Queue Analysis and Impact of Variability on Process Performance at Au Bon Pain

Olivia Cannizzaro, Yashitha Jeet, Matt Klein, Drew Moses

1. Statement of the Problem

In this analysis we will dissect the multiple processes at Au Bon Pain on Carnegie Mellon Campus. Au Bon Pain is a café that serves salads, sandwiches, and soups along with other quick treats. Many students on CMU campus have a common complaint regarding 'ABP': unreasonably, unpredictably long wait times. Certain times, you might find yourself strolling into an entirely empty café, where a group of idle staff members rapidly attend to your order. Other times the space is so packed that you could wait up to 45 minutes to check out with your food.

There are multiple problems apparent with the structure of Au Bon Pain. The high variability in waiting times and the café's susceptibility to very long queues is the clearest one. We also observe that when the store is not busy the staff has nothing to do and is being under-utilized, but when they are busy the staff is stretched too thin and significantly over-utilized. This leads us to a few, related questions:

A) How much money can Au Bon Pain save/make by adding x additional workers?

B) What other steps can management take to reduce their inefficiencies, i.e.

  i. Pooling
  ii. Priority rules in waiting time systems
  iii. Appointments
  iv. Employee training programs

C) By how much can we reduce queue times for customers by addressing these inefficiencies?

We believe it likely that a customer's wait time is inversely proportional to the likelihood of them to returning to ABP in the future; making this issue worth looking into for the café. All of these questions can be answered through queue analysis, and by understanding the impact of variability on process performance. Multiple necessary, simplifying assumptions will be made throughout the course of this report. In lieu of listing them all here, we will address them as they become relevant in the context of our analysis.

2. Approaches to Solve

2.1. Defining the Process

We start with our first question: how much money can Au Bon Pain save/make by adding x additional workers?

We will refer to the complete experience of getting food from ABP as a *process*. In order to answer our question, we must decide where within the process it would be best to add these workers. We will examine three main activities from the process:

1. Making the food
2. Delivering the finished order to the customer
3. Checking out the customer at the register

2.2. Analyzing the Current Process

Let's begin by looking at the current setup of the café. Through field observation, we discovered that there are typically 2 workers dealing with food preparation, one worker who gives finished orders to customers, and 2 workers at the cash registers. Notably, the employee who gives customers their finished orders is also the one who takes orders from incoming customers; this is why customers can end up waiting long times to receive their already-prepared food.

Each of our team members went to Au Bon Pain and ordered a variety of menu items, personally timing each of the three aforementioned activities. The data we ended up with is summarized in Table 1 below:

Table 1: Activity Times (minutes)

Food Preparation	Food Pick Up	Cash Register
5.60	0.25	4.57
15.25	0.57	1.53
8.85	0.15	1.70
6.00	0.15	1.70
0.43	0.15	3.42

Given that we are restricted to so few data points, we attempt to interpolate and fit them to a known distribution. Based off of observations, we felt that the normal distribution would be appropriate for all three. There seems to be a ”standard” experience of any one activity taking x amount of time, and despite high variability, the extremely long or extremely short wait times seemed much less common. Below are histograms for the data we collected, with the normal distribution curve overlaid.


Activity Time for Food Preparation
Density
0.00
0.02
0.04
0.06
0.08
0.10
0.12
0
5
10
15
20
Time (minutes)

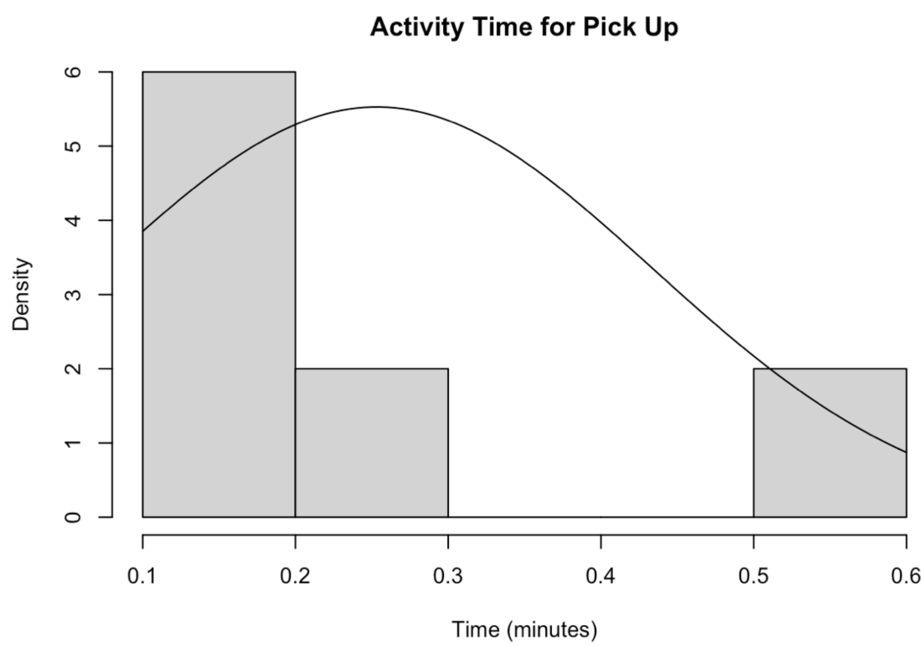
Activity Time for Pick Up
Density
0
1
2
3
4
5
6
0.1
0.2
0.3
0.4
0.5
0.6
Time (minutes)

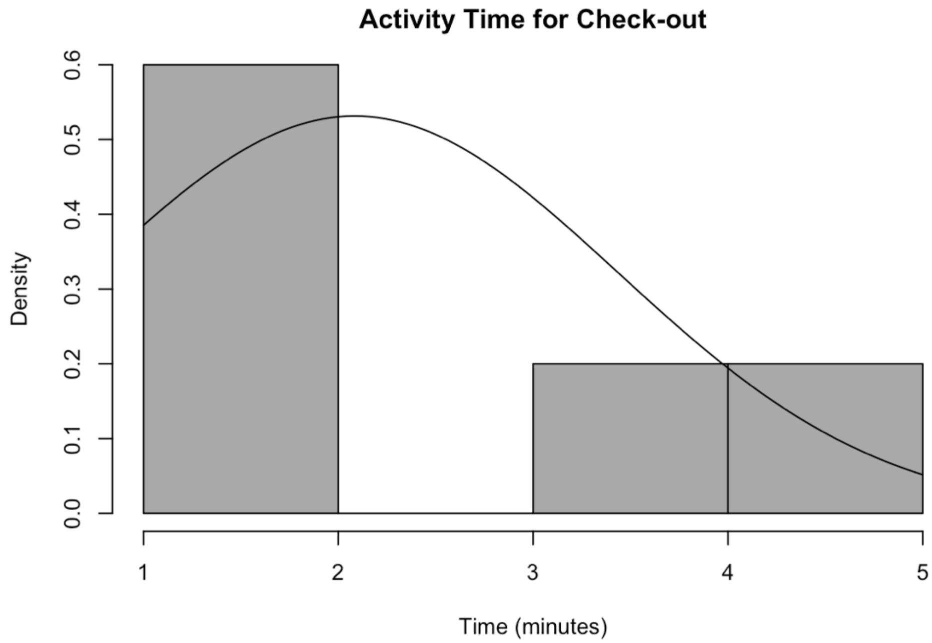

We can now summarize the average service times along with their standard deviations:

Table 2: Average Activity Times (minutes)

	Mean	Standard Deviation
Food Preparation	7.226	5.417
Food Pick Up	0.253	0.180
Cash Register	2.583	1.351

We define *utilization* as the proportion of time that the process resources (employees) spend doing actual work as opposed to experiencing idle time. Even though the employees experience intermittent periods of idle time, variability in demand means that customers still have to face long wait times at certain periods throughout the day.

2.3. Analyzing the Customer Flow

We know that the rate at which customers arrive is highly variable, and that inter-arrival times are commonly distributed following an exponential distribution. We will move forward using this assumption, meaning that the mean and standard deviation of the inter-arrival times are equivalent.

2.4. Dissecting the Waiting Time Formula

At this point we introduce the waiting time formula. This formula is defined for m parallel resources; in this case meaning that there are m employees working separately but in parallel on the same type of task. This is not two people working on the same sandwich, but rather two people both working on sandwiches. "Activity time" in this context will refer to the average time taken for any one of the three activities outlined in Table 2 to be completed.

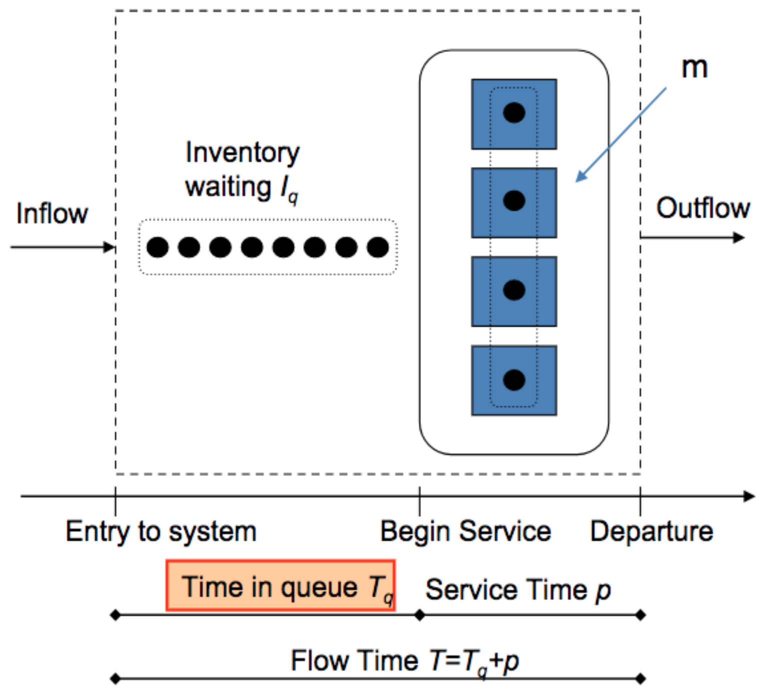

Figure 1: Model for waiting with m parallel resources

Formulaically, the equation is:

$$\textit{Time in queue} = \frac{\textit{activity time}}{m} * \frac{\textit{utilization}^{\sqrt{2(m+1)}-1}}{1-\textit{utilization}} * \frac{CV_a^2 + CV_p^2}{2}$$

We define

$$CV_a = \frac{\text{standard deviation(inter-arrival times)}}{\text{average(inter-arrival times)}}$$

and

$$CV_p = \frac{\text{standard deviation(activity time)}}{\text{average(activity time)}}$$

Au Bon Pain is open most days from 7:30AM until 2:00AM the following day, giving them an 18.5 hour workday. From observation we know that they have two main rushes, one at lunch and one at dinner, which each last around 2 hours. During this 4 hour cumulative period, the employees appeared to be working without break. We also observed that during late night hours, from around 12AM to 2:00AM, there are virtually no customers. During the other 12.5 hours of the day customers come in and out pretty variably but consistently, and we will estimate that the employees are busy roughly 65% of the time on average.

Using this, we can estimate utilization as $\frac{(1*4)+(0*2)+(.65*12.5)}{18.5} * 100 = 65.54\%$. We assume this utilization for each of the three activities, given that all three are always involved in any customer purchase.

From this formula, we can identify three different ways to reduce a customer's time spent waiting. Looking at the first term in the formula, we can easily decrease our activity time by increasing m, the number of workers. From the second term, we can see that by decreasing utilization, we can also reduce the time spent in the queue. Intuitively, this makes sense because employees would have additional idle time during which they could immediately service any incoming customers. The third term shows us that by decreasing the standard deviation of inter-arrival times, we might also be able to reduce time spent in the queue.

2.5. Assessing The Options

1. Increasing the number of workers

In order to determine if it would be worthwhile to add another worker, we will perform cost-benefit analysis. From researching Glassdoor for salary information, we discovered that cashiers make $9 per hour and the cooks make $8.50 per hour. Since our data showed that time spent picking up food is typically less than 30 seconds, we can just focus on food preparation and the check-out line.

Au Bon Pain's hourly cost of adding another worker at the food preparation step is then \$9. This step has the longest average waiting time, making it the *bottleneck* of the process. In order to increase capacity for the entire process, we want to focus on the bottleneck, so we will perform our calculations for the food preparation step. To assess the benefit, we can start by calculating the difference in the waiting time formula before and after the extra worker is added. Adding an extra worker also decreases the utilization (formulaically, $u = \frac{p}{a*m}$), so we have:

$$Time\ in\ queue_1 = \frac{7.226}{2} * \frac{0.6554^{\sqrt{2(2+1)}-1}}{1-0.6554} * \frac{1^2 + \frac{5.417}{7.226}^2}{2} = 4.4384$$

$$Time\ in\ queue_2 = \frac{7.226}{3} * \frac{0.6554 * 2/3^{\sqrt{2(3+1)}-1}}{1-(0.6554*2/3)} * \frac{1^2 + \frac{5.417}{7.226}^2}{2} = 0.7352$$

Since we start with only two workers, adding one additional employee represents a 50% increase in staffing. This is why adding just one worker leads to such a drastic decrease in a customer's expected time spent in the queue.

Considering the average experience, a customer spends the sum of average activity times plus queueing times in the process. We assume that the activities are carried out sequentially by the customer. We can calculate time in queue for food pick up and at the cash register using the waiting time formula as well. It is important to note that at the register we have m parallel processes as opposed to m parallel pooled resources. This idea of pooling is elaborated on later in the analysis, but for now, suffice to say that $m = 1$ for the cash register since each register has its own, independent waiting pool.

$$Time\ in\ queue_{\text{ register}} = \frac{2.583}{1} * \frac{0.6554^{\sqrt{2(1+1)}-1}}{1-0.6554} * \frac{1^2 + \frac{1.351}{2.583}^2}{2} = 3.128$$

$$Time\ in\ queue_{\text{ pick up}} = \frac{0.253}{1} * \frac{0.6554^{\sqrt{2(1+1)}-1}}{1-0.6554} * \frac{1^2 + \frac{0.180}{0.253}^2}{2} = 0.362$$

Now, we can calculate the total average time a customer spends in the process as $7.226 + 4.438 + 2.583 + 3.128 + 0.253 + 0.362 = 17.99 \approx 18$ minutes. Our decrease in wait time from adding an additional worker is $4.438 - 0.735 = 3.703$ minutes, meaning the new average time spent in the process is $18 - 3.703 = 14.297$ minutes. At this point, it is worthwhile to introduce Little's Law:

$$\textbf{inventory} = \textbf{flow rate} * \textbf{flow time}$$

Inventory, in this case, being defined as the total number of customers currently in the process. On the same day, if x customers are currently in the process, we can use this change in flow time to calculate a new flow rate. This new flow rate will tell us how many more customers we could service per hour with the decrease in flow time, assuming ample demand:

$$\textbf{flow rate}_1 = \frac{\textbf{inventory}}{\textbf{18}}$$

$$\textbf{flow rate}_2 = \frac{\textbf{inventory}}{\textbf{14.297}}$$

$$\textbf{flow rate}_2 * \frac{14.297}{18} = \textbf{flow rate}_1$$

$$\textbf{flow rate}_2 = \frac{18}{14.297} * \textbf{flow rate}_1 = 1.26 * \textbf{flow rate}_1$$

This implies that we would have 1.26 times as many customers flowing through the process per hour by adding an additional worker. For this to be profitable given the $9 cost per hour of the additional worker, assuming each customer purchases about $7 worth of food (the average cost of a sandwich at ABP), we would need:

$$[(x \text{ customers} * 1.26) - (x \text{ customers})] * 7 \geq 9$$

Where x is the number of customers typically flowing through the process in one hour. Solving for x, we get:

$$x \text{ customers} \geq 4.95$$

It is reasonable to assume that our current flow rate is at least five customers, so adding an additional worker at the food preparation step does indeed appear to be a profitable change for Au Bon Pain to implement.

2. Decreasing utilization

The best way to decrease utilization is by adding staff, the analysis for which is done above. Since Au Bon Pain already has an 18.5 hour workday it is unrealistic for them to extend hours, so decreasing utilization in this way does not appear to be a valid option. However, by incentivizing customers to come at non-peak hours, they might decrease utilization throughout the day making them more equipped to handle non-expected rushes.

In the **Given More Time** section at the end of this report, we detail how one might use promotional projects, sales or specials in order to even out their customer flow.

3. Decreasing the standard deviation of inter-arrival times

There are two main managerial responses to variability that are worth considering in this scenario. One is the idea of pooling: this is relevant to the check-out area. The current set-up of this area is a series of parallel lines at independent registers. Since there are typically two workers on staff, these are effectively two demand arrival processes being processed by two identical servers; a customer waits at the specific register line at which they've arrived.

The idea of pooling would be to combine these two systems into a single system with one waiting area and two identical resources. Au Bon Pain would designate a single waiting line, and the next person in line would go to whichever cashier is available first. This idea is illustrated in the figure below.

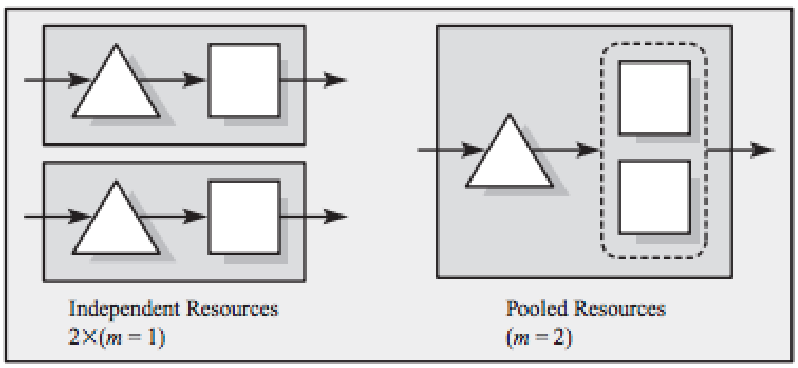

Figure 2: The concept of pooling

Formulaically, utilization is $\frac{\text{processing time}}{\text{arrival rate}*\text{m}}$. For a pooled process with m resources, arrival rate $_{\text{pooled}}$ $= \frac{1}{2} *$ arrival rate $_{\text{original}}$. Since our arrival rate is halved but our m is doubled (two processes each with $m = 1$ versus one process with $m = 2$), these two cancel out and we get that utilization is the same in either case. Therefore, for the separate, parallel processes we have:

$$Time\ in\ queue = \frac{2.583}{1} * \frac{0.6554^{\sqrt{2(1+1)}-1}}{1-0.6554} * \frac{1^2 + \frac{1.351}{2.583}^2}{2} = 3.128$$

But for the pooled process we have:

$$Time\ in\ queue = \frac{2.583}{2} * \frac{0.6554^{\sqrt{2(2+1)}-1}}{1-0.6554} * \frac{1^2 + \frac{1.351}{2.583}^2}{2} = 1.416$$

The pooled process uses the available capacity more effectively because it avoids the situation where one of the cashiers is idle while the other has an excess of customers to attend to. Pooling allows them to service the same number of customers with the same utilization, same processing time, and same number of workers, but in less than half the waiting time for the customer! The following figure illustrates the concept of decreases in waiting times as you increase the number of pooled resources.

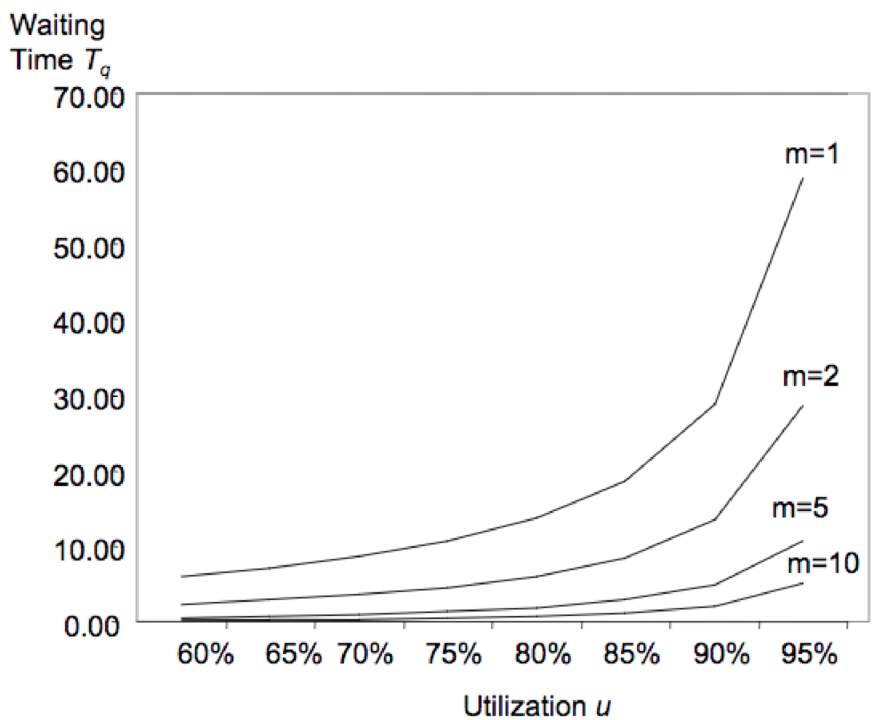

Figure 3: Wait times for differently sized resource pools

If we look at employee responsiveness versus efficiency, we also see that pooling would shift Au Bon Pain's efficiency frontier outwards without having to compromise either variable or add new resources. This cost-free system improvement is illustrated below:

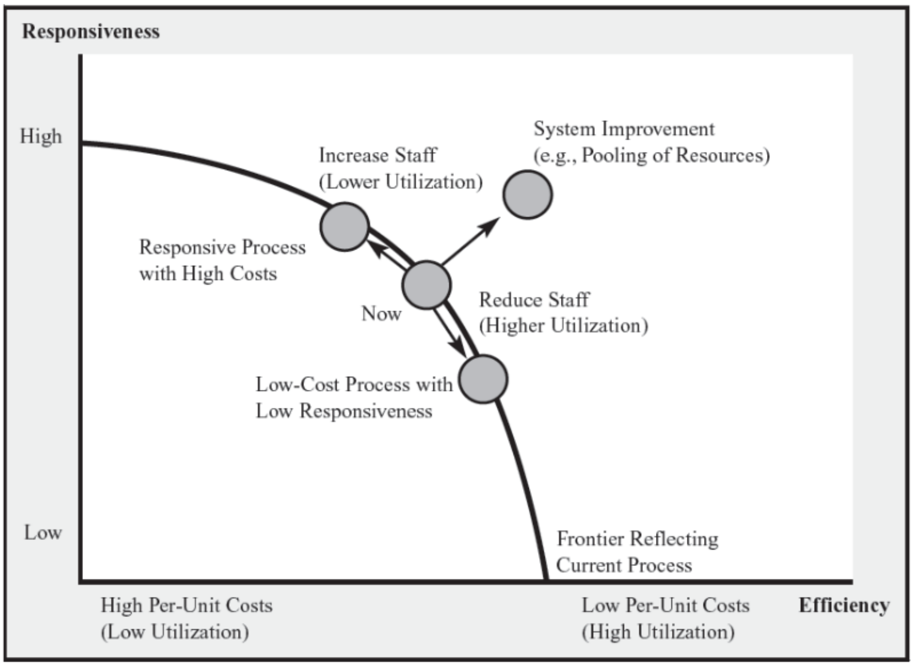

Figure 4: The efficiency frontier

The other potential response to variability would be to create an appointment system- this will be discussed in the **Given More Time** section.

3. Results Obtained and Recommendations

From this in-depth analysis, three key strategies emerge to not only drastically reduce customers' wait times, but also drastically increase Au Bon Pain's profits.

The first recommendation is to add another worker at the food preparation step. Under the assumption that there is ample demand to meet the increased capacity, and that the currently flow rate is at least 5 customers per hour, this addition would be profitable for ABP. Not only would this addition decrease average queue time, but it also would decrease utilization of the employees. This is dually beneficial, as it eases the burden on the customer and on the employees alike.

The next recommendation is to pool the waiting area for the registers into a single line. This is a definitively profitable decision for Au Bon Pain given that it increases their flow rate at zero cost to them. It shaves an additional $3.128 - 1.416 = 1.712$ minutes off the customer's flow time, making our flow rate $\frac{18}{14.297-1.712} = 1.43$ times its initial value (before any changes are implemented). If Au Bon Pain were, for example, to initially service on average 10 customers per hour, these improvements would allow them to service $10 * 1.43 = 14.3$ customers per hour. Assuming every customer earns Au Bon Pain $7, their net profit from these improvements is $[(14.3 - 10) * 7] - 9 = \$21.1$ per hour. In an 18.5 hour work day, this would be an additional $21.1 * 18.5 = \$390.35$ each day.

4. Given More Time

Despite the depth of analysis that we have conducted, our team was still constrained by our time and resources. Given a longer time frame and better access to Au Bon Pain's data, there are multiple other analyses that could have been explored.

The first additional analysis we would've liked to explore would be the waiting time to order food. In our analysis, we neglected to include the time it takes to place an order for prepared food. If we had analyzed this waiting time, we would have treated it like our other waiting times and measured the distribution of times then made a recommendation for improvement. As it is now, the same employee who takes orders also hands customers the prepared food, this seems like a potential bottleneck during peak hours. One possible solution to this problem could be to have an automated system that knows when a customer has paid and then have the cooks place the prepared on the counter. That way the employee who takes orders and hands out fresh meals would only have the former task.

The second additional analysis we would've liked to explore would be the use of equipment, specifically the use of ovens. Multiple customers in a row could order hot sandwiches and then the limiting factor in the system becomes oven space. Given a cost of installing an additional oven, rate of replacement, and rate of breakdown, we could have performed calculations to determine an ideal number of ovens and give a recommendation of how many to add in the case the existing ovens are not sufficient.

The third additional analysis we would've liked to explore would be methods of decreasing utilization. We could research how promotions which are offered at non-peak hours would affect customer flow. Then we could perform cost benefit analysis on the loss from the sale versus the gain from lowered utilization.

The fourth additional analysis we would've liked to explore would be improvement of data collection regarding waiting times. Two possible methods of improving data collection include: conducting more field research and CMU Dining/ ABP Management sharing data with our team. Given the time frame of the project, we could not conduct field research at a level of a professional consulting firm. Additionally, our requests for data were left unanswered by CMU Dining and ABP Management.

The fifth additional analysis would've been to recommend implementation of an appointment system to reduce variability of order flow. We could have researched possible implementation options, gauge user interest, and produce research on the affect on customer arrivals. Given these steps, we would determine if it's optimal to implement such a system and which variety of system would be ideal.

5. Works Cited

Sources:

Impact of Variability on Process Performance; PowerPoint presentation by Sham Kekre for 70371 Operations Management

Matching Supply and Demand, third edition; textbook by Cachon and Terwiesch