

Process Location and Product Distribution with Uncertain Yields

Felipe Caro, Kumar Rajaram

UCLA Anderson School of Management, Los Angeles, California 90095
{fcaro@anderson.ucla.edu, krajaram@anderson.ucla.edu}

Jens Wollenweber

Dr. Staedtler Transport Consulting, 90469 Nuremberg, Germany, wollenweber@web.de

We present a framework to analyze the process location and product distribution problem with uncertain yields for a large multinational food processing company. This problem consists of selecting the location of processes, the assignment of products, and the distribution of production quantities to markets in order to minimize total expected costs. It differs from the traditional facility location problem due to characteristics that are inherent to process industry sectors. These include significant economies of scale at high volumes, large switchover times, and production yield uncertainty. We model the problem as a nonlinear mixed-integer program. A challenging aspect of this problem is that the objective function is neither convex nor concave. We develop an exact approach to linearize the objective function. We present heuristics to solve the problem and also construct lower bounds based on a reduction of the constraint set to evaluate the quality of the solutions. This framework has been used to make process choice and product allocation decisions at the food processing company, and the estimated annual cost savings are around 10%, or \$50 million. In addition, the insights from the model have had a significant strategic and organizational impact at this company. Our framework and conclusions are relevant to other industrial sectors with similar characteristics, such as pharmaceuticals and specialty chemical manufacturers.

Subject classifications: processed food industry; facility location; uncertain yields; model linearization.

Area of review: OR Practice.

History: Received September 2011; revision received December 2011; accepted March 2012. Published online in *Articles in Advance* October 9, 2012.

1. Introduction

The worldwide production of processed food reached \$4.3 trillion in 2010 (Euromonitor 2011). This figure is poised to increase due to population growth and consumption patterns that tend toward higher value food products as the developing economies become richer (Gehlhar and Regmi 2005). In the United States, food manufacturing accounts for more than 10% of all manufacturing shipments (U.S. Dept. of Commerce 2006). Despite recent mergers and acquisitions, the industry remains competitive, and challenges abound, many of them involving a strong operational component (Scanlon 2006, Donovan 2009, Arend 2011). In recent years, several successful applications of operations research (OR) have been reported for short-term production planning in the processed food industry (Rajaram et al. 1999, Rajaram and Karmarkar 2002, Mehrotra et al. 2011). Another important problem that has received less attention is deciding where to open (or build) processes, choosing which product should be assigned to what processes, and the transportation of production quantities to markets in order to minimize total expected costs. This design/tactical problem is usually solved once or twice a year, and it also arises in other process industry sectors such

as pharmaceuticals, specialty chemicals, fertilizers, animal feed, paint, and petrochemicals.

We were exposed to this problem based on our interaction with a large multinational food processing company that had more than 400 plants in 63 countries.¹ On average, each plant had seven different *types* of processes such as refining, extraction, separation, modification, etc. This company had the option of locating processes and products across several of their manufacturing plants. In addition, the company often bought or built plants in several different locations. For example, in 2010 they acquired 10 plants with 65 processes in 5 countries. Thus, the top management had to rationalize the choice of where to open processes and decide which products to produce at these processes at least once every six months. These decisions were important because they significantly affected profitability. This company felt there could be improvements in the current procedures to make process location and product allocation decisions by using a more rigorous analytical framework. However, due to the challenges outlined below, they did not understand how best to develop such a framework. They approached the second author in this regard, which led to our involvement.

Deciding on where to open these processes, assigning which product to produce in which processes, and shipping production quantities to markets to minimize total costs is a challenging problem for several reasons:

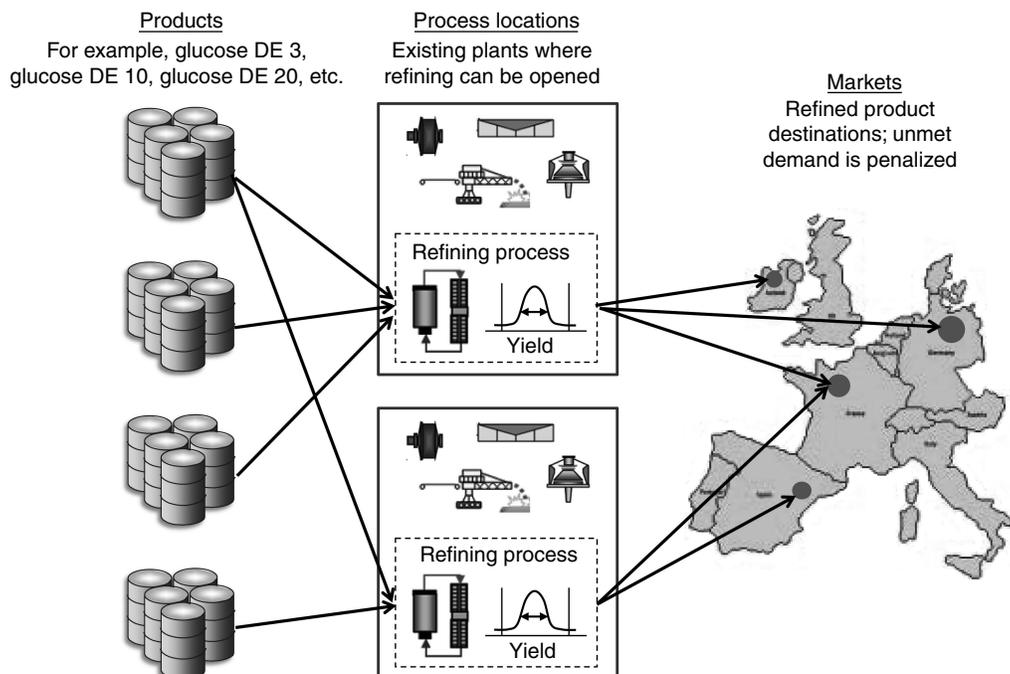
- Opening these processes requires major capital investments in equipment and installation costs. Depending upon the capacity, these costs range from \$15 up to \$60 million. Furthermore, customizing each process to include a particular type of product requires additional and significant product inclusion costs up to \$0.5 million, which typically involved configuring the technology and installing in a manner to match unique product characteristics. Thus to account for these setup and configuration costs, we need to choose the correct number of processes and the right number of products at each process.
- Due to the nature of the product, there are large economies of scale associated with high volumes of production, so it is desirable to maximize production of a product at any given process subject to its capacity limitations. For instance, unit costs at volumes corresponding to 90% process utilization levels typically are 30% lower than unit costs at 60% utilization.
- Because each process itself produces a variety of products, there are significant downtimes when switching between products. In some cases, these setup times could consume up to 20% of the available capacity.
- There is yield uncertainty that results in the actual output being between 70% to 99% of plant production, which leads to backorder costs and inventory holding costs across the product mix.

- There are significant distribution costs in moving the product from the process location to the markets. Depending on the market location, these costs typically account for about 10% to 20% of unit price.

Figure 1 shows a schematic description of the problem. Our model is for a single process type. In the application, the process type was refining, but the model could also be used to locate other types. The potential *process locations* were the existing plants the company owned. Note that these plants already had other processes in place, usually of another type. The specific *products* we were asked to consider were different grades of glucose products because these products constituted a significant proportion of the company’s sales volumes and profits. The production was allocated to *markets* given by countries divided into sales territories. For example, Germany consisted of seven territories, while Belgium and The Netherlands had four territories each. Demand was determined by long-term contracts that stipulated a penalty in case demand was not met.

Glucose products are distinguished by the concentration of sugar or dextrose equivalence (DE) levels. They are produced by refining starch slurry and serve as important inputs across a range of industries including pharmaceuticals, food processing, cosmetics, and paper manufacturing. The refining process consists of several unit steps or subprocesses, and each glucose level might require different equipment and a different configuration of unit steps, which correspond to the inclusion costs and setups in the model, respectively. The output of the refining process is affected by the yield uncertainty, which has two main causes. First, in contrast with manufacturing where

Figure 1. Description of the process location and product distribution problem with uncertain yields.



yield uncertainty usually can be reduced with more training and better procedures, here part of the uncertainty comes from the nature of the biochemical reactions in the process. This uncertainty cannot be eliminated simply by changing the procedures. Second, the raw materials. For instance, it is known that European wheat is different from Canadian wheat, but even two batches of Canadian wheat can be different depending on the location and when it was harvested. In the short run, these two sources of yield uncertainty are beyond the firm's control.

The problem we present naturally falls into the broad category of facility location problems. However, this problem differs from most problems in this category due to characteristics that are unique to the process industry sectors such as significant economies of scale at high production volumes, large switchover times, and yield uncertainty in the production processes. To incorporate these characteristics, we model this problem as a nonlinear mixed-integer program. As we show, these characteristics significantly complicate the solution methods for this model. We develop heuristics to solve this problem and construct lower bounds to evaluate the quality of these heuristics. We also report an extensive computational study using data provided by the large food processing company to validate the framework. Our model has been applied to make process choice and product allocation decisions at a large multinational food processing company, which has led to annual cost savings of 10%, or around \$50 million. In addition, our work has provided several managerial insights that have directly influenced the decision making of the top management and have consequently had significant strategic and organizational impact at this company. These insights include understanding the relationship between process choice and critical aspects within a process such as variability reduction, yield uncertainty, and switchover times.

There exists a wide body of literature for both the deterministic and the stochastic facility location problems. Comprehensive surveys are given by Aikens (1985), Brandreau and Chiu (1989), Drezner (1995), Owen and Daskin (1998), and recently Melo et al. (2009). Reviews exclusively concerning stochastic location problems are Louveaux (1986), Berman and Krass (2002), and Snyder (2006). Of particular interest for this work are publications dealing with multiple products, such as Geoffrion and Graves (1974), and nonlinear facility costs, such as Wollenweber (2008). Louveaux (1986) discusses different versions of the p -median and of the stochastic plant location problem (SPLP) in which demands, production, transportation costs, and selling prices can be uncertain. Louveaux and Peeters (1992) present a dual-based heuristic to solve the stochastic SPLP. Laporte et al. (1994) propose a branch-and-cut approach to solve the capacitated facility location problem (CFLP) with stochastic demands. Balachandran and Jain (1976) consider a capacitated version with general cost functions. Daskin (1983) and Ball and Lin (1993) consider coverage location models where

there is a chance that a server is busy and a probability of system failure, respectively.

Some of the more recent developments in stochastic facility location include Shen et al. (2003), who consider inventory decisions and variable demand; Shu et al. (2005), who build on the latter and exploit the special structure of the problem to solve it by column generation; Baron et al. (2008), who consider congestion and decisions on the number of servers at each facility; Drezner (2009) who looks at retail locations considering future demand scenarios; and Snyder and Daskin (2005), Qi and Shen (2007), and Cui et al. (2010), who consider reliability models for systems with facility disruptions. Using the taxonomy described by Aikens (1985), the problem we study can be considered a stochastic yield capacitated facility location problem (SYC-PLP) with nonlinear costs and side constraints. As mentioned in Aikens, this is a very complex and under-studied problem.

Despite the practical relevance and complexity of the SYCPLP, we found nothing in the academic or managerial literature that specifically addresses this problem for the process industry sectors. In this context, our objective is threefold. First, we establish a rich process location and product distribution problem, through which we model relevant characteristics of the process industry sectors like economies of scale, switchover times, and yield uncertainty. By creating awareness of the problem we hope to stimulate further research. Second, we develop an exact approach to linearize a nonlinear model, and we use it to develop bounds and heuristics. This approach serves as an alternative to the piecewise linear approximation commonly used in practice, and it represents our main methodological contribution. Third, to the best of our knowledge, this is the first reported implementation of process location optimization methods to the food processing industry. Several of the techniques we employ to solve the problem—such as the subgradient updating rule and our method to reduce the constraint set based on the law of iterated expectations—can provide guidance to practitioners facing similar challenges in other applications.

This paper is organized as follows. In the next section we formulate the model and formalize its complexity. In §3 we provide an exact approach to linearize the model. In §4 we provide lower bounds by reducing the constraint set based on the law of iterated expectations. In §5 we generate feasible solutions and upper bounds using three heuristics. The first is based on methods used by practitioners, and the other two are based on greedy and linear programming (LP) methods. In §6 we describe a computational study with data from the large food processing company. In §7 we discuss the implementation of our methods and also describe their strategic and organizational impact at this company. In the concluding section, we summarize our work and provide future research directions. An electronic companion to this paper is available as part of the online version at <http://dx.doi.org/10.1287/opre.1120.1087>. It contains the proofs and additional material.

2. Model Formulation

We formulate the problem of choosing where to locate processes with uncertain yields. It is important to reiterate that the problem we consider is *not* where to locate plants or facilities. These are decided on a long-term basis (i.e., 5 to 10 years) and for our purposes are taken as given. However, at each plant, processes are opened and products are allocated to processes and destinations depending on market conditions on a semi-annual basis. This decision is the focus of our work.²

We begin by introducing the model notation. We use indices $i \in I$, $j \in J$, and $k \in K$ to represent processes, markets, and products, respectively. The demand for product k at each market j is deterministic and equal to d_{jk} . Here d_{jk} is calculated as the total demand across a planning horizon of six months. Unmet demand is penalized by a factor g_{jk} per unit. The transportation cost of product k from process i to market j is given by c_{ijk} . On the production side, the fixed cost of opening process i is e_i and its capacity is s_i . Assigning product k to process i incurs a fixed inclusion cost f_{ik} and requires a setup that consumes b_k units of capacity. Also, let a_k be the amount of capacity required to process one unit of product k . The variable cost is given by the function $p_{ik}(x)$. To consider economies of scale in the production process, we assume that $p_{ik}(x)$ is increasing and concave with $p_{ik}(0) = 0$, and $\kappa_{ik}(x) := p_{ik}(x)/x$ is decreasing.³ Note that all the parameters are assumed to be nonnegative.

In terms of the decision variables, let z_i be a binary variable that equals one if process i is opened. Similarly, let y_{ik} equal one if product k is assigned to process i . We combine the processing and transportation decisions into one decision variable: x_{ijk} . This amount is the planned quantity of product k to be processed by process i and destined to market j , which is a continuous variable, and let w_{ik} be the random yield. Hence, the quantity actually shipped to market j is $w_{ik}x_{ijk}$. The precise value of w_{ik} is unknown in advance, but we assume that the probability distribution of the different possible values is available.

The SYCPLP can be represented by the following nonlinear, mixed-integer program:

$$\begin{aligned}
 \text{SYCPLP} \quad \min \quad & \sum_{i \in I} e_i z_i + \sum_{\substack{i \in I \\ k \in K}} f_{ik} y_{ik} + \sum_{\substack{i \in I \\ k \in K}} p_{ik} \left(\sum_{j \in J} x_{ijk} \right) \\
 & + \sum_{\substack{i \in I \\ j \in J, k \in K}} c_{ijk} x_{ijk} \\
 & + \sum_{\substack{j \in J \\ k \in K}} g_{jk} \mathbb{E} \left[d_{jk} - \sum_{i \in I} w_{ik} x_{ijk} \right]^+, \quad (1) \\
 \text{subject to} \quad & \sum_{j \in J} a_k x_{ijk} \leq (s_i - b_k) y_{ik} \quad \forall i \in I, k \in K, \quad (2) \\
 & \sum_{j \in J, k \in K} a_k x_{ijk} + \sum_{k \in K} b_k y_{ik} \leq s_i z_i \quad \forall i \in I, \quad (3) \\
 & x_{ijk} \geq 0 \quad \forall i \in I, j \in J, k \in K, \quad (4) \\
 & y_{ik}, z_i \in \{0, 1\} \quad \forall i \in I, k \in K. \quad (5)
 \end{aligned}$$

Objective function (1) minimizes total expected costs by the appropriate choice of x_{ijk} , y_{ik} , z_i and consists of five terms. The first term represents the fixed cost associated with opening each process. The second term is the inclusion cost incurred if product k is assigned to process i . The third term represents the variable production cost. The fourth and fifth terms are the expected transportation and penalty costs, respectively. The expectation is taken with respect to the random vector $w_k := (w_{ik})_{i \in I} \forall k \in K$, where $[x]^+ = \max\{0, x\}$. Constraints (2) couple flow and production because the binary variable y_{ik} has to be set to one if product k is assigned to process i . Constraints (3) ensure that production and switchover times do not exceed the capacity available at process i . Note that this constraint incorporates the possibility that each product can have different and multiple setups during the planning horizon. Finally, constraints (4) and (5) define variable domains.

Note that we formulate our model as a single-stage problem in which the location and product assignment decisions are made simultaneously with the processing and transportation decisions. We found this to be a reasonable assumption in our application and in other process industries because there is usually little volume or mix flexibility once a process has been opened and configured to produce a particular set of products. Furthermore, the shipping quantities to the different markets are contracted with the carriers before the realization of the actual yields. Indeed, the expected yield \bar{w}_{ik} for product k if assigned to process i is embedded in the transportation cost c_{ijk} for each market j .

When the support of w_{ik} has a positive lower bound w_{ik}^{\min} , the formulation above can be tightened by considering the constraint $w_{ik}^{\min} x_{ijk} \leq d_{jk} y_{ik} \forall i \in I, j \in J, k \in K$, which can either replace or be used simultaneously with constraint (2). It is also worth noting that our formulation can be easily extended to consider a holding or left-over cost h_{jk} that is incurred whenever the quantity processed exceeds demand due to the random yield. In that case, it suffices to redefine c_{ijk} and g_{jk} as $c_{ijk} + h_{jk} \bar{w}_{ik}$ and $g_{jk} + h_{jk}$, respectively. This transformation follows from the identity $x = [x]^+ - [-x]^+$. Using the identity $\min\{x, d\} = d - [d - x]^+$, it can also be verified that our problem is equivalent (up to a constant) to a problem in which the objective is to maximize profits.

As mentioned before, the objective function (1) of the SYCPLP is nonlinear, with the third and fifth terms being concave and convex, respectively, in variables x_{ijk} . This structure becomes a major obstacle to solve the problem in general. However, in practice, the random yield is either discrete—which was the case in our application (see Figure 2 in §6 for a representative example)—or has to be discretized for computational purposes. This fact can be used to reformulate the model as follows. Let T_{ik} be the set of yield scenarios for process i and product k , and let T_k be the cartesian product across processes. We use the index $t \in T_k$ to denote a scenario, where $t := (t_i)_{i \in I}$ and $t_i \in T_{ik} \forall i \in I$. Let $w_k^t := (w_{ik}^t)_{i \in I}$ denote the random

yield vector under scenario $t \in T_k$, and let $\Pr(w_k^t)$ be its probability. Then, in the objective function, the expected unmet demand for product k in market j , which is given by $\mathbb{E}[d_{jk} - \sum_{i \in I} w_{ik} x_{ijk}]^+ = \sum_{t \in T_k} \Pr(w_k^t)[d_{jk} - \sum_{i \in I} w_{ik}^t x_{ijk}]^+$, can be substituted by $\sum_{t \in T_k} \Pr(w_k^t)u'_{jk}$ and the following two constraints:

$$\sum_{i \in I} w_{ik}^t x_{ijk} + u'_{jk} \geq d_{jk} \quad \forall j \in J, k \in K, t \in T_k, \quad (6)$$

$$u'_{jk} \geq 0 \quad \forall j \in J, k \in K, t \in T_k. \quad (7)$$

The auxiliary variable u'_{jk} represents the unmet demand under scenario $t \in T_k$. This substitution is valid under the natural assumption that $g_{jk} \geq 0 \quad \forall j \in J, k \in K$. In fact, because we are solving a minimization problem, and given constraints (6) and (7), the variable u'_{jk} is guaranteed to be equal to $\max\{d_{jk} - \sum_{i \in I} w_{ik} x_{ijk}, 0\}$, $\forall j \in J, k \in K$ such that $g_{jk} > 0$.⁴ With this substitution, the objective function becomes concave in the x_{ijk} variables, and therefore an optimal solution is guaranteed to exist among the extreme points of the convex hull of the feasible set. This property allows the analysis of the resulting model as discussed in §3. From now on, we refer to the model in which the aforementioned substitution has been made as the SYCPLP.

As expected, the penalty parameter g_{jk} plays an important role in the SYCPLP. In fact, when g_{jk} is too low, it is optimal to leave all demand unmet. Indeed, if g_{jk} is lower than the production and transportation costs, then $u'_{jk} = d_{jk}$, $\forall t \in T_k$. On the other hand, when g_{jk} becomes sufficiently large, demand is met in all scenarios and $u'_{jk} = 0$, $\forall t \in T_k$. The latter is equivalent to solving a deterministic process location problem for the worst-case scenario. Clearly, for intermediate values of g_{jk} , which occur in practice, some—but not all—scenarios are penalized (i.e., $u'_{jk} > 0$ for some t). These are the cases for which our formulation is most relevant.

From the previous discussion, one would expect that in terms of complexity the SYCPLP is as hard as solving a location problem without yield uncertainty. That is the case as shown in the following proposition. Note that all formal proofs are provided in the electronic companion.

PROPOSITION 1. *The SYCPLP is NP-hard.*

The proof of Proposition 1 considers the special case with zero production cost, which is sufficient to show the result. In general, production costs are nonlinear, which becomes an additional challenge in the SYCPLP. In the next section we present an exact approach to linearize the model.

3. Model Linearization: An Exact Approach

In the presence of economies of scale, the variable cost functions $p_{ik}(\cdot)$ are (strictly) concave, which makes the problem nonlinear. To deal with this, a common approach

is to use a piecewise linear approximation to transform the problem into a mixed integer program (MIP). This approach requires some trial and error to determine the right number of breakpoints in the piecewise linear function. Moreover, it is an approximation that at best can provide a lower bound to the optimal value. In contrast, here we present an exact approach based on duality theory to linearize the SYCPLP so it can be solved as a MIP with no loss of optimality.

First, we introduce an auxiliary variable $v_{ik} \geq 0$ that represents the total amount of product k processed by process i , and we add the following linking constraints:

$$v_{ik} \leq \sum_{j \in J} x_{ijk} \quad \forall i \in I, k \in K, \quad (8)$$

$$v_{ik} \geq \sum_{j \in J} x_{ijk} \quad \forall i \in I, k \in K. \quad (9)$$

Clearly, with these two constraints we can replace $p_{ik}(\sum_{j \in J} x_{ijk})$ by $p_{ik}(v_{ik})$ in the objective function (1). The next step is to relax constraint (9) using multipliers $\mu_{ik} \geq 0 \quad \forall i \in I, k \in K$.⁵ Then we obtain the following relaxed problem that we denote $G(\mu)$:

$$\begin{aligned} G(\mu) := \min \quad & \sum_{i \in I} e_i z_i + \sum_{\substack{i \in I \\ k \in K}} f_{ik} y_{ik} \\ & + \sum_{\substack{i \in I \\ k \in K}} \left(p_{ik}(v_{ik}) + \mu_{ik} \left(\sum_{j \in J} x_{ijk} - v_{ik} \right) \right) \\ & + \sum_{\substack{j \in J, k \in K}} c_{ijk} x_{ijk} + \sum_{\substack{i \in I \\ k \in K}} g_{jk} \bar{u}_{jk}, \\ \text{subject to} \quad & (2)-(8), v_{ik} \geq 0 \text{ and } \bar{u}_{jk} = \sum_{t \in T_k} \Pr(w_k^t) u'_{jk}, \\ & \forall i \in I, j \in J, k \in K. \end{aligned}$$

The objective function of $G(\mu)$ still has the nonlinear term $p_{ik}(v_{ik})$. To make it linear, we use the fact that $p_{ik}(v_{ik})$ is concave and $G(\mu)$ is a minimization problem. Hence, at the optimum, the variable v_{ik} will be as large or as small as possible—i.e., it will “hit the boundaries”—so we can assume that $v_{ik} = 0$ or $v_{ik} = \sum_{j \in J} x_{ijk}$, $\forall i \in I, k \in K$. Intuitively, the case $v_{ik} = \sum_{j \in J} x_{ijk}$ occurs when $p_{ik}(\sum_{j \in J} x_{ijk}) < \mu_{ik} \sum_{j \in J} x_{ijk}$ and $v_{ik} = 0$ occurs when this last inequality is reversed. To distinguish between these two cases, we introduce a binary variable $q_{ik} \in \{0, 1\}$ and link it to the x_{ijk} and v_{ik} variables through the constraints

$$v_{ik} \leq B_{ik} q_{ik} \quad \forall i \in I, k \in K, \quad (10)$$

$$\sum_{j \in J} x_{ijk} \leq \kappa_{ik}^{-1}(\mu_{ik}) q_{ik} + B_{ik}(1 - q_{ik}) \quad \forall i \in I, k \in K, \quad (11)$$

with $B_{ik} := (s_i - b_k)/a_k$. Then in the objective function of $G(\mu)$, we replace the variable production cost $p_{ik}(v_{ik})$ by $\mu_{ik} \kappa_{ik}^{-1}(\mu_{ik}) q_{ik}$, where $\kappa_{ik}(x) = p_{ik}(x)/x$.⁶ Recall that $\kappa_{ik}(x)$ is decreasing by assumption, which implies that κ_{ik}^{-1}

is well defined and is also decreasing. With this MIP transformation we obtain the following modified dual function that is contingent on the value of the multipliers $\mu \geq 0$:

$$\begin{aligned}
 H(\mu) := & \min \sum_{i \in I} e_i z_i + \sum_{\substack{i \in I \\ k \in K}} f_{ik} y_{ik} \\
 & + \sum_{\substack{i \in I \\ k \in K}} \mu_{ik} \left(\kappa_{ik}^{-1}(\mu_{ik}) q_{ik} + \sum_{j \in J} x_{ijk} - v_{ik} \right) \\
 & + \sum_{\substack{i \in I \\ j \in J, k \in K}} c_{ijk} x_{ijk} + \sum_{\substack{i \in I \\ k \in K}} g_{jk} \bar{u}_{jk}, \\
 \text{subject to} & \quad (2)–(8), (10)–(11), v_{ik} \geq 0, \text{ and} \\
 & \quad \bar{u}_{jk} = \sum_{i \in T_k} \Pr(w_i^t) u_{jk}^t, \quad \forall i \in I, j \in J, k \in K.
 \end{aligned}$$

As expected, the multiplier μ_{ik} determines the variable processing cost of product k in process i . Note that if $v_{ik} = 0$, then μ_{ik} serves as a linear cost; whereas if $v_{ik} = \sum_{j \in J} x_{ijk} > 0$, then μ_{ik} determines the fixed cost $\mu_{ik} \kappa_{ik}^{-1}(\mu_{ik})$. However, when we maximize with respect to $\mu \geq 0$, the optimal solution satisfies $\kappa_{ik}^{-1}(\mu_{ik}) = \sum_{j \in J} x_{ijk}$ (see Proposition 2 below). Therefore, the optimal value is the same regardless of whether v_{ik} is positive or not.

Let Z^{NL} be the optimal value of the SYCPLP. From weak duality we know that $G(\mu) \leq Z^{\text{NL}}, \forall \mu \geq 0$, and because the SYCPLP corresponds to the minimization of a nonlinear concave function we should expect a duality gap. Note that $H(\mu)$ corresponds to a MIP problem. It can be shown that $H(\mu) \leq G(\mu), \forall \mu \geq 0$, so $H(\mu)$ also provides a lower bound for Z^{NL} . Remarkably, in our case the following proposition shows that this lower bound is tight and strong duality actually holds.

PROPOSITION 2. *Let $H(\mu)$ denote the optimal value of the MIP transformation and let Z^{NL} denote the optimal value of the nonlinear SYCPLP. Then, $\max_{\mu \geq 0} H(\mu) = Z^{\text{NL}}$ and there exists an optimal solution such that*

$$\mu_{ik}^* \sum_{j \in J} x_{ijk}^* = p_{ik} \left(\sum_{j \in J} x_{ijk}^* \right), \quad \forall i \in I, k \in K. \quad (12)$$

Proposition 2 shows that the nonlinear SYCPLP can be solved as a sequence of MIPs parameterized by the multipliers μ . In turn, the maximization with respect to μ can be done with a traditional subgradient method, although there is a caveat. In this method, the standard rules would dictate the following updating formula:

$$\mu_{ik}^{\psi+1} = \max \left\{ 0, \mu_{ik}^{\psi} - \tau^{\psi} \left(\sum_{j \in J} x_{ijk} - v_{ik} \right) \right\}, \quad (13)$$

where τ^{ψ} is a certain step size. Note that if μ_{ik} is sufficiently large, $\sum_{j \in J} x_{ijk} - v_{ik}$ equals 0, and the multiplier remains unchanged. At this point, we would not be able to improve the solution of SYCPLP using Equation (13).

An alternative is to leverage Proposition 2. Indeed, we can use Equation (12) to obtain an upper bound for the values of the multipliers. On the other hand, the multipliers should at least be comparable to the penalty for not meeting demand, so we can use the latter as a lower bound. Combining these two observations with Equation (13) provides a modified updating procedure that prevents the procedure from stalling and is also more robust because it does not depend on the starting values. The modified updating formula is the following:

$$\mu_{ik}^{\psi+1} = \min \left\{ \max \left\{ \min_{j \in J} g_{jk}, \mu_{ik}^{\psi} - \tau^{\psi} \left(\sum_{j \in J} x_{ijk} - v_{ik} \right) \right\}, \frac{p_{ik} \left(\sum_{j \in J} x_{ijk} \right)}{\sum_{j \in J} x_{ijk}} \right\}. \quad (14)$$

To apply Proposition 2 we need to compute $H(\mu)$ for any given vector of multipliers μ . This computation itself can be challenging because of the extremely large number of variables u_{jk}^t and constraints (6). Indeed, for the instances we consider in §6 where there are 7 yield scenarios and 10 potential process locations, the commercial ILOG CPLEX solver ran out of memory at the preprocessing stage (see §6 for a description of the computational resources used). Note that any other approach to deal with the nonlinear production cost—e.g., a piecewise linear approximation—would also face this issue. To deal with this, in the next section we describe the procedure we used to reduce the number of constraints.

4. Constraint Set Reduction

As mentioned, for real-size instances, the SYCPLP is computationally intractable due to memory requirements. In fact, the number of scenarios $|T_k|$ is $O(n_k^{|I|})$, where n_k is the average number of yield scenarios across processes for product k . In other words, the number of variables u_{jk}^t and constraints (6) increases exponentially with the number of processes $|I|$. To reduce the constraint set, we take advantage of the probabilistic structure of the problem. First, we aggregate constraints (6) over all scenarios for a given w_i^t , and then we divide the inequality by $\Pr(w_i^t)$. This transformation yields a reduced model with fewer constraints/variables denoted SYCPLP-R and for which Proposition 2 continues to hold. Below we provide the complete formulation of SYCPLP-R for a given vector of multipliers μ . Note that in this formulation it is convenient to redefine the variable x_{ijk} to be the fraction of demand d_{jk} processed by process i .

$$\begin{aligned}
 \text{SYCPLP-R} \quad \min \quad & \sum_{i \in I} e_i z_i + \sum_{i \in I, k \in K} f_{ik} y_{ik} \\
 & + \sum_{i \in I, j \in J, k \in K} (c_{ijk} + \mu_{ik}) d_{jk} x_{ijk} \\
 & + \sum_{i \in I, k \in K} \mu_{ik} (\kappa_{ik}^{-1}(\mu_{ik}) q_{ik} - v_{ik}) \\
 & + \sum_{j \in J, k \in K} g_{jk} d_{jk} \bar{u}_{jk}, \quad (15)
 \end{aligned}$$

$$\text{subject to } \sum_{j \in J, k \in K} a_k d_{jk} x_{ijk} + \sum_{k \in K} b_k y_{ik} \leq s_i z_i \quad \forall i \in I, \quad (16)$$

$$\sum_{j \in J} d_{jk} x_{ijk} \leq B_i y_{ik} \quad \forall i \in I, k \in K, \quad (17)$$

$$v_{ik} \leq \sum_{j \in J} d_{jk} x_{ijk} \quad \forall i \in I, k \in K, \quad (18)$$

$$v_{ik} \leq B_i q_{ik} \quad \forall i \in I, k \in K, \quad (19)$$

$$\bar{u}_{jk} = \sum_{t_i \in T_i} \Pr(w_k^{t_i}) u_{jk}^{t_i} \quad \forall j \in J, k \in K, \quad (20)$$

$$w_{ik}^{t_i} x_{ijk} + \sum_{l \in I \setminus \{i\}} \mathbb{E}[w_{lk} | w_{ik}^{t_i}] x_{ljk} + u_{jk}^{t_i} \geq 1 \quad \forall j \in J, k \in K, i \in I, t_i \in T_i, \quad (21)$$

$$x_{ijk}, \bar{u}_{jk}, u_{jk}^{t_i} \geq 0, \quad \forall i \in I, j \in J, k \in K, t \in T, \quad (22)$$

$$y_{ik}, q_{ik}, z_i \in \{0, 1\} \quad \forall i \in I, k \in K. \quad (23)$$

Note that the variables \bar{u}_{jk} and $u_{jk}^{t_i}$ in the formulation above have a direct probabilistic interpretation. Indeed, let U_{jk} be the random variable that represents the amount of unmet demand for product k at market j . Then \bar{u}_{jk} represents the expected unmet demand $\mathbb{E}[U_{jk}]$, and constraint (20) is an application of the law of iterated expectations with $u_{jk}^{t_i} = \mathbb{E}[U_{jk} | w_{ik}^{t_i}]$. Given the applied nature of our model, this reduction is of immense practical importance because it provides a substantial reduction in the number of constraints and variables. In fact, consider a product k ; then for each market j there are $|T_k|$ constraints (6) in model SYCPLP (which is equal to $\prod_{i \in I} |T_{ik}|$ if the yields are independent). On the contrary, there are only $\sum_{i \in I} |T_{ik}|$ constraints (21) in model SYCPLP-R. To illustrate the difference, suppose there are 10 independent processes and $|T_{ik}| = 5, \forall i \in I$. Then for product k , the cardinality of constraint (6) is $5^{10} \approx 10,000,000$ compared to only 50 constraints (21).

If required in an application, the reduction SYCPLP-R can be tightened even further by using alternative partitions of the scenario space T_k to obtain other valid cuts similar to constraints (20) and (21). This constraint generation can be done incrementally in a way that resembles the seminal cutting-stock problem (Gilmore and Gomory 1961). The approach is outlined in the electronic companion. Finally, the following proposition confirms that the reduced problem SYCPLP-R effectively provides a lower bound to the original subproblem.

PROPOSITION 3. *The reduced problem SYCPLP-R is a relaxation of SYCPLP. Therefore, the former provides a lower bound for the latter.*

In our implementation, we solve SYCPLP-R to obtain feasible solutions to the SYCPLP (see §5), and per Proposition 3 we have a lower bound to assess the performance

of the heuristic solutions. Computing this lower bound took at most eight hours for the real-size instances at the food processing company; see §6. However, because SYCPLP-R is still a facility location problem, for other applications it can be a hard to solve. In that case, further decompositions such as Lagrangian relaxation (Geoffrion 1974, Fisher 1981) might be necessary. The details on how to decompose SYCPLP-R are available from the authors.

5. Practitioners, Greedy, and LP-Based Heuristics

This section focuses on calculation of feasible solutions for the SYCPLP. Because the latter has a number of constraints that is exponentially large, the heuristics described below are applied to the reduced problem SYCPLP-R, and the solutions found are then evaluated in the objective function of the SYCPLP. This evaluation is easily done by recomputing the penalty in the scenarios with unmet demand. Because the SYCPLP is a minimization problem, these feasible solutions provide upper bounds.

In what follows we describe a heuristic used by practitioners at this food processing company. Then we develop a greedy approach that fixes the location decision and solves the remaining subproblem. Finally, we propose an LP-based heuristic, which starts with the solution of the LP-relaxation and provides a feasible solution by iteratively fixing the location variables. In the electronic companion these heuristics are formalized in pseudo-code.

5.1. Practitioners Heuristic

To decide which processes to open, the company uses a straightforward approach that we call the practitioners heuristic (PRAC) in which all processes are sorted by their capacity-to-fixed-cost ratios, and then processes are opened sequentially until total market demand is met. Then, to assign products and markets to a process, all market-product combinations $(j, k) \in J \times K$ are sorted by their demand. After that, these combinations are assigned to the process location with least transportation costs. If the remaining capacity $r_{\bar{i}}$ of the closest process location \bar{i} is not sufficient to cover demand in the worst yield scenario, then the portion unmet is assigned to the next closest process location. Note that the practitioners heuristic does not explicitly take into account economies of scale in production, penalties, or inclusion costs.

5.2. Greedy Heuristics: ADD and DROP

Well-known construction heuristics for facility location problems are the so-called ADD and DROP heuristics, which can be applied to the SYCPLP-R as well. Starting from a solution with all processes closed, the basic idea of the ADD approach is to open iteratively one new process and solve the remaining subproblem. The process that provides the highest gain in the objective is opened permanently. The algorithm will open one additional process in

each loop if this improves the best known solution and will terminate if no further improvement is possible. A second greedy approach, the DROP heuristic, could also be applied to the SYCPLP-R. In contrast to the ADD approach, it starts with all processes open and closes iteratively the one that provides the highest reduction in the objective value.

5.3. LP-Based Heuristic

One drawback of the greedy heuristics is that they can be computationally intensive when there is a medium to large number of processes that can be opened. We construct an LP-based heuristic as an alternative way to find upper bounds for the SYCPLP-R. In the first step of this heuristic we relax the integrality condition for the process variables z_i , then solve the remaining problem and provide a feasible solution by rounding the nonintegral z_i^* to 0 or 1. In the second step, we improve the solution through local search. These steps are described in detail below.

Step I. When the process binary variables z_i are relaxed, the remaining problem still has the product inclusion binary variables y_{ik} , but in our case this is not a challenge because the inclusion costs are relatively small compared to the fixed cost of opening a process. In fact, the ratio f_{ik}/e_i is in the order of 10^{-2} for all pairs (i, k) . Therefore, the relaxed problem can be solved using a commercial MIP solver (ILOG CPLEX in our implementation). Often, this solution is not feasible for SYCPLP-R. Therefore, it is crucial to identify a rounding rule that provides a good feasible solution. The simplest rounding rule is to set every $z_i^* \geq 0.5$ to 1 and the remaining z_i^* to 0. After that, the remaining subproblem can be solved and a feasible solution is obtained. Experimental evidence in Wollenweber (2008) shows that better results can be achieved using an iterative approach that fixes only one variable in each iteration. For that we suggest the following rounding rule: set the largest nonintegral variable z_i^* to 1 if its value is higher than a certain

threshold δ ; otherwise, the z_i^* with the smallest value is set to 0. The underlying assumption is that the fractional value of z_i^* is a good indicator of the “worthiness” of opening process i , and the cutoff δ determines the set of processes worth opening and those worth closing. The parameter δ should not be too high because any positive z_i^* that has not been fixed indicates there is still some unmet demand (in our numerical experiments we used $\delta = 0.1$). Because in each iteration one additional location variable is fixed, step I needs no more than $|I|$ iterations in total.

Step II. Once a feasible solution is obtained in Step I, we try to improve it by doing a local search. Specifically, given the incumbent solution, let the ADD-neighborhood be all the solutions with one additional process open. The local search consists in exhausting the ADD-neighborhood until a local optimum is reached, which then becomes the incumbent solution. For each new incumbent solution, a new ADD-neighborhood is generated, and the local search is repeated until no further improvements are made. In our implementation we tested other types of local variations of the incumbent solution, such as dropping one process or swapping two processes, but the additional computational effort did not pay off.

6. Computational Study

The first step in convincing the large food processing company to implement our methods was to test and validate our heuristics and bounds using real data provided by this company. The test data included all input data required by the model for a 10-process, 30-product, 20-market problem. The random yields for each combination of product and process were given by a discrete distribution with seven scenarios. The distributions were independent across processes and products. Figure 2 shows an example of the discrete yield distributions considered. This figure also illustrates the concave shape of the production cost

Figure 2. Examples for the yield distribution (left) and the nonlinear production cost function (right). The axes of the latter have been scaled.

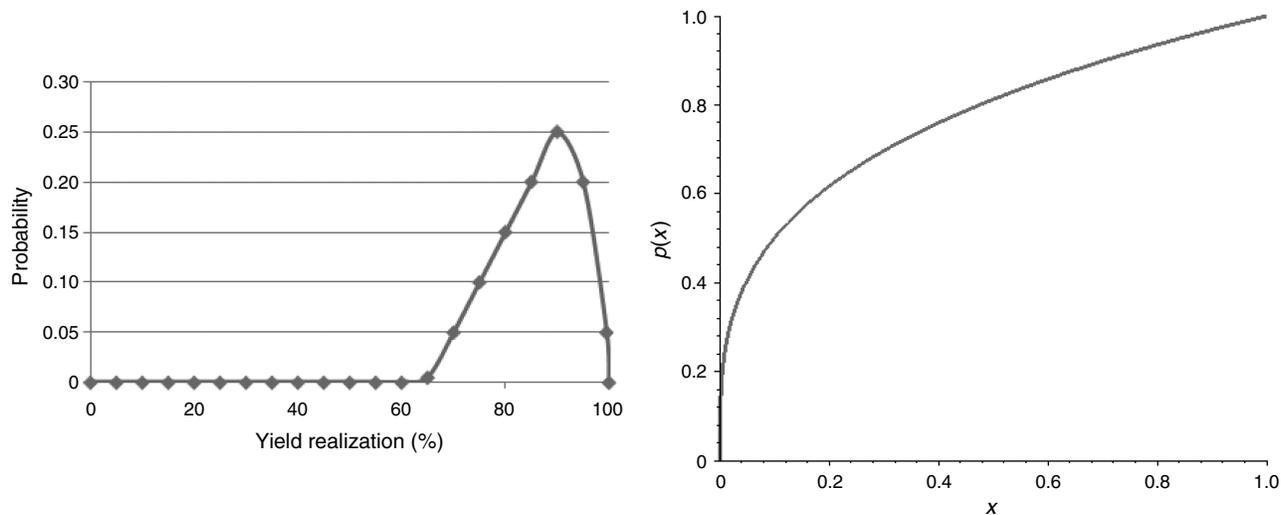


Table 1. Three data sets considered: small, medium, and large.

Instance	No. of processes, $ I $	No. of markets, $ J $	No. of products, $ K $
Small	4	10	8
Medium	10	10	8
Large	10	20	30

functions in the data set. The cost for opening a process e_i ranged from \$30 to \$60 million, while the inclusion cost f_{ik} ranged from \$100 K to \$500 K per product. The average transportation cost c_{ijk} ranged from \$75 to \$250 per unit depending on market destination, and the setups b_k ranged from 0.5% to 20% of average process capacity.

We also considered a small and a medium-size version of the original instance in which we reduced the number of processes, customers, and/or products. This way, we obtained three data sets that are described in Table 1. For each one of the three data sets we modified the capacity and penalty parameters in order to explore a wider range of instances. For that, we defined the following two ratios:

$$CR = \frac{\sum_{i \in I} s_i}{\sum_{j \in J, k \in K} d_{jk}}, \quad PR = \frac{\sum_{j \in J, k \in K} d_{jk} g_{jk}}{\sum_{i \in I} e_i} / \frac{\sum_{i \in I} e_i}{\sum_{j \in J, k \in K} d_{jk}}$$

The capacity ratio (CR) measures the capacity availability with respect to total demand. In the original data, the value of this ratio was 4.4. To test other possible values, we rescaled the capacities to obtain instances with CR equal to 1.5, 2.0, and 4.0. The penalty ratio (PR) measures the average penalty cost with respect to the average unit cost of capacity. Based on discussion with the food processing company, we adjusted the penalty parameters and considered three possible values for PR: 1.0, 1.5, and 5.0.

Given all the possible combinations described above, in total we obtained $3 \times 3 \times 3 = 27$ instances. The algorithms and procedures described in the previous sections were coded in C Sharp (C#). The computational tests were run on a workstation with a 2.4 GHz Intel Core2Duo T7700 processor, 2 GB of RAM, and Windows XP Service Pack 3 as the operating system. For the MIP subroutine calls we used ILOG CPLEX 11.1 with a maximum running time of 200 seconds per call. Furthermore, in the linearization approach we allowed for only 15 subgradient iterations.

Table 2 summarizes the results obtained for the 27 instances considered. The performance of the heuristics are presented in terms of the suboptimality gap with respect to the lower bound obtained from Proposition 3. The first three columns in Table 2 describe the instances. The fourth column reports the gap obtained with the practitioner's heuristic PRAC (see §5), which requires a negligible computational effort and represents the solution that was in practice. Columns five and six show the gap and running time of the ADD greedy heuristic, which consistently outperformed the DROP heuristic. Therefore, we do not report

Table 2. Suboptimality gap with respect to the SYCPLP-R lower bound for the three heuristics: PRAC, ADD, and LP-based.

Instance	PR	CR	PRAC (%)	ADD (%)	Time (sec)	LP-based (%)	Time (sec)
Small	1.0	1.5	36.3	4.6	19	4.1	47
	1.0	2.0	38.8	1.3	6	1.3	40
	1.0	4.0	50.0	15.5	2	15.6	124
	1.5	1.5	29.9	2.6	23	1.9	37
	1.5	2.0	33.1	5.3	18	5.6	48
	1.5	4.0	44.7	10.1	2	10.2	57
	5.0	1.5	26.7	5.6	242	11.1	103
	5.0	2.0	29.7	4.2	26	3.4	72
	5.0	4.0	40.2	7.0	2	7.0	33
		Minimum	26.7	1.3	2	1.3	33
		Maximum	50.0	15.5	242	15.6	124
		Average	36.6	6.2	38	6.7	62
Medium	1.0	1.5	43.8	0.6	24	0.8	69
	1.0	2.0	43.5	1.2	39	1.2	65
	1.0	4.0	41.6	3.6	22	3.9	133
	1.5	1.5	32.9	1.8	708	1.5	2,145
	1.5	2.0	33.6	2.1	622	1.8	162
	1.5	4.0	30.2	1.9	63	2.0	139
	5.0	1.5	27.5	5.1	18,069	5.4	3,015
	5.0	2.0	28.0	5.1	13,574	5.4	1,434
	5.0	4.0	23.6	11.9	778	13.0	55
		Minimum	23.6	0.6	22	0.8	55
		Maximum	43.8	11.9	18,069	13.0	3,015
		Average	33.9	3.7	3,767	3.9	802
Large	1.0	1.5	48.7	5.4	18,406	6.8	6,951
	1.0	2.0	48.5	7.2	10,503	5.9	4,987
	1.0	4.0	48.0	16.1	2,281	14.6	4,167
	1.5	1.5	43.6	5.9	23,240	6.3	8,408
	1.5	2.0	43.2	6.4	20,606	6.4	6,022
	1.5	4.0	42.4	14.9	2,371	13.4	3,239
	5.0	1.5	37.5	25.4	20,696	14.7	8,489
	5.0	2.0	37.0	7.8	23,840	8.0	8,538
	5.0	4.0	32.7	8.7	5,652	10.3	5,704
		Minimum	32.7	5.4	2,281	5.9	3,239
		Maximum	48.7	25.4	23,840	14.7	8,538
		Average	42.4	10.9	14,177	9.6	6,278

the latter. The gap and running time for the LP-based heuristic are shown in columns 7 and 8. In each row, the suboptimality gap in boldface represents the lowest one among the PRAC, ADD, and LP-based heuristics.

Several observations can be made from Table 2. First, although the practitioner's heuristic PRAC is computationally very fast, the suboptimality gaps are large, especially compared to the other two heuristics. Second, for the small and medium-size instances, the ADD heuristic achieves slightly better gaps than the LP-based, but many times it is at the expense of a significantly higher computational effort. Third, for the large instances, the LP-based heuristic seems to dominate the ADD approach because on average, the LP-based heuristic achieves smaller gaps than the ADD approach, and it is also faster. Recall that the LP-based heuristic performs no more than $|I|$ iterations in Step I plus a few local interchanges in Step II, whereas the ADD

greedy approach solves a number of subproblems that is anywhere in between $|I|$ and $|I|^2/2$. That explains why the running time for the LP-based heuristic spans a much narrower interval than the ADD approach. The fact that for large instances the LP-based heuristic also achieves a smaller gap confirms that the LP relaxation, i.e., removing the integrality constraint for the z_i variables, provides a good starting point to generate feasible solutions in location problems. Similar observations have been reported in the literature (see, for instance, Barahona and Chudak 2005, Wollenweber 2008). For the three data sets, the running times in Table 2 tend to increase as PR increases and CR decreases. In other words, as the penalty increases and the capacity is tightened, the process location problem becomes more challenging. The LP-based heuristic is again noteworthy because across the 27 instances in Table 2 it achieved a gap lower than 16% in less than 2.5 hours.

To reduce the suboptimality gap below 10% and convince the management at the food processing company about the efficacy of our methods, we used a branch-and-bound (B&B) procedure that was able to narrow the gap in all the large instances with an average and maximum final gap of 5.4% and 8.1%, respectively (the details of the B&B procedure are available in the electronic companion). The small gaps achieved indicate that the feasible solutions obtained are near optimal, and it also means that the lower bound is tight. This observation was critical in providing confidence to the management that our approach not only could improve upon current procedures but also had a measure of optimality attached to it. This performance assessment in turn convinced them to implement the LP-based heuristics and its insights as described in the next sections. The small gaps also show that the constraint aggregation method based on the law of iterated expectations presented in §4 was quite effective for the 27 instances considered above. For these instances we did not have to strengthen the reduced problem SYCPLP-R with the inclusion of additional cuts. However, it could be necessary for larger penalty values because the unmet demand converges faster to zero in the reduction SYCPLP-R than in the original problem SYCPLP, which adds to the suboptimality gap. This situation is illustrated through a small numerical example in the electronic companion.

To conclude this section, in Table 3 we compare the solutions obtained with the PRAC and LP-based heuristics for the large instance with PR = 5.0 and CR = 4.0. Because

the capacity ratio in the original data was approximately 4.0, the PRAC heuristic applied to this scenario provides the solution in practice prior to this project. The second column in Table 3 shows that both heuristics need to open only three processes to accommodate production. The reason is that the processes in this instance have ample capacity as measured by the CR ratio. Interestingly, there is only one process in common among those that are opened by each heuristic. In fact, while the PRAC heuristic opens the three processes with the lowest capacity-to-fixed-cost ratios, the LP-heuristic opens processes with slightly higher ratios but lower capacity so the total fixed cost decreases. Although the overall number of processes is the same in both cases, the third column in Table 3 shows that the LP-based solution assigns less products per process than the PRAC solution. This allocation is because the PRAC heuristic does not incorporate economies of scale, and also because it aims to cover demand in all possible yield scenarios, as shown in the fourth column. The last five columns show the breakdown of the total cost savings obtained by the LP-based solution. Because the penalty costs increase compared to the legacy PRAC solution, it has a negative (−12%) contribution to savings. However, this is largely offset by the reduction in all the other cost components. In particular, 13% of savings come from lower inclusion costs, 84% comes from lower production cost. Remarkably, the LP-based solution also achieves savings in terms of shipping costs. At first sight, this could be counterintuitive because concentrating production means that some products travel a longer distance to reach market. Although this does happen, the total volume being shipped also decreases, as it attempts to cover only 98% of demand on average (here the average is across yield scenarios). The reduction in volume outweighs the higher per-unit cost, which explains the overall decrease in shipping costs.

Despite the increase in penalty costs, allowing for a very small fraction of unmet demand and concentrating production to take advantage of economies of scale pays off, as shown by the performance of the LP-based heuristic in Tables 2 and 3. This trade-off is absent in the PRAC heuristic because it aims to cover demand in the worst yield scenario, which happens with very low probability. The results of the computational study just described motivated the company to adopt the LP-based heuristic. The outcome of this implementation is described next.

Table 3. Comparing the PRAC and LP-based solutions for the large instance with PR = 5.0 and CR = 4.0.

Heuristic	Processes open	Products per process	Expected unmet demand	Total savings breakdown of LP-based solution (cost component savings as a fraction of total savings) (%)				
				Opening	Inclusion	Production	Shipping	Penalty
PRAC	3	22.4	0.0	—	—	—	—	—
LP-based	3	11.3	2.0	8	13	84	7	−12

Notes. The last five columns show the cost component savings as a fraction of the total savings achieved by the LP-based solution. Penalty costs increase, so it has a negative contribution.

7. Application and Discussion

7.1. Model Implementation

The methods described in the paper were applied at the large multinational food processing company described in §1. The problems considered by the company were comparable to the large instance described in §6. The model was run using ILOG CPLEX to execute the first step of the LP-based heuristic and a specialized C++ program to execute the second step. These programs were integrated using Microsoft Visual Basic, which also provided an Excel-based interface to input the data and export the results. The model was maintained by the strategic operations department at this company and run by the senior managers, such as the director and assistant director at this department, who had been involved in all aspects of this project.

The LP-based heuristic was used to solve the model for several critical product families and processes chosen by the management. Based on this analysis, several products were reassigned to different processes, some processes were closed and consolidated to other plants. For example, before this project there were 14 refining processes in all of the 14 plants in Northern Europe. Based on the recommendations of this model, 5 smaller processes were moved to the larger plants, so that they currently operate with 9 processes across these 14 plants. Also, 15 products were reassigned across processes based on the results of the model to reduce switchover times and achieve greater economies of scale. Further product/process assignments to several markets were suitably changed. It is worth noting that the company ran the LP-based heuristic under different scenarios and then selected a solution that was modified to accommodate business conditions not captured by the model. For instance, the management decided to open 5 more processes to have some redundancy in case of disruptions or because it envisioned certain growth opportunities in the short term. The company estimates the overall annual cost savings in using our approach to be at least 10%, or \$50 million. These numbers were calculated using estimates of the total costs one year before and one year after the implementation of the recommendations of this model at each plant and then aggregating these costs across all plants.

7.2. Sensitivity Analysis: Managerial Insights on Process Choice

In addition to the implementation aforementioned, we used the LP-based heuristic to generate some managerial insights on process choice, which have directly influenced decision making of the top management and have led to significant strategic and organizational impact. To better understand the impact of *process variability on process choice*, we split the choice of processes under two categories: high variability processes with coefficient of variation (COV) of the yield distribution at 33%, and low variability process

with the COV at 3%. To isolate the impact of variability, we made sure the split between the high and low variability processes were not correlated to the capacity levels and, we changed only the standard deviations of the yield distributions. We found that when the per-unit capacity costs were the same across processes, then low variability processes were preferred to those with high COV. When the per unit capacity costs increased with process capacity, we found that processes with higher per-unit capacity cost and low variability were preferred over processes with lower per-unit cost but high COV. This happens because the reduction in penalty costs for unmet demand compensates for the increased per-unit capacity costs, and it illustrates the importance of variability reduction in these processes. Finally, if the goal is to reduce the total number of processes or fixed costs, then variability reduction should be targeted at the highest volume processes. This insight was particularly useful to the management of this company because this showed a direct link between decreasing variability and fixed cost reduction, a relationship that is useful in justifying further variability reduction schemes such as increased automation and more effective process control.

To further assess the impact of *capacity costs on process choice*, we used the same yield distribution across all the processes and considered the case when per-unit capacity costs decreased with process capacity. Our analysis showed that the processes with highest capacity and lowest per-unit costs were preferred. This observation was consistent with the management view of developing high-volume processes to the extent possible. However, in some plant locations, due to scarcity of existing space, adding another process would entail significant investments in real estate. This fact leads to per-unit capacity costs that increase with process capacity. Under this scenario, we would still pick the processes with lowest per-unit costs, but as expected, more processes must be opened to meet the same demand level.

We next wanted to analyze the impact of *product switchover times b_k on process choice*. We found that as switchover times increased as a percentage of average process capacity, more facilities are opened to compensate for these switchover times and avoid backorder penalty costs. Furthermore, this also led to consolidation of products to processes to the extent possible.

To understand the impact of *economies of scale on product choice*, we systematically increased the degree of concavity of the production curve for a particular product across all the processes, which led to lower marginal costs at higher volumes or higher economies of scale. As expected, this showed consolidation of products across processes. However, what was interesting is that the number of open processes itself did not change. Management was intrigued by this insight because they had always felt the production characteristics were a key driver to the number of processes. This analysis showed that one could be responsive to the market if necessary by producing products with lowered economies of scale without increasing

the fixed costs of product allocation, as long as one is careful about how these products are allocated to processes.

7.3. Qualitative Impact

Our sensitivity analysis on how process variability affects process choice had a significant impact on the policies of this company. Although the benefits of variability reduction for a single process were well known to the company, in part due to the influence of lean production and Six Sigma, its impact on process choice was less well understood. After our project, variability reduction initiatives have been prioritized to start with high volume processes in order to reduce total fixed investments at this organization. Furthermore, process improvement and product development teams now work closely in an integrative manner as the importance of reducing process variability concurrently with new product introduction became apparent after this research. In the medium term, management has also asked process engineering to reduce switchover times at key processes. They now believe this is critical to reducing backorder and distribution costs. Additionally, this allows this organization to be more responsive to customers, which is crucial in maintaining customer service levels, market share, and long-term profitability. Overall, management estimated that the return on investments resulting from these changes was on average 4%, varying from 3% to 5% across various product families. This improvement is significant given that products are commodities with market-defined prices and small margins, and consequently, making continuous improvements matters to remain competitive in the long run.

While management was intrigued by our analysis relating to the impact of capacity costs on process choice, they felt that it might not be feasible to achieve such decreases in per-unit capacity costs in the medium term because that would require substantial investments in new technology or increases in existing capacity. Their preference was to reduce process variability and do more with less capital investment in the process. Similarly, after discussing with R&D, they found that it might not be feasible to alter production characteristics of the product without substantial investments. Therefore, they decided not to alter the product design decisions at this time.

The organizational impact of this model has been significant. Initially, decisions on process location and product and market choice were based on subjective opinion, seniority, and anecdotal evidence. Furthermore, there was considerable skepticism about whether a structured and rigorous approach would improve on current practice, mostly because the managers at this organization were not familiar with OR-based methodologies. Furthermore, they felt that an analytical approach could not sufficiently capture all the aspects of what they considered to be an immensely complicated problem. A key step in convincing them about the efficacy of our method was to run the model in parallel to their manual approach. This process enabled them

to understand how and why our model outperformed their methods and to more broadly accept that OR as an effective discipline to deal with complexity. As a direct consequence of this project, the strategic operations department at this company is actively identifying complicated problems across various functions in this organization that can be modeled and solved using OR-based methodology.

Going forward, the management at this company has decided to use our model at least once in six months to compare the existing decisions on which process are opened and the allocation of products to processes and markets. The purpose is to compare the costs of the model and existing decisions and see how this could be reconciled. This procedure will also be repeated at all major change points that involve introducing a new product or acquisition of a new company. We expect this organization to sustain and improve the use of our framework several years into the future.

8. Conclusions

In this paper, we consider the process location and product distribution problem with uncertain yields typically found in several process industry sectors. In this problem, we choose the location of processes, the assignment of products to these processes, and the transportation of production quantities to markets in order to minimize total expected costs. This problem is very challenging due to characteristics that are unique to the process industry sectors such as economies of scale, large product switchover times, and yield uncertainty.

We model this problem as a nonlinear mixed-integer program. The complexity of this problem precludes its solution using even the most powerful commercial solvers. Therefore, we develop three heuristics to solve this problem. The first is a practitioners heuristic that opens processes with the least ratio of capacity to fixed costs and allocates products to process locations that minimize transportation costs. The second heuristic is a greedy procedure, and the third one is LP-based.

We develop a lower bound to assess the quality of the heuristic solutions. To circumvent the nonlinearity of the objective function, we derive a strong duality result that allows us to solve the nonlinear problem iteratively as a sequence of linear MIPs. Because this problem is still hard to solve due to the exponential number of yields scenarios, we propose a technique using the law of iterated expectations to aggregate the constraint set.

To test our methods, we used real data from the leading food processing company that motivated this project. Our computational study is encouraging. For the large instances we obtained solutions within 8.1% of the lower bound in the worst case and within 5.4% on average. The computational running times were within seven hours, which is reasonable given the complexity of the problem.

The LP-based heuristic has been used by this company to assess their product choice and product allocation

decisions. In particular, this has led to products being reassigned, processes being closed at some plants and consolidated at other plants. In addition, there has been change in the assignment of process/product combination to markets, which has led to annual savings of 10%, or \$50 million. In addition, this heuristic has generated several managerial insights. These include understanding the relationship between process choice and critical aspects within a process such as process variability due to yield uncertainty and switchover costs, and the relation between process choice, capacity costs, and product allocation decisions. Some of these insights have directly influenced the decision making of the top management and have consequently had significant strategic and organizational impact.

This paper opens up several opportunities for future research. First, further work could be done to improve the heuristics and lower bound to reduce the suboptimality gap. Second, although not relevant in our application context, this problem could be extended to include the location of warehouses between processing plants and markets. Third, another extension would include multiple time periods with inventory constraints at the different locations, which could be important in other settings. Undoubtedly, both of these extensions would require significant modifications to the computational methods to solve the general problem. Finally, our model could be applied in other industrial contexts. We are currently exploring this possibility in the pharmaceutical industry.

Electronic Companion

An electronic companion to this paper is available as part of the online version at <http://dx.doi.org/10.1287/opre.1120.1087>.

Endnotes

1. For confidentiality due to the strategic nature of our work, the company has not been disclosed. However, verification of this work has been provided by the company to the department editor.
2. There is a lower decision level that deals with the production sequence within a six-month time window; see Rajaram and Karmarkar (2002). However, this decision level is too granular for our model's purpose.
3. Most variable cost functions used in practice satisfy these conditions, e.g., Cx^a with $0 < a < 1$, which follows from a Cobb-Douglas production model with the sum of technology factors greater than one.
4. From this property we can assume that u_{jk}^t is bounded by $d_{jk} \forall j, k, t$.
5. To be precise, constraint (9) needs only to be dualized for those pairs (i, k) for which $p_{ik}(\cdot)$ is nonlinear.
6. Because $p(v) = \kappa(v)v$, it follows that $p(\kappa^{-1}(\mu)) = \mu \kappa^{-1}(\mu)$.

Acknowledgments

The authors are grateful to the large food processing company for introducing them to this problem, describing the practitioner's heuristics, providing the data used in the computational study, and sharing the implementation results. The authors thank Srinivas Bollapragada, Charles Corbett, Uday Karmarkar, Rakesh Sarin, an

associate editor, and two anonymous referees for their constructive comments that helped improve the paper. The authors also thank the UCLA Academic Senate Grant and the Easton Technology Leadership Program at UCLA Anderson School of Management for providing funding for this project. The third author thanks RWTH Aachen University, where a significant part of this work was done while he held the Deutsche Post Chair of Optimization of Distribution Networks.

References

- Aikens CH (1985) Facility location models for distribution planning. *Eur. J. Oper. Res.* 22(3):263–279.
- Arend M (2011) Site plan features active cultures. *Site Selection*. Accessed September 6, 2011, <http://www.siteselection.com/issues/2011/jul/food-processing.cfm>.
- Balachandran V, Jain S (1976) Optimal facility location under random demand with general cost structure. *Naval Res. Logist. Quart.* 23(3):421–436.
- Ball MO, Lin FL (1993) A reliability model applied to emergency service vehicle location. *Oper. Res.* 41(1):18–36.
- Barahona F, Chudak FA (2005) Near-optimal solutions to large-scale facility location problems. *Discrete Optim.* 2(1):35–50.
- Baron O, Berman O, Krass D (2008) Facility location with stochastic demand and constraints on waiting time. *Manufacturing Service Oper. Management* 10(3):484–505.
- Berman O, Krass D (2002) Facility location problems with stochastic demands and congestion. Drezner Z, Hamacher HW, eds. *Facility Location: Applications and Theory* (Springer, Berlin), 331–373.
- Brandreau ML, Chiu SS (1989) An overview of representative problems in location research. *Management Sci.* 35(6):645–674.
- Cui T, Ouyang Y, Shen Z-JM (2010) Reliable facility location design under the risk of disruptions. *Oper. Res.* 58(4):998–1011.
- Daskin MS (1983) A maximum expected covering location model: formulation, properties and heuristic solution. *Transportation Sci.* 17(1):48–70.
- Donovan D (2009) Food processing location trends. Industry report. Accessed September 6, 2011, <http://tradeandindustrydev.com/industry/food-processing/food-processing-location-trends-635>.
- Drezner T (2009) Location of retail facilities under conditions of uncertainty. *Ann. Oper. Res.* 167(1):107–120.
- Drezner Z, ed. (1995) *Facility Location: A Survey of Application and Methods* (Springer, New York).
- Euromonitor International (2011) Global market information database. Accessed September 6, 2011, <http://www.portal.euromonitor.com>.
- Fisher ML (1981) The Lagrangian relaxation method for solving integer programming problems. *Management Sci.* 26(1):1–17.
- Gehlhar M, Regmi A, eds. (2005) New directions in global food markets. Electronic report from the Economic Research Service, United States Department of Agriculture. Accessed September 6, 2011, <http://www.ers.usda.gov/publications/aib794/aib794.pdf>.
- Geoffrion A (1974) Lagrangian relaxation for integer programming. *Math. Programming Stud.* 2(1):82–114.
- Geoffrion AM, Graves GW (1974) Multicommodity distribution system design by Benders decomposition. *Management Sci.* 20(5):822–844.
- Gilmore PC, Gomory RE (1961) A linear programming approach to the cutting-stock problem. *Oper. Res.* 9(6):849–859.
- Karmarkar US, Rajaram K (2001) Grade selection and blending to optimize cost and quality. *Oper. Res.* 49(2):271–280.
- Laporte G, Louveaux FV, van Hamme L (1994) Exact solution to a location problem with stochastic demands. *Transportation Sci.* 28(2):95–103.
- Louveaux FV (1986) Discrete stochastic location models. *Ann. Oper. Res.* 6(2):23–34.
- Louveaux FV, Peeters D (1992) A dual-based procedure for stochastic facility location. *Oper. Res.* 40(3):564–573.
- Mehrotra M, Dawande M, Gavirneni S, Demirci M, Tayur S (2011) Production planning with patterns: A problem from processed food manufacturing. *Oper. Res.* 59(2):267–282.

- Melo MT, Nickel S, Saldanha-da-Gama F (2009) Facility location and supply chain management—A review. *Eur. J. Oper. Res.* 196(2): 401–412.
- Owen SH, Daskin MS (1998) Strategic facility location: A review. *Eur. J. Oper. Res.* 111(3):423–447.
- Qi L, Shen Z-JM (2007) A supply chain design model with unreliable supply. *Naval Res. Logist.* 54(8):829–844.
- Rajaram K, Karmarkar US (2002) Product cycling with uncertain yields: Analysis and application to the process industry. *Oper. Res.* 50(4):680–691.
- Rajaram K, Jaikumar R, Behlau F, van Esch F, Heynen C, Kaiser R, Kuttner A, van de Wege I (1999) Robust process control at Cerestar's refineries. *Interfaces* 29(1):30–48.
- Scanlon C (2006) Food processing prospers. Industry report. Accessed September 6, 2011, <http://www.areadevelopment.com/industryReport/aug06/foodProcessing.shtml>.
- Shen Z-JM, Coullard C, Daskin MS (2003) A joint location-inventory model. *Transportation Sci.* 37(1):40–55.
- Shu J, Teo C-P, Shen Z-JM (2005) Stochastic transportation-inventory network design problem. *Oper. Res.* 53(1):48–60.
- Snyder LV (2006) Facility location under uncertainty: A review. *IEEE Trans.* 38(7):537–554.
- Snyder LV, Daskin MS (2005) Reliability models for facility location: The expected failure cost case. *Transportation Sci.* 39(3):400–416.
- US Department of Commerce (2006) Industry outlook, Food manufacturing NAICS 311. Accessed September 6, 2011, http://ita.doc.gov/td/ocg/outlook10_food.pdf.
- Wollenweber J (2008) A multi-stage facility location problem with staircase costs and splitting of commodities: Model, heuristic approach and application. *OR Spectrum* 30(4):655–673.

Felipe Caro is an assistant professor of decisions, operations, and technology management at the UCLA Anderson School of Management. His research interests span retail operations, supply chain management, carbon footprint, and natural resources, with a strong emphasis on practical applications.

Kumar Rajaram is a professor of operations and technology management at the UCLA Anderson School of Management. His current research interests include improving operations in the process manufacturing sectors including food processing, pharmaceuticals, and the petrochemical industry. He has focused on developing analytical models of complicated systems, and his methods have been implemented at several large-scale industrial processes worldwide.

Jens Wollenweber is head of the freight audit department at Dr. Staedtler Transport Consulting, Nuremberg, Germany. He holds a doctoral degree in operations research from RWTH University, Aachen, and a diploma in information systems from Cologne University. His research interests are models and heuristics for facility location, routing, and scheduling.