Interfaces
Vol. 37, No. 6, November–December 2007, pp. 553–565
ISSN 0092-2102 | EISSN 1526-551X | 07 | 3706 | 0553

informs®

DOI 10.1287/inte.1070.0314

Operations Research Helps Reshape Operations Strategy at Standard Register Company

Sanjay L. Ahire
Department of Management Science, Moore School of Business, University of South Carolina, Columbia, South Carolina 29208, ahire@moore.sc.edu

Michael F. Gorman
Department of MIS, Operations Management and Decision Sciences, University of Dayton, Dayton, Ohio 45469, michael.gorman@udayton.edu

David Dwiggins, Oleh Mudry
The Standard Register Company, Dayton, Ohio 45408
{david.dwiggins@standardregister.com, oleh.mudry@standardregister.com}

Standard Register (SR) Company is a 93-year-old firm leading in high-volume print production of forms and print stationery for major US firms. SR is facing the strategic challenge of minimizing the total landed costs to offer competitive pricing in the highly competitive traditional print market segment. We applied a trio of operations research (OR) techniques to help SR optimally allocate the production orders across its production-distribution network for minimizing the total landed cost: (1) regressions to estimate the cost and time efficiency attributes of various printing presses on print jobs of different types; (2) optimization modeling to determine the optimal order-routing strategy; and (3) simulation modeling of the production-distribution network to assess the effectiveness of optimal and heuristic allocation strategies under uncertainty of customer orders and equipment performance. With an estimated potential annual savings of over $10 million across SR's major product segments in the high-volume rotary production business, the study has resulted in a strategic shift in SR's capacity-allocation policies. SR's executive-leadership team has launched system-wide production-distribution improvement initiatives and expedited efforts to build real-time supply chain decision-support capabilities to support this philosophy.

Key words: production-distribution networks: assignment; simulation.
History: This paper was refereed.

Standard Register Company (NYSE: SR) is a leading traditional print and digital information-capture and management firm. It has more than 90 years of experience in its traditional document and label business solutions, and is a leader in innovation in its newer digital products such as electronic forms and document security. SR and the rest of the printed business-documents industry is currently faced with oversupply due to conversion from traditional print documents to new product lines using digital formats. Thus, SR is faced with making challenging operational decisions to maintain competitiveness in its traditional print business. One of SR's major challenges in the traditional print-production business is to reduce total landed costs to offer competitive pricing to customers in a highly competitive market.

The "rotary" product in the print industry generally focuses on high-volume, low-cost print jobs that are run on a continuous-feed roller and are subsequently cut, folded, and trimmed to customer specifications. Examples of such jobs are monthly statements in the financial services, bills in the health-care industries, and customized target-marketing mailers in retail. Approximately 150 presses exist across 13 locations in SR's network of facilities, each with different capabilities and efficiencies. In some cases, a job cannot be performed on a machine because of its requirements (e.g., a four-color job is required and a press has a two-color capability). In other cases, the press might not be well suited for the job, such as when its setup time and cost for configuring the press are high, or because the cost per unit is high on a less efficient press and the job size is large.

SR recognized that deciding which of its thousands of customer orders (jobs) went to which print presses each year was critical to realizing efficiencies in an increasingly competitive market. Although SR has established its order-allocation rules centrally, it frequently makes tactical decisions on the fly based on available local press capacity, geographic proximity, and marketing incentives. SR wanted to reassign customer orders of different sizes and requirements to presses at different locations and of different capabilities to reduce total costs and improve service, or both. We successfully applied a trio of classic operations research (OR) techniques to help SR drastically change its production-orders allocation logic to the contemporary strategy of minimizing total landed cost for the production-distribution network of the firm (SR has coined this approach as *performance-based routing strategy*).

First, we used regression analysis to estimate the cost and time efficiency attributes of various printing presses on print jobs of different types, as well as to evaluate the confidence of optimization results based on these estimates. Second, we employed optimization modeling to determine recommended job-allocation strategies (production, transportation, and materials costs) subject to capacity (hours) constraints of each press in the production network. Finally, we developed a comprehensive simulation model of the production-distribution network and used it to assess the likelihood of success of allocation strategies recommended from optimization as well other heuristics, within the uncertainty of customer orders and equipment performance.

Existing Production-Distribution Network

The focus of our project was SR's traditional high-volume lithographic print business. Its current production network for this part of its business consists of six plants located across the United States and with annual revenues of about $350 million. Each of these plants processes two types of orders. Figure 1 displays

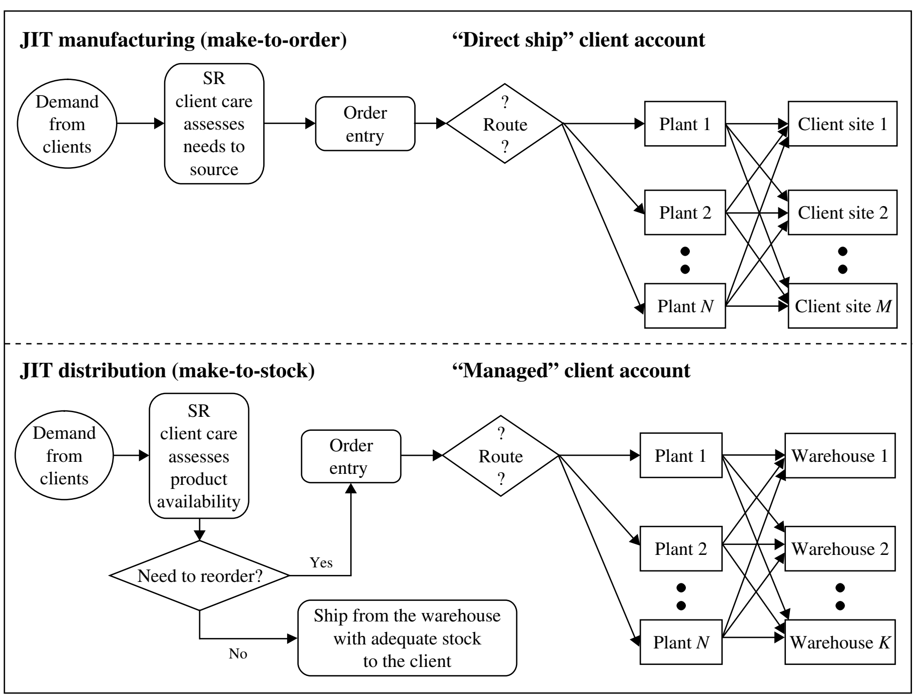

Figure 1: Standard Register (SR) Company's production-distribution network logic depends on whether the production run is make-to-order or make-to-stock.

the current logic of assigning orders to individual and production equipment.

For the make-to-order jobs (see the JIT manufacturing section of Figure 1), customers place the orders with sales associates either out in the field or at the corporate office. The orders are entered into the order-entry system, and then routed based on the sales associate's best estimate of which facility in his/her geographical proximity will be able to process the order within the promised due-date time frame. Often, a client places orders with a sales associate in one region for delivery to remote sites. In this case, it would make sense to produce the orders at a plant closer to the actual demand destination. This policy of allowing a sales associate to determine where the order should be produced consistently led to suboptimal order assignment to production locations and individual presses. For make-to-stock orders (see the JIT distribution section of Figure 1), products are produced in anticipation of expected demand from steady customers (managed accounts). Thus, at the time of the order entry, the availability of on-hand inventory is checked, orders are shipped, and reorder (production) decisions are implemented subject to the inventory position falling below the reorder point. The plant for production of the large batch-size jobs is derived using general guidelines like "most efficient production" or "freight-cost minimization." The make-to-stock order is stocked at one or more of multiple company warehouses.

Figure 2 describes the general printing process at these rotary facilities.

Martinich (1997) describes the production process of printing forms at one of SR's plants in great detail. The production process can be divided into three phases: preprint operations, printing operations, and postprinting operations. Depending on the product specification (roll or cut sheet or its variations, number of colors, trifolds, etc.), SR schedules the order for a specific printing press at a specific location. There are more than 150 presses across the six locations across which an order is scheduled; however, the subset of feasible presses depends upon the technical specifications of the order.

The presses may vary in their cost effectiveness for different volumes of a specific job type. Differential setup costs and variable operating costs make the choice of an economical press a function of the order size. The rotary facilities currently process orders of at least 10,000 sheets (smaller orders are produced at another set of smaller-volume, custom-print SR facilities or are outsourced). A production run can be as high as hundreds of thousands of sheets.

Although preprint and postprint operations are important, the printing presses represent one of SR's most significant equipment investments. Because

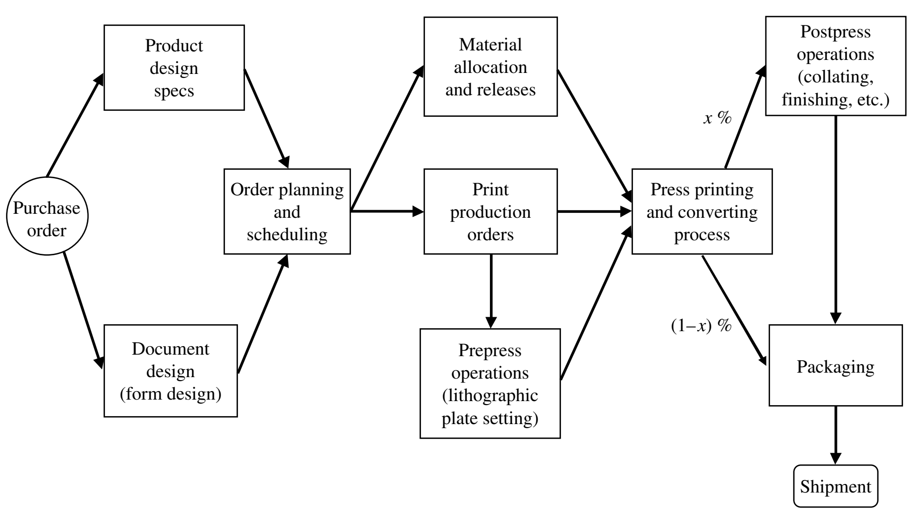

Figure 2: The efficiency of SR's general printing production process depends heavily on the press printing and converting process, which utilizes its key printing press resource.

cost-effective production of orders depends on the assigned press capacities and capabilities, our focus was on helping SR rethink the logic of allocation of orders to presses across the firm's production-distribution network.

Problem Scope

From the outset, we decided to quickly demonstrate to SR's senior management the opportunity for revamping their operations strategy to reflect a more supply chain optimization paradigm and the value of OR in such an effort. Therefore, we carefully managed the scope of the project to reduce the risk of overinvestment without results. Our focus was on a set of tools that would demonstrate the potential value of a change in job-allocation strategy for a meaningful and significant segment of the business. Also, we had to define a segment of SR's business that was representative, high value, and separable from other business segments so that results based on that segment alone were meaningful and implementable as a partial solution to SR's broader challenge.

We took multiple passes at determining which segment of SR's business constituted a meaningful yet manageable portion of the business. For example, we looked at a market segment based on one type of product: print orders to be supplied as rolls. However, we quickly realized that this segment, although high volume and unique in its customer requirements, shared press capacity with numerous other product lines; thus, we could not extricate it easily from other SR product lines. We also considered selecting only the jobs with more than four colors. However, presses with high-color capability are often used to produce low-color work; thus, this segmentation was also undesirable for analyzing press-capacity utilization. We finally settled on segmenting the business by "page width" (e.g., seven-inch checks, or 8.5-inch standard width or 11-inch standard length page). Although some presses can produce products of different widths by switching the roll the paper is on, the cost of changeover is prohibitively high and is done only on an exception basis.

We also had to establish the time frame over which to model. Job-allocation strategy is established over a time interval such as a year or more. However, selecting an analysis time frame that is too long would increase both the data requirements and the modeling complexity. On the other hand, selecting a time frame that is too short would underestimate the impact of busy presses from large jobs in previous time periods. We settled on quarterly analyses, run in sequence (e.g., four quarterly runs identified the year's work). Because SR's business shows regular monthly and quarterly patterns, but relatively little seasonal difference between quarters, we captured the essence of the problem without quadrupling its size. We built smaller-scale models for each quarter and product width. These analyses could be easily replicated for separable lines of business and time intervals.

We also had to understand the impact of postpress operations on press choice to model SR's problem well. Because our focus was press capacity, we did not explicitly model postpress production operations, but introduced these considerations as a constraint into the analyses. Although a press may be qualified to print a job, if postpress capabilities were not available to meet the customer specifications, we disallowed the assignment. We also decided to make production batch sizes exogenous and used historical batch sizes as an input to our analyses. Finally, we considered how to include the "make-buy" decision, and given SR's general overcapacity in these product lines, we treated outsourcing as the option of last resort. Thus, we steered order allocations toward retaining as much work in-house as possible. These complications to press assignment reflect the process of considering the impacts of important but secondary considerations on press assignment, and the challenge of judiciously managing the project scope to facilitate quick turnaround for the project.

Analysis Methodology

Phase 1: Statistical Analysis

We used regression analysis to predict three critical parameters for the optimization and simulation models: production times, production costs, and transportation costs. We identified simple intercept and slope regression estimates, where the intercept captured the fixed cost of setup for each job and the slope expressed the cost and time per unit in the order. However, we had to estimate each relationship separately for each job type and press combination. Job type depends on the features of each job (such as width, length, grade, color, and finish of the paper), number of colors in the job, and job postpress processing requirements (such as cut sheet, trifold, or rolled). We estimated press operating times and operating costs separately. These two relationships are similar, but not identical because of different facility cost structures, materials usage, and equipment operating costs.

We used three years of historical operations data to estimate these relationships. Despite the long time horizon, we still had numerous machines that did not have enough usage with jobs of various features from which to base a regression. One reason is that some presses cannot perform some operations; jobs with such requirements must be excluded from those presses (e.g., a two-color press cannot perform a four-color job). Another reason is the vast number of features and feature combinations that are possible caused so many divisions of the regression data that we ended up with relatively few observations for each feature combination. To rectify this problem, we turned to "feature groups"—sets of job features with comparable operating characteristics on a press. For example, the color of the paper on the press has no impact on its setup or run rate, and there may be dozens of shades of white that the customer might consider as different features. By pooling these features, we were able to estimate more "feature group—press operating regression relationships" and to do so with higher confidence.

We based transportation costs on SR's shipment-rate matrix from its preferred provider. Transportation costs can be accurately estimated based on the zip code of the facility where the press is located and the final shipment-destination zip code (the equivalent of our feature group—press combination in the other regressions). The slope term in this case is the overall weight of the shipment in pounds (rather than sheets) because the paper grade can substantially impact the weight and thus the cost of shipment.

We identified numerous outliers in the historical data that affected production cost and time regression estimates. From casual inspection, these outliers seemed to be the result of special circumstances of the particular order, such as a split order, rework, or data-entry error. Therefore, we considered these outliers to be poor information for planning and decided to remove them. Because of the large number of regressions in the model, we took a two-stage approach to estimating operating costs and time to automate the elimination of these outliers. We first estimated a regression with the full data set. Then, we removed any observation from the data set if it was more than three standard deviations from our predicted regression value and re-estimated the relationship. We removed approximately two percent of the data through this process. Finally, we did not use some predictive regressions for lack of explanatory power or data. We threw out any regression with an r-squared under 0.2 or fewer than five observations. As a result, our regressions were generally quite good. The average r-squared value was 0.83; all f-statistics for equations' predictive abilities and most t-statistics for variable and fixed costs were significant at the five percent level.

Given that our focus in this project was improved assignment of jobs to presses, we wanted to understand whether the estimated intercepts and slopes for each feature group were significantly different between presses. For example, if feature group 1 could be assigned to press A or press B, we tested whether the estimated intercepts and slopes were different for the two presses through the introduction of dummy variables and running paired regressions to compare press performance. Figure 3 illustrates these trade-offs.

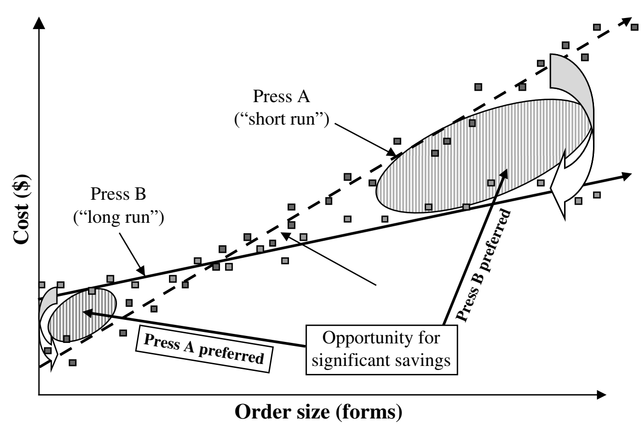

Figure 3: The significance of total cost differences between presses depends directly on the fixed costs and variable costs of each job. However, we found that we did not need to specify the fixed and variable costs of each job; rather, we tested for significant difference in total costs.

Generally, we found that two-thirds of the slope estimates and half of the intercepts were significantly different. Thus, with this analysis, we are able to discern real differences in the presses' capabilities and costs. However, we quickly realized that these comparisons were very costly, and in many cases not necessary for our purposes. If a feature group has 11 press options, each of which requires a paired comparison of both intercept and slope, it leads to 110 comparisons (55 × 2 comparisons!). In fact, we only wanted to compare the significance of the difference of the historically assigned total cost and the model-recommended assignment cost. Further, we were not concerned whether that difference was due to intercept difference, slope difference, or both, but only if the estimated total costs were different for two presses. Finally, this difference may be important for a large job, but not for a small job (or vice-versa); therefore, we did not want to compare the two presses' ability in general, but only for a particular job type and size.

For example, Figure 3 shows a significantly different intercept and slope for a "short-run" press (press A with low fixed cost and high variable cost) and for a "long-run" press (press B with high fixed cost and low variable cost). For very small jobs, the short-run press is significantly better in terms of total cost; for large jobs, the long-run press is significantly better. However, for mid-sized jobs, they are roughly equivalent, and do not result in a statistically significant total cost estimate. Thus, rather than evaluating machines based on their cost profiles for all jobs, we evaluated the statistical difference of the resulting cost on each press of a job of specific size.

We thus opted for a simpler and more relevant comparison: the predicted total cost of the historically assigned press with the total predicted cost of the recommended press. (Of course, this comparison is only relevant if the model recommends a change in press.) We tested the significance of these two predicted costs given the standard error of estimate of each regression. In this way, we gauged the confidence of each of our savings estimates from the models we describe in the next sections.

Phase 2: Optimization Analysis

Our problem falls into the category of "classic production location and distribution cost minimization" with the constraint that each order must be assigned to a particular press, and all orders must be processed. We determined the optimal order assignment to one among multiple presses at different locations to the set of orders over a one-quarter planning horizon. In terms of the goal of minimizing total landed costs across a production-distribution network, the problem we address is similar to the enterprisewide optimization of total landed cost for a grocery retailer reported by Erhun and Tayur (2003). However, their work involved setting up a decision-support system based on replenishment policies for grocery-store products using demands at individual stores. Our problem is about allocating incoming customer demands to the entire set of presses across various locations to minimize total annual system (production and transportation) costs. From our review, we found only one article that integrates order allocation among different machines with subsequent process scheduling and equipment load planning in paper mills (Keskinocack et al. 2002). Our work differs from their study in that, for SR, the rotary printing presses form the primary capital equipment dictating the production costs and bottleneck resources in completing printing orders. Hence, we decided to focus our attention on press assignments without explicitly incorporating capacity constraints of other processing steps.

It was apparent from our statistical analysis that in many cases, the press with the lowest predicted production cost was not utilized in the historical data. However, there are many rationales for this. First, the transportation costs from the most cost-efficient press to the customer location may be prohibitive and overtake any production-cost savings. Second, it is not clear that the most cost-effective press has available capacity to handle all the jobs for which it is most effective.

Hence, we developed an optimization model (see the appendix) based on the classic assignment integer-programming problem (Lawrence and Pasternack 2002) using the following general logic:

(1) For a set of orders to be assigned, identify all feasible presses for each order (from a technical perspective). If there are m jobs and n presses, the assignment matrix array size is $(m \times n)$.

(2) Identify the total cost of production (total material and variable and fixed costs) based on job-press

regressions and transportation costs based on origin-destination shipment-weight regressions for each feasible press for each order.

(3) Assign the orders to the candidate presses to minimize the total system cost subject to capacity constraints of the individual presses, while ensuring all candidate orders are fulfilled.

The optimization is intended to show at a high level that "closest press" and "cheapest press" are not always optimal to provide an alternative recommendation, and to indicate roughly what the cost of such heuristics were. We used the results of the optimization assignment model to determine policy variables, which we tested and validated in the simulation model. To model the fixed-cost portion of the production and transportation-cost function accurately would require a binary integer variable for each job and press pair, and for each origin-destination pair. Given that we had thousands of such possible assignments, the problem was impossible to solve efficiently with binary variables. Hence, we opted for expediency and replicability over precision by eliminating the binary variables and solving the optimization as a straightforward assignment problem. The risk of this approach was that the optimization model may opt to split a job between two presses, and incorrectly charge a portion of each press's fixed costs and setup time. This misstatement would be costly if the occurrence of split jobs was widespread. However, given that each press works on scores of jobs during the quarter, we observed empirically that the model rarely split a job between two presses. (We saw this in less than one percent of the jobs.) Further, in practice one press could be operated overtime, eliminating the split job with minimal adjustment to the model's recommendations. Finally, the primary purpose of the optimization was only to identify strategic rules for job allocation that we would later test in our simulation setting and in practice; therefore, a level of precision that eliminated all occurrences of job splitting was not necessary.

We minimized the material, operating, and transportation costs for a set of customer orders while ensuring that all customer orders were processed without exceeding press production capacities. All production costs and times were the regression-based estimates where available. We allowed the model to consider for assignment only presses with sufficient historical experience with a feature group. The exception to this rule was that we always allowed historical assignment even if that press did not have a cost-prediction equation for that feature group (we used the actual reported historical cost). Even if a feature group had no options for reassignment because of lack of history, we kept the job in the model to accurately reflect the press capacity that the job consumed. We allowed for press overtime assignment or outsourcing production to a third party at a 50 percent production cost premium.

The optimization model, implemented in What'sBest® (version 7.0, Lindo Systems, Inc., http://www.lindo.com), reevaluated historical assignments and described the greatest opportunities for improvements. The model output indicated total planned production, materials, and transportation costs; job assignments to machines and facilities; and total setup and run times and costs for each machine and location in the network. We demonstrated the changes from historical assignments, and the impact of these changes on setup costs, run costs, and transportation costs for a set of orders. We provided reports at the facility, press, and job level.

Finally, as we discussed in the previous section, we tested our optimization results against the statistical uncertainty from our regressions. We evaluated the probability that the estimated savings are indeed significantly different than zero by comparing the standard error of the total cost regression against the estimated cost savings. We further estimated the confidence in the optimization model results by checking the standard error of the total press-time-regression results against the allowable increase and decrease of the press-time constraints, which we found to be very robust because SR generally has excess capacity. But, having capacity, and having it when it is necessary, are quite different. The simulation analysis of the next section describes how we put our optimization recommendations to the test in a stochastic environment.

Phase 3: Simulation Analysis

The optimization model identified the optimal assignment of customer orders in a deterministic environment within an aggregated level of press capacity. We wanted to ensure that:

(1) Optimization results could be distilled into simple, practical business rules that would deliver real savings for SR,

(2) The effectiveness of alternative business rules for allocation of orders could be evaluated along cost-savings and customer-service dimensions, and

(3) These rules were robust in a highly variable and uncertain environment.

To address these questions, we developed a comprehensive simulation using ProcessModel® (version 5.1, http://www.processmodel.com) to model the dynamic order routing, production, and shipment through the order-fulfillment process. The simulation used much of the same input data as the optimization model: (1) machine time constraints, (2) job run-time requirements, (3) total jobs to run, and (4) transportation and production-cost profiles for various job-to-machine assignments.

This simulation allowed for the dynamics of: (1) customer-order arrivals, (2) press availabilities and down time, and (3) the constraint of meeting variable expected ship dates given press outages and availability.

We modeled SR's current operating practices to benchmark actual performance. We started with historical order arrivals to compare closely with historical cost and customer-service levels. In our simulation, each order had certain attributes such as job size, number of colors, sheen of paper, etc., which affect the performance of different presses for the job. These attributes determined both disallowed presses (e.g., a two-color press cannot perform a four-color job), and cost estimates for other presses that had different setup and run times and costs as our regression estimates predicted.

Due to the relatively high setup costs, we did not preempt or split jobs. Although, in practice, SR has some ability to look ahead to short-term forecasts of customer-order arrivals, we assigned each job using a FIFO rule because this is how they are usually processed in practice. In some cases, SR may face a short-run capacity shortage because of a spate of orders or machine downtime, which prevents meeting customer delivery schedules. In these cases, SR will outsource an order to a subcontractor.

The simulation logic entailed assigning an incoming job to a press based on six alternative business rules:

(1) *Historical only*—our baseline scenario explicitly follows historical assignment,

(2) *Least cost only*—the lowest predicted cost press only,

(3) *Least cost only, then subcontract*,

(4) *Lowest cost available, then subcontract*—starting with least cost, moving to successively higher cost presses. If no presses are available, then subcontract,

(5) *Optimization recommendation only*—optimal press recommendation only,

(6) *Optimization recommendation first, then lowest cost available, then subcontract.*

In our baseline simulation, Scenario 1, we allowed only the historical press to be selected. If the historical press could not meet customer delivery dates, no other alternative was considered and customer service suffered. In our Scenario 2, we allowed only the least-cost machine to process a job; thus, by definition, the least-cost solution was attained. In Scenario 3, we looked to the least-cost machine first; if it was not available, we subcontracted the work to an outside service provider to keep internal production capacity free. In Scenario 4, we sent it to the lowest-cost press available and looked to subcontracting only as a last resort; therefore, the order might go to a relatively high-cost press in busy times. Scenario 5 allowed the job to go only to the optimization-based recommended press. Scenario 6 was very similar to Scenario 4; however, instead of lowest-cost press first, the sequence of assignment rules used was: seek the optimum press first; if it is not available, then consider alternative low-cost presses.

In each scenario, we simulated two quarters of historical orders, tracking performance on the final quarter after a one-quarter warm-up period to develop process queues. The simulation results predicted the total production and transportation costs in a consistent way with the optimization model. In addition, the simulation tracked dynamic variables such as order-cycle time, percent of orders late, average lateness of orders, and the number of orders assigned to the first-choice press. We compared each of these results against the baseline to evaluate the effectiveness of the business rule employed.

Results and Recommendations

We applied our suite of analytic tools on the 11-inch width business segment, which represents about eight percent of SR's total revenues. Regression estimates were based on data from the 2001Q3–2004Q2. In our statistical analysis, we developed over 700 regressions expressing cost and run-time relationships for feature groups and presses.

Optimization Results

We ran three optimization models for the historical quarters 2003Q4–2004Q2. In each optimization run, we assigned approximately 1,300 jobs to 33 presses at four locations. The optimization results were consistent from quarter to quarter, making the same general recommendations in each run.

The optimization model identified an opportunity for SR to save 3.5 percent of its production, materials, and transportation costs from its current job allocation. Most of those savings came from labor costs from improved allocations (10 percent savings); some came from material cost savings (one percent savings). Transportation costs were projected to increase one percent in these model runs because the optimization model traded more efficient production for longer shipping distances. Thus, we identified a $1.5 million dollar opportunity within this segment of the business. These results translate into a conservative estimate of over $10 million in annual savings across SR's entire business base.

The output reports also indicated to SR which presses gained or lost work, and which production facilities were most affected in terms of total work to be done and the mix of order types the facilities received. Because we included transportation in the model, we could also glean planned shipment patterns from each facility.

Finally, we also tested the statistical significance of the savings estimate of each job reassignment based on the standard error of estimate of the regression and the estimated cost savings. We found that over 90 percent of the estimated savings opportunity was significantly different than zero at the 90 percent level of confidence. This boosted our confidence that, although savings estimates are subject to estimation error, they are highly likely to be profitable changes to the SR cost structure. The confidence estimation surrounding our savings estimates was critical to gaining the credibility of SR management.

Simulation Results

We wanted to assess if the optimization-derived savings estimates were achievable in a dynamic applied setting in which meeting customer commitments is a major priority. In addition, it was important to see if our optimization-based recommendations were superior to simpler rules of thumb (such as lowest-cost machine only or lowest-cost machine then subcontract) in the face of uncertain order arrivals and press availability. We used the order-arrival patterns and press availability from historical data as the basis of our simulation, but allowed for randomness to test the true robustness of our recommendation.

We thus evaluated each of our six business rules for effectiveness. From Figure 4, we first note that the dynamics of customer-order arrivals, machine downtime, and availability drive a significant need for flexibility. Thus, following Strategy 1, 2, or 5 (historical, least-cost, or optimum only) is not feasible from a customer-satisfaction perspective and demonstrates the need for flexible order-assignment rules. We thus focused on the more flexible rules that allow customer-service goals to be met.

Scenario 3 (least–cost, then subcontract) was designed to drive work to the lowest-cost press, but

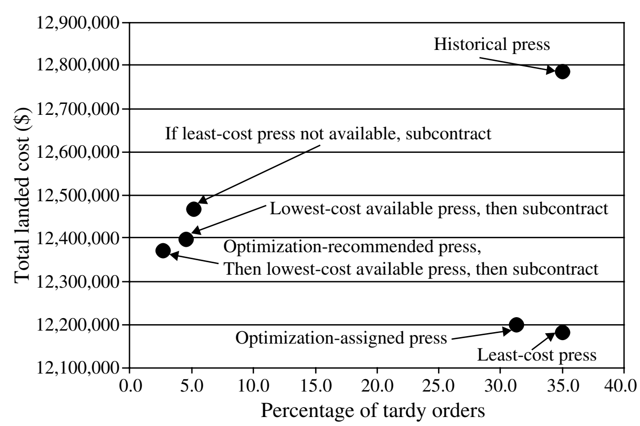

Figure 4: The simulation outcomes for a quarter using six different order assignment rules for a sample SR business segment showed that optimization results achieved improvements in cost, but simulations revealed optimization rules must be augmented with heuristic assignment rules to maintain service levels.

avoid the domino effect that may result from choosing other presses if the first-choice press is unavailable (as in Scenarios 4 and 6). This approach proved to be more costly than keeping the work in-house because SR's subcontractors tend to be more expensive than in-house production. Because SR and the printing industry are in a condition of general oversupply, the domino effect proved to be minimal. Scenario 4 outperformed Scenario 3 in terms of both cost and service. As a result, subcontracting should be considered only as a last resort because of its higher cost and detrimental effect on SR's return on investment in press capacity.

Scenarios 3 and 4 did not consider the capacities of the presses nor the demand for their services; thus, the lowest-cost press was often unavailable. Although the lowest-cost press was the goal in these scenarios, the resulting cost tended to be higher than the optimization-based rule because there were 45 percent more exceptions to the first-choice press. Because the optimization-based recommendation allocated work to a press considering availability, total customer demand, and the opportunity cost of a poor assignment, fewer and less costly exceptions occurred. This ability to follow a more feasible plan (or rule of thumb) allows both slightly lower costs and better customer service.

Using the optimization-based recommendations of Scenario 6, we identified a 2.6 percent savings opportunity. This savings opportunity is less than that identified in the optimization model because numerous orders are routed to less-preferred presses on a cost basis because of unavailability of the desired press prior to customer due-date constraints. However, viewed another way, 70 percent of the savings opportunity identified by the optimization was verified in simulation. Interestingly, using a totally different methodology, simulation, we found approximately the same realized savings as our "high-confidence" savings estimate based on optimization and statistical analysis. This confirmation further enhanced our confidence in the results.

Optimization-recommended assignment (Scenario 5) and least-cost assignment (Scenario 2) had similar results because of the excess capacity that is common in the industry. However, to the extent they are different, optimization-based recommendations improved on the least-cost assignment heuristic by, in effect, saving room on some low-cost presses for orders that arrive later. This was evident from a higher rate of first-choice presses being available in the optimization-recommended press-assignment scenario (with a 100-percent first-choice press assignment) versus least-cost press-assignment scenario (with a 94 percent first-choice press assignment).

Client Impact and Implementation

We formally presented the detailed findings from the multidimensional analyses to SR's top-level Executive Leadership Team (ELT). The ELT members evaluating our recommendations included the CEO, CFO, COO, Vice President for Business Excellence, and the Vice President of Marketing. The scientific basis of analyses, the communication of results in a managerial language, and demonstration of the evidence of the accuracy and power of the OR tools we used to support the recommendations enabled the ELT to approve this "re-thinking." This represented a dramatic departure from their traditional overall operations strategy and use of production infrastructure. This was very important within the context of the organization with 90-plus years of heritage in leading in this industry and now being challenged to maintain its prominence.

The ELT received our recommendations enthusiastically; as a result, it has revised SR's strategic approach to its business. Prompted by the opportunity we identified, SR has changed its business processes in several ways over the last two years. Its organization and policies, plant configuration, and computer and data systems have all changed as a function of its shift in strategy.

SR modified both its organization and policies as a result of its new strategy. First, the decision for job assignment (or job sourcing) including outsourcing decisions on individual jobs is no longer made in the sales organization. Jobs are now based on centralized knowledge of (internal and external source) capabilities, total landed costs, and press availability across the network. As a signal of its new emphasis on managing its total landed costs, SR has established a new executive position, Chief Supply Chain Officer, who reports to the CEO and whose mission it is to oversee and control SR's order-routing decisions across the entire organization (beyond the division studied here). Finally, guidelines for job allocation based on total landed cost efficiency have been established between SR's traditional print business and its newer digital print business to improve its cost structure.

Improved routing created additional capacity which allowed SR to shut down plants and reallocate presses. We identified a plant and set of presses that we highly underutilized in our optimization that assumed all presses were fixed in location. As a result, a four-day per week plant was closed and its presses consolidated into a 7 by 24 plant, increasing capacity and lowering fixed costs. SR has reassigned presses under one roof based on common products and capabilities to reduce the organizational challenges of allocating jobs between plants. From an organizational perspective, a single plant manager can make better trade-off decisions when not faced with protecting his workloads for fear of "losing the job to another plant." While centralizing of these presses may in some cases create longer shipments and higher transportation costs, these costs pale in comparison to the production cost efficiencies gained.

In part justified by the results of this project, SR's production control system (PRIMAC) is being upgraded to include a new million-dollar, real-time decision-support system for job assignment, outsourcing, overtime, etc., to optimize and balance the workflow for enterprise-wide lowest landed production and transportation cost. Finally, SR is developing new automated data sources (i.e., "dashboards") to help plant managers more easily see their performance relative to other plants to communicate and benchmark production efficiency.

As a result of these changes, SR's ability to meet its customer's needs has improved while its cost structure has fallen. Although it is impossible to isolate the impact of a change in strategy in an ever-changing business environment, we note SR's net income was up over $31 million in 2005 compared to 2004 on sales that increased only $10 million; most of the improvement in performance was due to cost savings. SR now keeps more of its print work in-house, which is a testimonial to the effectiveness of improved order routing

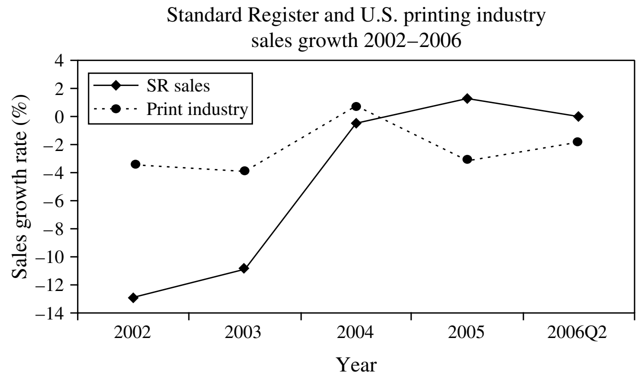

Figure 5: SR sales lagged the print industry until 2004; since then, SR has grown while the rest of the industry has contracted.

on SR's own equipment. SR has been able to maintain or expand its profit margins despite downward price pressures, which has made it more competitive. Figure 5 shows SR's traditional print business. It was shrinking faster than the market before 2004. It has started to grow while the print industry's total sales continue to shrink.

Implementation Challenges

Perhaps our greatest challenges were organizational and behavioral. Centralizing the job-assignment decision faces resistance. Marketing and sales personnel have a relationship with local plant managers where their jobs have traditionally been assigned. Giving the job to a distant lower-cost press, or one with more current or projected capacity, removes some of the interpersonal comfort that had governed such decisions. In addition, it is not always clear to operational personnel at the local-plant level why central planners would "take work away from us," especially if it clearly results in more transportation costs. The answer is in a more global perspective on production costs and capacity; however, this is challenging to present. It is more challenging to overcome subjective objections. Our analyses and demonstrated net benefits convinced the top management and enhanced SR's global operational perspective. Because of the scope of the project and its organization-wide implications, implementation was gradual and time consuming and required our extended involvement over many months. During this time frame, we had to navigate organizational-structure changes and transitions in key management positions to ensure the sustained momentum of implementation of the recommendations. On the positive side, implementation of the strategy, despite changes in key management positions, provided additional evidence of the real value of the new approach to SR.

Concluding Remarks

SR is a prominent player in the print industry, which is facing competitive pressures for low-cost production and pricing in the print-production business. Employing a trio of OR techniques, we demonstrated the value of rethinking SR's operations strategy to the new paradigm of a systematic order allocation to minimize total aggregate landed cost of their production-distribution network. The systematic and comprehensive analyses and their presentation in managerial language helped us convince the executive leadership team at SR of the value of this new strategy. The project was not sold as an OR project replete with OR tools but rather as a solution to a real management challenge in operational terms. The discussion of this project's evolution and its broad influence on SR's senior leaders and operations managers makes a strong case for the value of effective application and dissemination of seemingly intractable OR tools to solve strategic managerial challenges (Ahire 1997).

In the last decade, the OR/MS field has come under heavy criticism for its inability to adequately prove its strategic contributions to the competitiveness of organizations and industries (Corbett and Van Wassenhove 1993, Geoffrion 1992, Meredith 2001, Samuelson 1996). In our view, our effort contributes to the OR practice literature to address these challenges in many ways. First, it provides insights into why careful scoping of an actual operations-strategy challenge is critical to yielding meaningful problem granularity and how it can be effectively accomplished by working closely with the client organization's key participants to understand the goals of analysis. Second, it illustrates the need to apply "needed" OR methods to evolve usable recommendations. Finally, it highlights the positive outcomes of the diligence in convincing senior executive leadership of the worth of OR as a strategy facilitator for competitive performance improvement.

Appendix

Notation

TLC—total landed cost of assignments to be minimized.
N—total jobs to complete; from a period in history.
M—total presses available.
i—index on jobs.
j—index on machines.
a_{ij}—decision variable; whether to assign job i to press j.
PC_{ij}—predicted production cost of assigning job i to machine j.
TC_{ij}—predicted transportation cost of assigning job i to machine j.
PT_{ij}—predicted production time for job i on press j.
TT_j—total processing time available for press j.

Analyses

In Phase 1, we used regressions to predict PT_{ij} and TC_{ij} from historical data.

In Phase 2, we ran an optimization model for job assignments to presses as follows:

$$\text{Min} \quad TLC = \sum_{i=1}^{N}\sum_{j=1}^{M} PC_{ij} * a_{ij} + TC_{ij} * a_{ij}$$

$$\text{subject to} \quad \sum_{i=1}^{N} PT_{ij} * a_{ij} \le TT_j,$$

$$\sum_{j=1}^{M} a_{ij} = 1.$$

In Phase 3, we simulated the process, predicting TLC and service levels within the context of uncertainty of order arrival rates and timing, PT_{ij} and TT_j.

References

Ahire, S. L. 1997. Management science–total quality management interfaces: An integrative framework. *Interfaces* **27**(6) 91–105.

Corbett, C. J., L. N. Van Wassenhove. 1993. The natural drift: What happened to operations research? *Oper. Res.* **41** 625–640.

Erhun, F., S. Tayur. 2003. Enterprise-wide optimization of total landed cost at a grocery retailer. *Oper. Res.* **51** 343–353.

Geoffrion, A. M. 1992. Forces, trends, and opportunities in OR/MS. *Oper. Res.* **40** 423–445.

Keskinocack, P., F. Wu, R. Goodwin, S. Murthy, R. Akkiraju, S. Kumaran, A. Derebail. 2002. Scheduling solutions for the paper industry. *Oper. Res.* **50** 249–259.

Lawrence, J. A., Jr., B. A. Pasternack. 2002. *Applied Management Science: Modeling, Spreadsheet Analysis, and Communication for Decision Making*, 2nd ed. John Wiley & Sons, New York.

Martinich, J. S. 1997. *Production and Operations Management: An Applied Modern Approach*. John Wiley & Sons, New York, 79–87.

Meredith, J. R. 2001. Reconsidering the philosophical basis for OR/MS. *Oper. Res.* **49** 325–334.

Samuelson, D. A. 1996. Promising future, disappointing present. *OR/MS Today* **23**(April) 42–45.

Steve McDonell, VP Supply Management and Engineering at Standard Register Company, wrote in a memo dated December 15, 2006:

"Two years ago Drs. Ahire and Gorman completed a project that validated an opportunity for us to realize significant cost savings through improved customer order assignment to our presses. David Dwiggins and Oleh Mudry (co-authors on the paper submitted to you) were Standard Register's coaches for the project. We were impressed with the analysis and the opportunity the results identified.

"Standard Register is undergoing a significant change in direction in its approach to handling customer orders. Our sales force and production facilities were organized geographically, with our sales force working closely with local production managers on delivering orders to our customers. This effort was one of many that led to the creation of a Chief Supply Chain Officer Position to formalize the new approach organizationally.

"Our initial experience has led us to combine presses of like capabilities in a single plant so we can better manage production capabilities and capacities. Although this configuration leads to higher transportation costs, the production efficiencies gained outweigh these costs.

"Finally, we are building new dashboards for our plant managers and a new production job scheduling system, 'PRIMAC,' to help us manage our customer orders efficiently centrally. These systems improvements will create the right information and decision capabilities to minimize the total landed costs while maintaining customer service levels.

"Standard Register's financial performance has improved over the last two years, despite declining industry sales and revenue statistics. The efficacy of the approach demonstrated by Drs. Ahire and Gorman made a convincing argument for multimillion savings across major segments in the high-volume rotary production business. We believe our new strategy on our approach to operations instigated by this project is a key contributing factor to achieving improvement and is pivotal to our continued success."