP is weakly communicating. Theorem 2 holds for an N -product K -plant system and general probability function, as long as $P(M) \neq 0$.

PROOF OF THEOREM 3:

PART 1

First we divide the state space \mathcal{S} into T subsets $\mathcal{S} = \mathcal{S}_0 \cup \mathcal{S}_1 \cup \dots \cup \mathcal{S}_{T-1}$, where

$$(t_A, t_B, d_A, d_B, y_A, y_B) \in \begin{cases} \mathcal{S}_0 & \text{if } (t_A, t_B) \in \{(1, 1), (2, 2), \dots, (T, T)\} \\ \mathcal{S}_1 & \text{if } (t_A, t_B) \in \{(1, 2), (2, 3), \dots, (T, 1)\} \\ \vdots & \\ \mathcal{S}_{T-1} & \text{if } (t_A, t_B) \in \{(1, T), (2, 1), \dots, (T, T-1)\}. \end{cases} \quad (14)$$

We focus on one subset \mathcal{S}_k ($k \in \{0, 1, \dots, T-1\}$). Since state space \mathcal{S}_k and action set \mathcal{A} are finite, by Theorem 9.1.8 of Puterman (1994), there exists a deterministic stationary optimal policy.

Now we prove that the MDP defined by (6) is weakly communicating. We know that all the states with $t_A + d_A - 1 > T$ or $t_B + d_B - 1 > T$ are transient since they are not reachable from any other states under all policies. For an arbitrary pair of states $s = (t_A, t_B, d_A, d_B, y_A, y_B)$ and $s' = (t'_A, t'_B, d'_A, d'_B, y'_A, y'_B)$ that are not transient, define $\Delta t_A = t'_A - t_A$ and $\Delta t_B = t'_B - t_B$. From the way \mathcal{S}_k is defined in (14), we have either $\Delta t_A = \Delta t_B$ or $|\Delta t_A - \Delta t_B| = T$.

For the $\Delta t_A = \Delta t_B$ case, define $\Delta t \triangleq \Delta t_A = \Delta t_B$. Consider a Δt -step transition from state s , which updates the age of the two products from (t_A, t_B) to (t'_A, t'_B) . Within these steps both product A and B are refreshed exactly once. Choose action y'_A when product A is refreshed, and choose action y'_B when product B is refreshed. Then with probability $P(d'_A + t'_A - 1)P(d'_B + t'_B - 1) > 0$, state s' is reached.

For the $|\Delta t_A - \Delta t_B| = T$ case, without loss of generality, assume $t_A > 0$. Consider a $(T + \Delta t_A)$ -step transition from state s , which updates the age of the two products from (t_A, t_B) to (t'_A, t'_B) . Within these steps product A is refreshed twice and product B is refreshed once. Choose action y'_A when product A is refreshed the second time, and choose action y'_B when product B is refreshed. Then with probability $P(d'_A + t'_A - 1)P(d'_B + t'_B - 1) > 0$, state s' is reached.

Therefore s' is accessible from s under a stationary deterministic policy d^∞ as defined in the above. Since s and s' are arbitrarily chosen, this completes the proof. Then, by Theorem 8.3.2 of Puterman (1994), there exists a stationary optimal policy with constant gain.

From Section 8.5.4 of Puterman (1994), we know that through a simple transformation, all policies can be made aperiodic. Then by Corollary 9.4.6 of Puterman (1994), the value iteration algorithm converges and the span can be used as a stopping criterion.

Since the arguments above are for arbitrary k , we have proved that for each of the T subsets, there exists a stationary optimal policy for the MDP defined in (6) that has a constant average gain, and the corresponding value iteration algorithm converges.

PART 2

Since state space \mathcal{S} and action set \mathcal{A} are finite, by Theorem 9.1.8 of Puterman (1994), there exists a deterministic stationary optimal policy.

We now prove that the model is communicating. For an arbitrary pair of states $s = (d_A, d_B)$ and $s' = (d'_A, d'_B)$, under action $a = (r, r)$, state s reaches state s' with probability $P(d'_A)P(d'_B) > 0$ (from the property of binomial distribution probability function). Therefore s' is accessible from s under a stationary deterministic policy d^∞ with $d(s) = (r, r)$. Since s and s' are arbitrarily chosen, this completes the proof. Hence, by Theorem 8.3.2 of Puterman (1994), there exists a stationary optimal policy with constant gain.

From Section 8.5.4 of Puterman (1994), we know that through a simple transformation, all policies can be made aperiodic. Then by Corollary 9.4.6 of Puterman (1994), the value iteration algorithm converges and the span can be used as a stopping criterion.

PART 3

Since state space \mathcal{S} and action set \mathcal{A} are finite, by Theorem 9.1.8 of Puterman (1994), there exists a deterministic stationary optimal policy.

We now prove that the model is communicating. For an arbitrary pair of states $s = (d_A, d_B, y_A, y_B)$ and $s' = (d'_A, d'_B, y'_A, y'_B)$, we will show that there exists a deterministic stationary policy under which s' is accessible from s . Starting from state $s = (d_A, d_B, y_A, y_B)$, by a one-step transition, the probability of reaching state $(\max(d_A - 1, 1), \max(d_B - 1, 1), y_A, y_B)$ equals

$$\begin{cases} 1 & \text{under action } (0, 0) \text{ if } \mathcal{A}(s) = \{0\} \times \{0\}, \\ P(\max(d_B - 1, 1)) & \text{under action } (0, y_B) \text{ if } \mathcal{A}(s) = \{0\} \times \{1, 2, 1\&2\}, \\ P(\max(d_A - 1, 1)) & \text{under action } (y_A, 0) \text{ if } \mathcal{A}(s) = \{1, 2, 1\&2\} \times \{0\}, \\ P(\max(d_A - 1, 1))P(\max(d_B - 1, 1)) & \text{under action } (y_A, y_B) \text{ if } \mathcal{A}(s) = \{1, 2, 1\&2\} \times \{1, 2, 1\&2\}, \end{cases}$$

which is positive. Hence, starting from s , as the state space is finite, within finite steps state $s'' = (1, 1, y_A, y_B)$ is reached with positive probability. And since $pol(1, 1) = (r, r)$, $\mathcal{A}(s'') = \{1, 2, 1\&2\} \times \{1, 2, 1\&2\}$. Under action $(y'_A, y'_B) \in \mathcal{A}(s'')$ state $s' = (d'_A, d'_B, y'_A, y'_B)$ is reached from state s'' with positive probability $P(d'_A)P(d'_B)$. Therefore s' is accessible from s under a stationary deterministic policy d^∞ as described above. Since s and s' are arbitrarily chosen, this completes the proof. Hence, by Theorem 8.3.2 of Puterman (1994), there exists a stationary optimal policy with constant gain.

From Section 8.5.4 of Puterman (1994), we know that through a simple transformation, all policies can be made aperiodic. Then by Corollary 9.4.6 of Puterman (1994), the value iteration algorithm converges and the span can be used as a stopping criterion.

PROOF OF PROPOSITION 2:

For a symmetric system, the linear programming model \mathbf{P}_π ((1)-(4)) can be written as

$$\begin{aligned} & \pi(d_A, d_B, y_A, y_B) \\ & = \max_{\mathbf{x}} \left\{ m \left\{ x_{A1} + x_{A2} + x_{B1} + x_{B2} - (c_{OT}/m) \left[(x_{A1} + x_{B1} - K_{reg})^+ + (x_{A2} + x_{B2} - K_{reg})^+ \right] \right\} \right\} \end{aligned}$$

subject to:

$$\begin{aligned} x_{A1} + x_{A2} &\leq K_{reg}(f(d_A)/K_{reg}), & x_{B1} + x_{B2} &\leq K_{reg}(f(d_B)/K_{reg}), \\ x_{A1} + x_{B1} &\leq 1.5K_{reg}, & x_{A2} + x_{B2} &\leq 1.5K_{reg}, \\ x_{A1} &= 0 \quad \text{if } y_A = 2, & x_{A2} &= 0 \quad \text{if } y_A = 1, \\ x_{B1} &= 0 \quad \text{if } y_B = 2, & x_{B2} &= 0 \quad \text{if } y_B = 1. \end{aligned}$$

So for a given demand value function $f(\cdot)$, and fixed ratios $f(M)/K_{reg}$ and c_{OT}/m , using the notation $\pi(m, K_{reg}, d_A, d_B, y_A, y_B)$ to emphasize that the net revenue is a function of parameters m and K_{reg} , we have $\pi(\alpha \cdot m, \beta \cdot K_{reg}, d_A, d_B, y_A, y_B) = \alpha\beta \cdot \pi(m, K_{reg}, d_A, d_B, y_A, y_B)$.

On the other hand, for fixed ratios $c_{add,ded}/c_{add,flex} = c_{retool,ded}/c_{retool,flex}$ and $c_{add,ded}/c_{retool,ded} = c_{add,flex}/c_{retool,flex}$, using the notation $c_{tool}(c_{total}, y_A, y_B, a_A, a_B)$ to emphasize that the tooling cost is a function of the parameter c_{total} , we have $c_{tool}(\alpha \cdot c_{total}, y_A, y_B, a_A, a_B) = \alpha \cdot c_{tool}(c_{total}, y_A, y_B, a_A, a_B)$. Thus if we further fix the ratio $c_{total}/(m * K_{reg})$, in the optimality equation (5), the cost structure for the MDP is fixed, that is, the relative value of different actions does not change. Therefore, the optimal policy is determined by the five parameters. Similar arguments hold for the decoupled decision model.