Case Analysis: Analyses of the National Cranberry Cooperative—1. Tactical Options

EVAN L. PORTEUS

Graduate School of Business
Stanford University
Stanford, California 94305-5015

We initiated this section in the November-December 1989 issue of Interfaces with the publication of the National Cranberry Case and an invitation to submit analyses by May 1, 1990. We promised to "review the reports, select and edit the best sections from them, prepare a commentary, and publish the results." In this first part of those results, I discuss results that consider the environment fixed and focus on the tactical options available. I include edited analyses of the submissions from Khashoggi, Chabas, and Bakken [1990], Juergens, Hoyt, and Swenson [1990], and Langelo [1990]. In the second part, I will expand the analysis to consider changes in the environment and implementation issues, and I will include edited analyses of three additional submissions.

In February 1971, considered the present for purposes of this analysis, Hugo Schaeffer, vice-president of operations at the National Cranberry Cooperative (NCC), faces two primary problems that arise at receiving plant No. 1 (RP1) during the annual autumn cranberry harvesting season: (1) trucks and drivers spend too much time waiting to unload process fruit at the receiving plant, and (2) overtime costs and absenteeism are out of control. There is also a secondary problem: half of the berries graded as top quality and awarded a 50-cent premium per barrel (bbl) are not top quality and do not deserve the premium.

The Management Science/Operations Research (MS/OR) approach to problem solving is to seek more than one cause, to seek alternative ways to alleviate the problems, and to seek numerical measures to evaluate the alternatives. It also emphasizes looking at the big picture. For example, one possible cause of the buildup, especially as the truck drivers and those in the receiving area see it, is the limited number of Kiwanee dumpers. So last year, the cooperative purchased a fifth Kiwanee

0091-2102/93/2304/0021$01 25
This paper was refereed

INDUSTRIES—AGRICULTURE/FOOD
QUEUES—APPLICATIONS

dumper for $75,000. But the analyses in this article reveal that, unless there were other reasons for buying the additional dumper, the money could have been spent more wisely on other things.

MS/OR problem solving 20 to 40 years ago emphasized systems: Be sure you understand how the whole system works and how each little piece affects everything else before you run out and buy another dumper. However, we have often made the different mistake of trying to build one best model of the system; anything that might be relevant is included in the model. Such an approach leads to large complicated models that even the analysts may not fully understand. Successful managers do not base their decisions on analyses they do not understand [Little 1970; Woolsey 1978]. Our challenge, therefore, is to become masters of model building and analysis. We must build the simplest models that encompass the important parts of the big picture. There are no set rules. We must practice. We may even need a little something that can't be taught. But we must keep it simple, so that we obtain usable insights and conclusions. In this article and its sequel, I try to illustrate the art of model building and analysis. In particular, I emphasize examining a variety of models, taking various perspectives into account, and drawing different insights, rather than presenting complete analyses or illustrating an efficient way to reach managerial conclusions pertinent to the case at hand. I also try to describe the serendipity that frequently accompanies a good analysis. Understanding how something works and what is critical often leads to new, creative ideas. I present a number of probably unworkable options to emphasize this important and underappreciated aspect of good analysis.

Analysis 1: Systemwide Optimization

NCC is a cooperative, which means that any surplus (or loss) gets returned to (or absorbed by) the growers themselves. Thus, any analysis should account for the consequences to the growers as well as those to RP1. For example, truck waiting costs and grower net revenues should be taken into account whenever they are affected. If the managers of the cooperative are reluctant to think in this way, then the growers should insist on it. For example, if the managers are evaluated only on the performance of the cooperative and not on grower returns, then their performance evaluation system should be revised. This idea is not reserved for cooperatives. There is no law against an organization cooperating with its suppliers and customers to find ways of operating that yield the largest total gain and sharing that gain so all parties are better off.

The chief berry receiver grades truckloads of berries on arrival at RP1, using color pictures as a guide, into grades 1 (poorest), 2A, 2B, or 3 (best). When in doubt between grades 2B or 3, he usually selects grade 3. The 50-cent premium for grade 3 berries was paid on about 450,000 barrels of berries in 1970, yet only about half of them deserved that premium, amounting to over $100,000 in undeserved premium payments. Schaeffer is thinking of installing a light meter system costing $10,000 and requiring another skilled operator to eliminate this problem. A first analysis may indicate that such a system should be installed, as so much can be saved. However, accounting for the consequences to the growers can lead to a different conclusion. If the growers share in the surplus of RP1 in proportion to the total volume they deliver and they each receive the undeserved premium on the same percentage of berries they deliver, then the current system merely pays out cooperative returns to the growers early. If all the berries were graded correctly, the cooperative would show a surplus that is larger by the amount of the undeserved premium payments, which would then be distributed back in exactly the same amounts to the growers. The growers would lose the use of those funds until they received them later. Adding the light meter system and an additional skilled operator would reduce the surplus and the nets to the growers. The problem is exacerbated if the growers are taxed at a higher effective marginal tax rate on funds they receive out of RP1 surplus than on funds they receive directly.

Under the current system, some other buyer of berries who grades berries correctly may be able to offer a higher premium than RP1 for top-quality berries. Growers able to identify their top-quality berries may divert them to this other buyer to obtain the higher premium. RP1 would then receive fewer top-quality berries. It is unclear whether such growers would benefit from the diversion, as fewer of their undeserving berries would receive the premium from RP1. Possibly the remaining growers would benefit: Reducing the volume of berries flowing through the plant would reduce congestion at RP1 and might improve performance. If RP1 anticipated a major reduction in volume, it might be able to cancel or postpone its planned expansion.

Currently, growers may not receive the undeserved premiums equitably. For instance, growers who use wet harvesting (flooding the bogs, mechanically shaking the bushes, and collecting the floating berries) may get disproportionately more of the undeserved premiums than growers who use dry harvesting (the traditional method of handpicking the bushes), because wet harvesting is apt to be less selective in getting only the top-quality berries.

Any growers who do not realize that the chief berry receiver is grading the berries too generously might be inappropriately encouraged to harvest berries either before or after the time they are at top ripeness quality, to avoid peak congestion periods at RP1. They could believe that their berries would continue to be graded as No. 3 berries yet fewer grade 3 berries would be delivered, perhaps leading to lower total returns.

I recommend that each truck's berries be graded according to the percentage that are of each color grade, rather than rating the whole truckload as one color grade or another. If that is infeasible, it may be worth introducing another berry category, between 2B and 3, with a smaller premium. In either case, RP1 would need new color charts. If both plans are infeasible, the chief berry receiver could be instructed to be less generous in grading berries in 1971.

There is no law against an organization cooperating with its suppliers and customers.

Preseason training and creating new, perhaps larger, color charts should be far cheaper than the proposed light meter system and nearly as accurate. Granting graders incentive pay for accurate grading is worth considering in conjunction with any of the above changes.

Analysis 2.1: Base Case Forecasts of Peak Season Delivery Volumes

I select a day as the time unit of analysis, because it is the shortest plausible cycle. Each day should begin with no trucks waiting to be unloaded and no berries in the temporary holding bins, and so can be analyzed somewhat independently. Significant dependencies exist among the different time periods within a day. I determine how long it takes to finish processing all the berries that arrive on a day and estimate the resulting overtime and truck waiting costs.

My approach is to break out the daily volumes over the 20-day peak season into a few categories and analyze the plant on each such day. If volumes were to resemble those of last year (1970), I would assume 18,000 barrels will arrive on nine days, 16,000 will arrive on six days, 14,000 on four, and 12,000 on one day. However, evidence in the case suggests that daily volumes will be higher than that.

Over the years, US cranberry harvest productivity has increased (case Table 1). The yield per acre has increased by about 44 percent over the past three years, which corresponds to a compound growth rate of nearly 13 percent per year. Some of the recent gain is attributed to water harvesting, with the statement in the case that "water harvesting could result in yields up to 20 percent greater than those obtained via dry harvesting. . . ." So one simple approach is to assume that process fruit volume will increase by another 13 percent in 1971.

A more precise approach is to examine the volume of process fruit delivered to RP1 over the last few years. Across the US, process fruit has recently represented about 80 percent of the total of process fruit and fresh fruit sales (case Table 1). Applying that percentage to the total volumes at RP1 for 1967 through 1969 (case Figure 1), I estimate the process fruit volumes at RP1 to be 387,000, 426,700, and 542,500 barrels in 1967 through 1969, respectively. RP1 received 610,000 barrels of process fruit in 1970 (case Exhibit 2). Because the case indicates that 10 percent of the crop was set aside in 1970, I estimate the actual process berry crop at 610,000/.9 ≅ 678,000 barrels, which is about a 25 percent increase over 1969. The compound growth rate from 1967 through 1970 is therefore about 20.5 percent. Some, but clearly not all, of that increase is due to an increase in water harvesting. Thus, the RP1 area seems to be achieving better than average productivity gains. I shall therefore assume that the process fruit volume will increase by another 20 percent in 1971.

An even more precise approach would be to try to break out the productivity gains in the RP1 area to gains from water harvesting and from other changes. Schaeffer asserts in the case that the percentage of water harvested berries will increase to 70 percent in 1971 from 58 percent in 1970, so this more precise approach is apt to indicate a percentage increase around 20 percent as well.

To estimate the distribution of daily volume of process fruit for the 1971 peak season, I add 20 percent to each of the volumes observed in 1970 and group the volumes in increments of 2,000 barrels per day. I conclude that the 1971 peak season will see seven 22,000 days, six 20,000 days, four 18,000 days, two 16,000 days, and one 14,000 day. Equivalently, I estimate the probability that the volume on a peak season day will be 22,000 is 7/20, and so forth.

Peak day volumes in 1970 may have been limited by the processing capacity at RP1: if RP1 had enough capacity to prevent trucks from waiting long, then they could return to the fields more quickly and deliver more berries on peak days. Thus, if RP1 increases its capacity for 1971, it may see peak days that are more than 20 percent higher than the peaks of 1970. On the other hand, the RP1 area may not continue to see productivity gains in the 20 percent range and peak days may be less than that.

Analysis 2.2: Base Case Model

Figure 1 shows the process flow for RP1, including the processing capacity of each processing point and the storage capacity of each storage point. This diagram is due to Jerry Miller and is nearly equivalent to one given by Hayes and Wheelwright [p. 175, 1984]. The diagrams differ on where the output of the bulk bin and bulk truck stations can go. The case is not clear on this point. If the difference is important, it can be resolved on site. Capacities either come directly from the case or are estimated using simple assumptions. For example, the case says it usually takes a Kiwanee dumper from 5 to 10 minutes to service a truck. Using an average of 7.5 minutes to unload a truck carrying an average of 75 bbls, I get a capacity of 600 bbl/hr.

I assume that berries flow continuously and at a constant rate throughout the process: Berries (on trucks) start arriving at 7:00 AM and continue to arrive at a constant rate throughout the day, until 7:00 PM. Although trucks stop arriving during the lunch hour (case Exhibit 1), and it is not difficult to refine the model to include this characteristic, I do not, to keep the model as simple as possible. The actual (constant) processing rate at a work station is determined by the staffing level at that station and the maximum rate at which berries can be processed through other work stations at that time. Buildup or draw down at a storage node also occurs continuously, at a constant rate. If all work stations are active, I assume the time it takes for a berry to flow from receiving through shipping is on the order of a few minutes and can be ignored in the analysis.

Consider a day in which 18,000 barrels arrive, 70 percent of which are wet and the rest dry. That is 1,050 bbl/hr of wet berries and 450 bbl/hr of dry will flow into the plant during its 12 hours of receiving berries.

I decompose the process into two independent continuous flow processes, one for wet berries and one for dry. To do so, I assume that work-station capacities can be allocated in any fraction desired to either berry type. This assumption is justified if the changeover times for work stations that can process both types are insignificant.

I make a preliminary allocation of the shared capacity at each work station and storage point, attempting to allocate 70

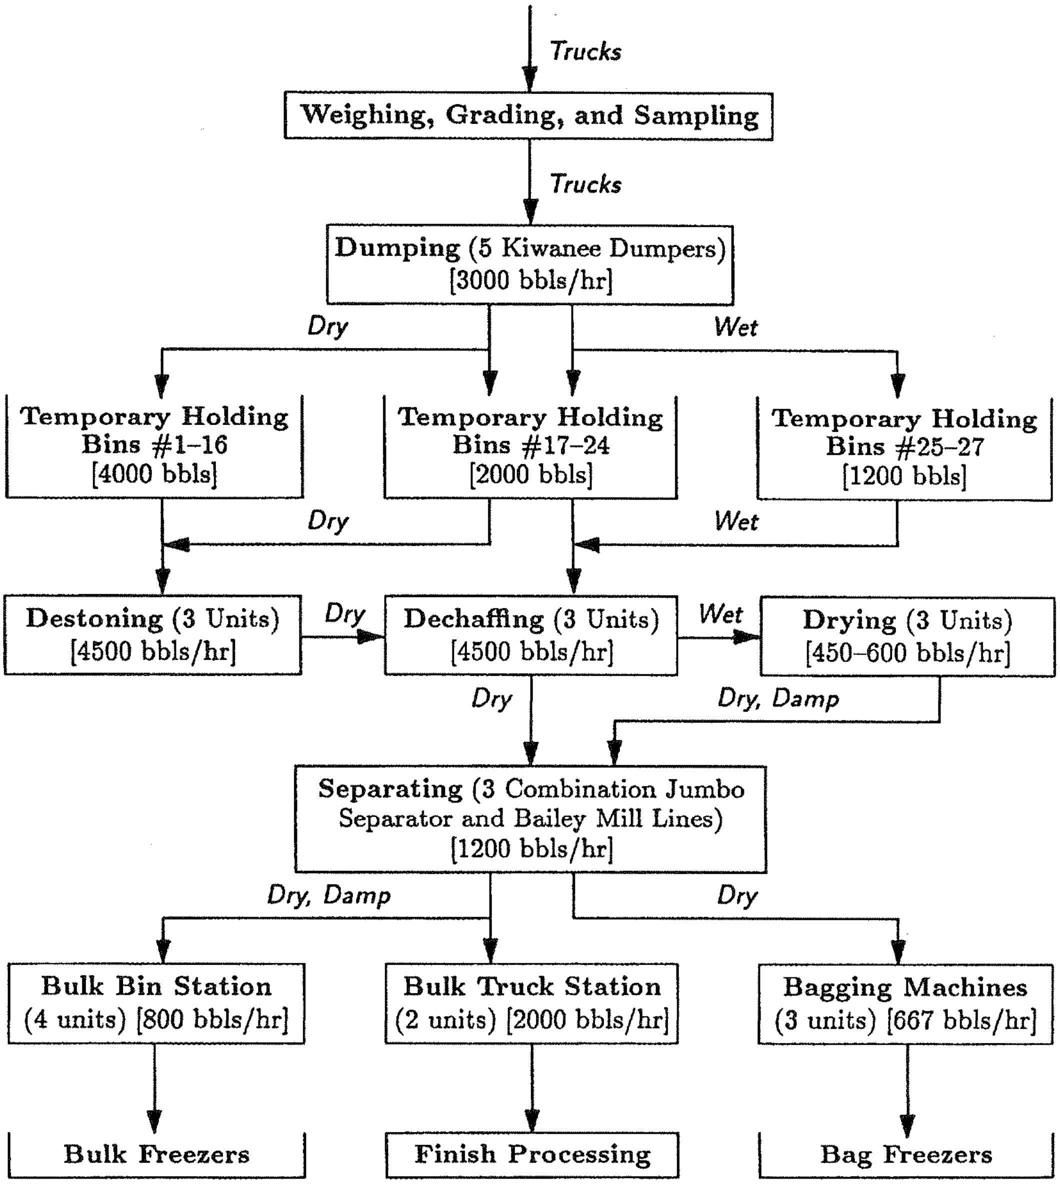

Figure 1: This process flow diagram shows the flow of berries through RP1, including truck arrivals, segregation into wet and dry berries, and the option to partially dry wet berries, creating damp berries. Capacities of work stations, shown as rectangles, are given in barrels per hour, and capacities of storage points, shown as uncovered rectangles, are given in barrels.

percent of the capacity there to wet berries and 30 percent to dry. For example, I tentatively allocate 840 bbl/hr of separator capacity to wet berries, and 360 to dry. At most, 3,200 barrels (less than 70 percent) of the holding bin capacity can be allocated to wet berries, so I do that. The work station with the least capacity for a product is its bottleneck. The bottleneck for dry berries is separating (360 bbl/hr), while the bottleneck for wet is drying, at 450–600 bbl/hr, depending on whether the berries are to be bagged or not. (Bagged berries need to be drier than those shipped in bulk, which can be damp.) To relieve this bottleneck, as suggested by Juergens, Hoyt, and Swenson [1990], I assume that wet berries are only partially dried during the peak season and are shipped only by bulk. Thus, wet berries can flow through the plant at 600 bbl/hr. Since wet berries have more separating capacity than they can use, I reallocate the excess 240 bbl/hr to dry berries (Table 1).

Under the high volume schedule (case Figure 5), all five Kiwanee dumpers and all 27 holding bins are available starting at 7:00 AM when the berries start arriving, but destoning, dechaffing, drying, separating, and shipping do not start until 11:00 AM when their crews arrive. Thus, from 7:00 to 11:00 AM, 4,200 bbl of wet berries and 1,800 bbl of dry arrive. The holding bins cannot hold all these wet berries, so the excess 1,000 barrels must wait on arriving trucks, creating anger and frustration. At 11:00 AM, RP1 begins processing the dry berries at 600 bbl/hr, which is faster than the 450 rate at which they are arriving. The 1,800 barrels in bins decline at the rate of 150 bbl/hr, diminishing to 600 at 7:00 PM when arrivals stop, and they are cleaned out in one additional hour, at 8:00 PM (Figure 2).

Wet berries are a nightmare. RP1 begins processing them at 11:00 AM, also at 600 bbl/hr, which is far less than the 1,050 rate at which they are arriving. The truck queue keeps growing until 7:00 PM, when 7,800 barrels of wet berries are in the system, 3,200 in bins and the rest in trucks (Figure 2). The trucks continue to unload until 2:40 AM and processing continues until 8:00 AM. The workers are scheduled to return to work again in three hours, the required two hours per day for cleaning and maintenance may be difficult to schedule, and this isn't even the peak day!

If during the peak season, RP1 scheduled the work force to arrive at 7:00 AM so that processing began when the berries started arriving, there would be no buildup of dry berries and wet berries would build up at a rate of 450 bbl/hr over the 12 hour day, reaching a total of 5,400 barrels at 7:00 PM, 2,200 of which would be on trucks. The trucks would be emptied by 10:40 PM and the bins by 4:00 AM (Figure

Description	Work Station (W) or Storage (S)?	Capacity Wet Only	Capacity Dry Only
Kiwanee Dumpers	W	2,100	900
Holding Bins	S	3,200	4,000
Destoning	W	NA	4,500
Dechaffing	W	3,150	1,350
Drying	W	600	NA
Separating	W	600	600
Shipping	W	2,427	1,040

Table 1: Roughly 70 percent of processing and holding capacity is allocated to wet berries. The maximum possible holding bin capacity is given to wet. The bottlenecks are drying for wet berries, at 600 bbl/hr, and separating for dry, also at 600 bbl/hr.

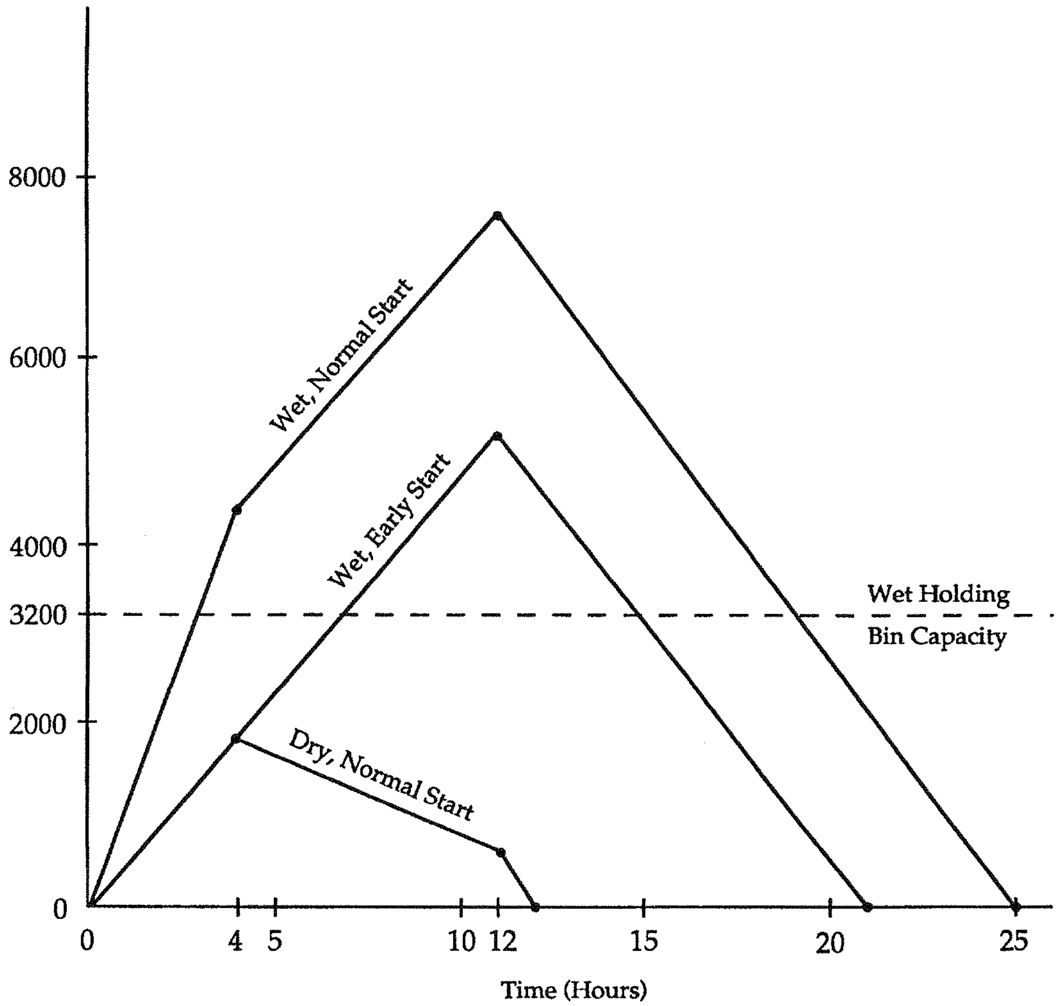

Figure 2: The buildup of wet and dry berries over an 18,000 barrel day at RP1 depends on when processing operations begin, either a normal start at hour 4 (11:00 AM) or an early start at hour 0 (7:00 AM). Wet berry buildup above 3,200 barrels is on waiting trucks. The area between the wet berry curve and the 3,200 line gives the truck-hours of waiting on such a day.

2). This "early start" would reduce the average truck wait from 3.4 hours to 45 minutes. The number of truck-hours of waiting during the day is given by the area below the buildup curve and above the holding bin capacity. For example, for wet berries and an early start, there are $(5400 - 3200) \times (15.67 - 7.11)/(2[75]) \cong 125.5$ truck-hours of waiting (at 75 barrels per truck), which amounts to about 45 minutes per truck, since $0.7(18{,}000)/75 = 168$ trucks arrive with wet berries over the day. I assume an early start in the following analyses, unless indicated otherwise.

Analysis 2.3: Base Case Recommendations

I develop a simple cost model to evaluate various options. I assume that the primary consequences of any options considered will occur during the 20-day peak season and that the primary costs are for trucks waiting and labor. In the absence of data, I arbitrarily assume that each truck-hour of waiting costs $10. If the results of the analysis are sensitive to this assumption, NCC can use a more accurate estimate.

I assume 15 people work in receiving from 7:00 AM to 7:00 PM and 37 others work in other areas (a control room operator, a destone/dechaff/dry operator, 15 people in separating, and 20 in shipping) as long as the plant is operating (case Figure 5). The case indicates that staffing can be reduced for operations after 11:00 PM, but I leave such fine tuning for a further analysis. My labor cost measure is total labor hours over the peak season multiplied by $3.375, the overtime rate for temporary workers. Although using this rate does not yield the correct total for labor, it is useful for estimating the changes in that total due to various options. Thus, my simple single day cost model consists of $10 for each truck-hour of waiting, $50.625 for each hour that the receiving crew is working, and $124.875 for each hour that the rest of the plant is operating.

Clearly, the status quo is unacceptable, and something needs to be done. The bottleneck is processing wet berries through the dryers. An obvious option is to add one or more dryers, as suggested in the case. I use the model of Analysis 2.2 to predict the inputs needed for the simple cost model.

If one dryer is added, then the dryer capacity increases to 800 bbl/hr, and, to utilize it, separator capacity must be reallocated to wet berries, leaving a capacity of 400 bbl/hr for dry berries. Thus, on an 18,000 day, wet berries build up at a rate of 250 bbl/hr over the day, ending at 3,000 barrels at 7:00 PM. They are cleared out at 10:45 PM. The dry berries build up at a rate of 50 bbl/hr to 600 at 7:00 PM. They are cleared out by 8:30 PM. Thus, no truck waiting is predicted. Receiving operates for 12 hours and the plant operates for 15.75 hours. Total peak season costs amount to $69,349 (Table 2).

If two dryers are added, the nominal dryer capacity increases to 1,000 bbl/hr. However, a new problem arises. For the dryers to operate at 200 bbl/hr capacity, their output must be shipped by bulk. While the direct bulk stations can handle 2,000 bbl/hr, their output can go only to the local process plant, which handles only 700 bbl/day. The bulk bin stations, when

Daily Volume (000s):	22	20	18	16	14	Total
Receiving Hours	15.25	13.5	12	12	12	
Plant Operating Hours	19.25	17.5	15.75	14	12 25	
Truck-Hours Waiting	149.57	38.18	0	0	0	
Daily Costs	4,672	3,250	2,574	2,356	2,137	
Days in Season	7	6	4	2	1	20
Labor Cost	22,231	17,212	10,297	4,712	2,137	56,589
Waiting Cost	10,469	2,291	0	0	0	12,760
Cost Measure	32,701	19,503	10,297	4,711	2,137	69,349

Table 2: The costs of waiting and labor over the 20-day peak season with one new dryer amount to $69,350. Most of the costs are incurred during the peak days.

fully staffed, can process only 800 bbl/hr. I therefore assume that the output of a second additional dryer would be directed to the bagging stations, and therefore, this dryer would be able to operate at only 150 bbl/hr capacity. The drying capacity would therefore be 950 bbl/hr, so 250 bbl/hr of the separator capacity would be allocated to dry berries. The output of the fifth dryer would be added to the flow of dry berries through one of the separator lines, destined for bagging. The first four dryers' output would be processed on the two other separator lines, in a flow to the bulk stations. The bulk bin stations would operate at full capacity, 800 bbl/hr, and the direct bulk stations would handle 700 barrels per day and act as excess capacity to cover for variability in the processing rate of the bulk bin stations. If, at any time during a day, the flow of wet berries through the plant is not sufficient to fully utilize the separators, then additional dry berries could be processed and added to the flow of wet berries headed for bulk shipping.

On an 18,000 day, the wet berries build up at a rate of 100 bbl/hr to 1,200 at 7:00 PM and they are cleared out at 8:16 PM, when all separator capacity can be reallocated to dry berries. Dry berries build up at 200 bbl/hr, to 2,400 at 7:00 PM. The level drops to 2,084 at 8:16 PM. They are cleared out at 10:00 PM. The only benefit to plant operations of the second added dryer on an 18,000 day is to reduce clearout time to 10:00 PM from 10:45 PM, a total of 45 minutes. Essentially, the bottleneck determining closing time of the plant has shifted to the separators, which can process only 1,200 bbl/hr. However, adding the second dryer reduces truck waiting substantially and receiving labor costs moderately on higher volume days (Table 3). Total peak season costs amount to $54,489. Thus, the second dryer, costing $25,000, saves nearly $15,000 in its first year. It appears to be a worthwhile investment.

Another option is to convert one or more of the dry holding bins to hold wet berries as well, at $5,000 per conversion. This option would reduce waiting times for trucks with wet berries and receiving crew hours. Such conversions would not affect the time the plant must operate over a day. My previous analysis indicates that it is worth spending $25,000 to purchase a second dryer. However, I must determine if that $25,000 is better spent on converting five bins. Comparing the effects of the two op-

Daily Volume (000s):	22	20	18	16	14	Total
Receiving Hours	12.84	12	12	12	12	
Plant Operating Hours	18.33	16.67	15	12	12	
Truck-Hours Waiting	17.29	0	0	0	0	
Daily Costs	3,112	2,689	2,481	2,272	2,106	
Days in Season	7	6	4	2	1	20
Labor Cost	20,572	16,134	9,923	4,544	2,106	53,279
Waiting Cost	1,210	0	0	0	0	1,210
Cost Measure	21,782	16,134	9,923	4,544	2,106	54,489

Table 3: The costs of waiting and labor drop to $54,491 when a second new dryer is purchased. Waiting costs almost disappear.

tions shows that it is not (Table 4). I must also determine whether it is worth converting any bins in addition to adding two dryers. I estimate that converting one bin would save $731 over the season, and converting a second bin would save an additional $431. Converting one or two bins might be worth doing, if I consider the possible benefits in subsequent years, but the decision is not clear cut. For example, the Kiwanee dumpers may have a moderate amount of storage capacity, which might exceed that of a single bin in aggregate, reducing the benefits of a conversion. Furthermore, the Cranberry Marketing Order of 1968 bases grower's post-1973 allotments on their yields through 1973, giving them an incentive to temporarily increase their yields, even if these increases would not be economically justified in the post-1973 environment. The volume of the harvests may actually decrease starting in 1974. Investments made this year to cope with high volume harvests probably ought to have a three-year payback or less. I recommend deferring a positive decision on converting bins for another year to get a better handle on how actual operations unfold.

I therefore propose the base case, against which other options will be compared, as consisting of adding two dryers, converting no bins, and an early start.

Analysis 3.1: Accounting for Variability

It is well known that variability can play a significant role in an analysis. Indeed, the base case accounts for inter-day variability in arrivals. However, it assumes exactly 70 percent of all berries will be wet on each peak day. Yet case Exhibit 2 indicates that, during the 1970 peak season, only 48 percent of the deliveries were wet, even though the average over the entire season was 58 percent wet. Furthermore, the percentage wet during the first few heavy days of the peak season was even less than 48 percent. While 70 percent of the volume this year may be wet, the percentage of berries delivered wet during the critical peak season may be less than 70 percent, and much less during the first few days of that season. If wet harvesting tends to occur later than dry harvesting, the trend toward primarily dry berries at the beginning of the peak season and primarily wet at the end may be more pronounced. Because most of the highest volumes occur early in the peak season, recognizing this effect in the analysis may reduce the need for new dryer capacity.

On the other hand, accounting for variabilities in intra-day truck arrivals, percentage of berries arriving wet, and processing times is likely to add support to the need

Daily Volume (000s):	22	20	18	16	14	Total
Receiving Hours	13.69	12	12	12	12	
Plant Operating Hours	19.25	17.5	15.75	14	12.25	
Truck-Hours Waiting	40.33	0	0	0	0	
Daily Costs	3,500	2,793	2,574	2,356	2,137	
Cost Measure	24,501	16,756	10,297	4,711	2,137	58,402

Table 4: The costs of waiting and labor drop only to $58,404 when five bins are converted instead of adding a second new dryer.

for additional drying capacity. Juergens, Hoyt, and Swenson [1990] assess the effect of such variabilities by building a Lotus 1-2-3 spreadsheet simulation model in which they break up the day into discrete periods one hour long, starting at 6:30 AM. They use case Exhibit 1 to determine the volume of arrivals during each such hour on September 23, 1970. They assume that the arriving volume in each hourly period of a 1971 peak season day is a random variable that is uniformly distributed over the interval from 80 percent of the volume that hour on 9/23/70 to 120 percent of that volume. Thus, heavy arrivals occur from 9:30 to 10:30 AM and light arrivals from 11:30 AM to 1:30 PM. They also assume that the percentage of the arriving berries that are wet during an hourly period is uniformly distributed, going from 60 percent to 80 percent wet. They have 26 random variables for each day: one each for volume and percentage wet for each of the 13 delivery hours of the day. They assume all these random variables are statistically independent. Thus, on any particular 1971 peak season day, the arriving volume can be quite different from the 18,340 that arrived on September 23, 1970, and the percentage wet can also vary from 70 percent. It is worth noting, though, that the expected volume on such a day will be 18,340 barrels, which is less than the 19,700 average suggested by Analysis 2.1. They continue to assume that processing rates are deterministic; for example, 600 bbl/hr of wet berries can be processed each hour with the existing three dryers.

Khashoggi, Chabas, and Bakken [1990] recommend using 12-minute time periods. Excluding lunch time, they use case Exhibit 1 to determine the empirical probability distribution, rounded to multiples of 10 barrels, of the arrival volume over the various remaining active arrival time periods on September 23, 1970. They assume that these are the only periods in a day when arrivals can take place and that this distribution is representative of all such periods during each day of the 1971 peak season. Assuming statistical independence, they sample from this distribution to determine a simulated volume of arriving berries in each period. They assume that the entire volume of arrivals during a time period is either wet, with probability 0.70, or dry.

With either of the above models, the base case and other pertinent proposals can be evaluated by simulating an appropriate number of days. No completed analyses using these models are offered here. However, performance of the system will be degraded, further supporting the need for additional dryer capacity. Such an analysis may conclude that buying a third new dryer is justified. However, the bottleneck already shifted to the separators with purchase of the second new dryer, so the benefit of a third dryer will be modest. The major effect of these variabilities appears to be on staffing issues, which are considered in detail in subsequent sections.

Analysis 3.2: Receiving Crew Staffing

On a peak day, when 22,000 barrels arrive at RP1, assuming no arrivals occur during lunch, berries will arrive at a rate of 2,000 bbl/hr. Since five dumpers can process berries at an average rate of 3,000 bbl/hr, the dumpers may have excessive capacity. With the deterministic, continuous flow model of Analysis 2.2, RP1 would need only four dumpers and four crews, with a capacity of 2,400 bbl/hr. However, because the variability in both the truck interarrival times and dumper processing times put additional demands on dumper capacity, I conduct a queuing analysis to assess the need to staff the fifth dumper.

I conduct a steady state analysis for each of the five different representative days during the peak season, using an M/M/K queue with a mean service time of 7.5 minutes, and setting K equal to the number of crews. For example, for a 22,000 day, I use a mean interarrival time of 2.25 minutes (based on 75 barrels per truck, yielding 26.67 trucks per hour). Continuing to assume truck waiting costs $10 per hour, I find that truck waiting costs of nearly $15 per hour can be saved by staffing the fifth Kiwanee dumper (Table 5). Staffing that dumper for an hour requires three workers, say at $2.25 each, for a total of $6.75, which is less than the monetary benefit.

However, the variability in arrivals and in dumper times may not be as much as assumed in the analysis. If not, the savings from staffing the fifth dumper will be smaller and possibly less than the $6.75 per hour needed to justify it. I therefore attempt to estimate better the variance of the interarrival times and dumper times. Excluding the first arrival of the day, the first arrival after lunch, and all arrivals recorded at 7:00 PM (which are likely to have occurred later), the mean and standard deviation of the interarrival times on September 23, 1970 are 3.8 minutes and 3.35 minutes, respectively, so the coefficient of variation (standard deviation divided by the mean) is $c_A \cong 0.88$, compared to 1.0 for the case of Poisson arrivals. If I assume that the dumper time is uniformly distributed over five to 10 minutes, then the coefficient of variation of service times is $c_S \cong 0.19$. Thus, actual variability appears to be less than assumed by the M/M/K analysis, and I seek a more refined analysis to evaluate the benefits of staffing the fifth dumper.

Whitt [1983a, b] (see also Chen et al. [1988]) suggests an approximation formula for the expected delay (prior to service) in a GI/G/1 queue that amounts to taking the M/M/1 result and multiplying by $(c_A^2 + c_S^2)/2$, which equals 0.409 in this case. If I apply the same idea in this setting, we estimate that the savings per hour for staffing the fifth dumper is 0.409(14.73) $\cong$ $6.03, which suggests that staffing the fifth dumper is not justified. However, Whitt suggests the correction for single server queues, not for the multi-server queues I have, and I am interested in transient results, not steady state results. (Every day starts with an empty system.) To get a sense of the predictive validity of this correction, I conduct a transient analysis of

Daily Volume (000s):	22	20	18	16	14	Average
Trucks/Hour	26.67	24 24	21.82	19.39	16.97	23.76
Probability of Such a Day	0.35	0.30	0.20	0.10	0.05	
Average Wait/Truck (min), 4 Dumpers	7.40	4.04	2.36	1.40	0.82	4.46
Average Wait/Truck (min), 5 Dumpers	1.47	0.93	0.58	0.34	0.19	0.95
$ Savings/hour from fifth Dumper	26.37	12.56	6.50	3.43	1.76	14.73

Table 5: The savings in truck waiting costs from staffing the fifth Kiwanee dumper average $14.73 per hour over the peak season, assuming Poisson arrivals and exponential service times.

the waiting time of all arriving trucks on a peak 22,000 day. I assume trucks arrive at a mean rate of 26.67 per hour and stop the simulation after nine hours of a day: The deterministic approach of Analysis 2.2 indicates that trucks will begin to queue about then, because of downstream bottlenecks in the system, rather than lack of available dumpers. For each simulated day, I assume the interarrival times are independent and consist of the sum of 0.25 minutes plus a negative exponential random variable with mean 2 minutes, so that the coefficient of variation of the interarrival times is approximately equal to the 0.88 I estimated above. I also assume that the service times are independent and uniformly distributed over five to ten minutes. I simulate 300 independent days for each crew size. The average wait in the queue (excluding service time) per truck is 2.73 minutes for four crews (with a standard error of 0.098), and 0.49 minutes for five crews (0.016). These are 80 percent to 90 percent of the figures obtained by use of the correction factor (3.03 minutes and 0.60 minutes, respectively). While these two examples do not justify the use of the correction factor from a scientific perspective, they do support a managerial decision to use it in this case, when the cost of additional analyses must also be taken into account.

Growers may not receive the undeserved premiums equitably.

My analysis therefore indicates that staffing the fifth dumper during the peak season will cost about $0.70 per hour using the correction factor and possibly twice that or more if an additional correction is made for transient effects. However, I recommend that all five dumpers be staffed during the peak season. Truck waiting costs don't need to be much above $10 per hour to justify staffing five. Even at $10 per hour, the error of staffing on the high side will cost no more than $500 for the season. This is worthwhile insurance against gripes from truck drivers when they have to wait. After all, didn't we just buy the fifth dumper last year, and now we're saying we don't need it? Reducing the staffing to four crews might get some people fired. Furthermore, if employees are absent elsewhere in the plant, workers from one of the dumper crews can be reassigned without substantial negative consequences.

An interesting situation arises on the anticipated seven peak days in which Analysis 2.2 predicts that trucks with wet berries will begin waiting at 4:36 PM. At this point, at most 1,200 bbl/hr (950 wet and 250 dry) can be processed within the plant, so dumper crew staffing should be based on that (smooth) rate, rather than the larger, variable 2,000 bbl/hr rate at which berries are arriving. Without conducting a careful analysis, it appears that two of the dumper crews can be sent home or reallocated to other needs in the plant once the wet berry bins fill up. The remaining capacity of 1,800 bbl/hr should be more than enough to avoid starving the dryers. However, the political cost of sending two dumper crews home just when substantial lines form in front of them is apt to exceed the expected $2(\$6.75)(2.4)(7) \cong \$227$

annual savings, so we recommend that all five dumpers be staffed for all 12 hours from 7:00 AM to 7:00 PM during the peak season. Two crews can be sent home at 7:00 PM on days when there are still trucks left to be unloaded.

Analysis 3.3: Staffing Issues for Other Crews

Clearly a control room operator, a destone/dechaff/dry operator, and all three separator crews, of five people each, are required from 7:00 AM until the facility shuts down every day of the peak season. I recommend that workers be reassigned in the shipping area. The case is inconsistent in this regard, but it appears that the local processing plant can handle 700 bbl/day, delivered in bulk on trucks loaded at one of the two bulk stations. While the case suggests that one worker can run both stations, at a rate of 2,000 bbl/hr total, so that less than one worker hour might be needed theoretically, I assign one worker for eight hours to this activity to cover for the possibility of limited trucking capacity between RP1 and the processing plant. To achieve the planned 950 bbl/hr capacity through the dryers, I recommend that four dryers be devoted to processing wet berries heading for bulk shipment, and one to wet berries to be bagged. I also recommend that all four bulk bin stations be staffed to operate at 800 bbl/hr. Thus, only 400 bbl/hr will flow to the bagging stations. Since the three bagging stations can handle 667 bbl/hr, only two need to be staffed, leaving a capacity of 444 bbl/hr; this capacity should be enough over the incoming flow to cover for any variabilities in the process.

The nonreceiving portions of the plant perhaps should be scheduled to start work later than 7:00 AM because of the likelihood of idleness at the beginning of the day until enough wet berries arrive for processing. For instance, on September 23, 1970, no wet berries arrived until 7:39 AM. Thus, starting a little later in the morning would reduce labor costs and would not significantly increase truck waiting costs. Lode Li [1991] suggests that an optimal stopping model (in this case, an optimal starting model) can be used to exercise this trade-off between labor and truck waiting costs. On the other hand, early morning arrivals may not always consist solely of dry berries loaded the previous day. As growers become better informed of anticipated buildups, they may send some trucks with wet berries earlier in the morning to beat the buildup, changing the pattern of arriving trucks and suggesting that the nonreceiving portions of the plant should be scheduled to arrive around 7:00 AM after all.

The base case assumes that the plant operates over lunch, and operating performance is degraded if the plant closes down for the lunch hour. However, it may be possible to keep processing wet berries at full capacity over lunch. Suppose that plant operators are given a half-hour lunch break. Since trucks appear to stop arriving over lunch, the receiving crews are not needed then and can take some of that time to substitute for workers processing wet berries who need to take their half-hour lunch break. Those processing dry berries in separating and bagging can do the same.

Analysis 3.4: Design of Peak Season Shifts

For any operator position that must be

staffed for at least 12 hours per day, the ideal arrangement is to schedule two shifts, the first eight hours with the full time operators plus any temporaries needed to fill out the staff, and the second lasting at least four hours, and up to eight, with staff being paid only for the hours they are needed. However, I assume that at least eight hours must be scheduled for every shift.

The receiving crews must work 12 hours every day of the peak season, so the current plan of one shift, with four hours overtime (paid at time and a half) every day, is cheaper than adding a second shift. Occasionally, three crews must work a little additional overtime to empty out the waiting trucks.

Langelo [1990] suggests that RP1 not constrain itself to using standard eight-hour days: The case indicates that overtime is paid for hours worked over 40 per week, but not necessarily for hours above eight on any particular day. RP1 should therefore consider changing the schedules. Consider the receiving crews that must be staffed for 12 hours per day. Under the current plan, each position would require four hours of overtime every day, which, when paid at time and a half, amounts to six hours of effective (regular time) pay, for a total of 14 hours per day, and 98 effective hours per week. Another approach is to have one shift work a four-day week at 10 regular hours per day, plus two hours of overtime, resulting in 52 effective hours of pay per week, and another shift work a three-day week at 12 hours per day. This approach results in 88 effective hours paid per week. A better approach is to have two shifts each work three 12-hour days and one six-hour day per week, resulting in 43 effective hours per week per crew, and a total of 86 effective hours. Although this schedule does not achieve the minimum possible 84 effective hours per week, it is close, and schedules that achieve the minimum amount may have other negative ramifications. The fact that the receiving crews must occasionally work additional time (beyond the scheduled 12 hours) to empty awaiting trucks does not materially affect the comparison among the plans.

Wet berries are a nightmare.

Except for the single eight-hour shift for the bulk truck worker, I recommend that all other positions in the plant have two shifts scheduled throughout the peak season. Even with two shifts, there will be overtime for the second shift on most days, although there will be some days when the plant closes down before 16 hours of operation.

RP1 is also having difficulty with absenteeism. Introducing two shifts will reduce the incredible number of hours employees work per day during the peak season and should help reduce absenteeism. Elizabeth Schwerer [1990] suggests that if problems persist, and RP1 must pay workers for full shifts even if they are not needed, then a news-vendor model can be used to determine the optimal number of extra workers to schedule for each shift to cover for unexpected no-shows. I do not offer a detailed analysis, but it appears that the first three workers short would come from the fifth receiving crew.

Analysis 3.5: Annual Bagging Requirements and Bulk Freezer Capacity

One difficulty of the base-case plan is that it may not provide enough frozen bagged berries to meet the needs of the processing plants during the rest of the year, or, viewed slightly differently, the plan calls for a large increase in berries directed to the bulk freezers, which may not be able to handle the additional load. If that is the case, RP1 may need additional bulk freezer capacity. In either event, it makes sense to implement the suggestion made by Juergens, Sherman, and Swenson [1990] to bag an appropriately higher percentage of the berries in the nonpeak season. Another option is to convert to plastic bags from the paper or burlap currently used, so that the moisture from the wet berries will not seep through and cause the bags to stick together when frozen. Other considerations, such as technological feasibility and additional materials and handling cost, may play a role in the evaluation of this option.

There is also another way to bag more berries during the peak season: On higher volume days, the wet berries will be cleared out before the dry ones, at which point all facilities will devote themselves to processing dry berries. The third bagging station can then be put into service. Two of the five workers required can come from one bulk bin station and the others can come from the receiving crews who are no longer needed in receiving. This approach would require additional cost and should be considered only if conditions warrant it.

Analysis 3.6: Tactics to Reduce Truck Waiting

Regardless of the plan implemented, a queue of trucks may form on some days, which suggests that two new tactics be considered.

The first tactic is to load wet berries temporarily into holding bins designed for dry berries only and to process those berries first. Holding wet berries in such bins may be no worse on the berries than holding them in trucks. This tactic would almost completely eliminate truck waiting on a 22,000 day.

The second tactic, which would be needed only if the first tactic could not be implemented, is to reserve one Kiwanee dumper for unloading dry berries. Chances are that the bottleneck is within the plant and there is enough storage capacity for dry berries. However, if a dumper is not reserved for dry berries, trucks carrying them will have to wait to unload as do trucks carrying wet berries, because all of the dumpers will be tied up with trucks carrying wet berries. Implementing this tactic should be resisted, because the logic behind it would be difficult to explain to the growers.

Analysis 3.7: Peak versus Nonpeak Season Issues

Extending the receiving hours to 8:00 PM each evening of the peak season might be productive. Apparently trucks arrive after 7:00 PM under the current system: 12 truck deliveries were listed at exactly 7:00 PM on September 23, 1970: these trucks probably arrived after 7:00 PM. This policy might allow the berry arrivals to be spread out over a longer period of the day, relieving buildup. However, I would first need to ensure that extending receiving hours for another hour would not accentuate the peak days and exacerbate the situation.

Perhaps just recognizing that some trucks currently arrive after 7:00 PM would allow for lower projections of buildup, without any formal policy change.

Extending the length of the peak season may allow daily volume over the peak season to be reduced. If 1970 is a reliable indicator, starting the peak season earlier than September 20 has the advantage that the volume starts out high on the first day of the season; some growers may be delaying their harvesting and delivery until peak season staffing begins to avoid truck delays. However, the percentage of high quality berries is low on the first two days of the season, suggesting that many of these berries have been prematurely harvested. Starting the season earlier may induce further premature harvesting, which would seem to be a bad idea.

In 1970, volumes had already dropped off by the end of the peak season, so little benefit would be obtained from ending the peak season later. However, wet harvesting seems to take place later than dry harvesting, so, with the shift to additional wet harvesting in 1971, the late peak season volumes may warrant extending the season by a day or so.

Introducing a nonpeak-season period with an intermediate level of staffing appears worthwhile, especially right after the peak season: RP1 should plan to start drying, separating, and shipping before the wet berry buildup exceeds 3,200 bbls and provide enough staffing to prevent any truck waiting. For example, if 6,000 bbls of wet berries are to arrive on a given day, then the simple model predicts that trucks will begin waiting at 1:24 PM if processing has not yet started.

Acknowledgments

I am very grateful to Mike Harrison, David Kreps, Avi Mandelbaum, Jerry Miller, for the insights I gained from their analyses of this case. I am particularly grateful to the anonymous author of the unpublished 1977 Harvard Business School Teaching Note, 5-677-197, who presented the essence of the base-case model of Analysis 2.2. Helpful comments and suggestions were made by Xavier de Groote, Albert Ha, Marty Lariviere, Lode Li, Christoph Loch, Bill Lovejoy, Ick Hyun Nam, Andy Tsay, Seungjin Whang, and especially Fred Murphy, who was instrumental in getting this section into print. Thanks go to Steven Gordon and Don Schneider for encouraging their students to submit contributions to this section, and to Joni Podolsky for preparing Figure 2. Special thanks go to Steve Wheelwright who went well out of his way to facilitate the publication of the case, and to Mary Haight for her profound editorial help.

References

Chen, H.; Harrison, J. M.; Mandelbaum, A.; Van Ackere, A.; and Wein, L. M. 1988, "Empirical evaluation of a queueing network model for semiconductor wafer fabrication," *Operations Research*, Vol. 36, No. 2, pp. 202–215.

Hayes, R. H. and Wheelwright, S. C. 1984, *Restoring Our Competitive Edge. Competing Through Manufacturing*, John Wiley and Sons, New York.

Khashoggi, Khalid; Chabas, Bruno; and Bakken, Jan 1990, "Analysis of the National Cranberry Cooperative case," Babson College, Babson Park, Massachusetts.

Juergens, Kenneth G.; Hoyt, Sherman; and Swenson, Chris 1990, "Analysis of the National Cranberry Cooperative case," Babson College, Babson Park, Massachusetts.

Langelo, Judy 1990, "Analysis of the National Cranberry Cooperative case," Babson Col-

lege, Babson Park, Massachusetts.
Li, Lode 1991, written communication, July 10.
Little, J. D. C. 1970, "Models and managers: The concept of a decision calculus," *Management Science*, Vol. 16, No. 8, pp. B466–B485.
Schwerer, Elizabeth 1990, written communication, October 5.
Whitt, W. 1983a, "The queueing network analyzer," *Bell System Technical Journal*, Vol. 62, No. 9, pp. 2779–2815.
Whitt, W. 1983b, "Performance of the queueing network analyzer," *Bell System Technical Journal*, Vol 62, No. 9, pp. 2817–2843.
Woolsey, R. E. D 1978, "How to do integer programming in the real world," Chapter 13 in *Corporate Planning Models*, ed., T. N. Naylor, Addison Wesley, Reading, Massachusetts.

Note: Future of the Case Analysis Section

The case analysis section is intended to help bridge the gap between theory and practice in the MS/OR field. The field has developed a great deal of theory and many methods. The world is full of practical problems, and cases are a convenient medium for representing those problems in print, albeit with certain weaknesses. We need to build more connections between theory and practice, so each can benefit from the other. By publishing case analyses, I hope to help build those connections.

I am changing the format of the case analysis section, to increase interest and activity. Henceforth, authors who wish to publish a case and their analysis of that case are encouraged to submit both to me. If they are accepted, *Interfaces* will publish the case in one issue and the analysis soon thereafter. The analysis should illuminate the process of model building, deriving insights, and developing practical recommendations.