EE126 Project: Cafe Management Optimization

Lucine Oganesian and Alex Lin

October 22, 2015

1 Introduction

On college campuses, there is always a high demand for places to work. Cafes are popular spots for students and faculty alike. Many members of the campus community prefer working at cafes, where they can procure a cup of coffee and a settle down to work for an indefinite period of time, amidst the inspiring murmurings of other cafe goers[1].

However, it is not uncommon for potential cafe customers to walk into a cafe and quickly scope out seating availability. If there are no available seats, the customer will depart, with neither customer nor cafe owner benefiting from the exchange. Thus, it is in the best interest of both business owners and customers for cafes to optimize their available seating based on the expected behavior of their clientele.

The goal of this project was to create two models for cafe owners to optimize the number of tables in their stores to maximize their profits and, as a byproduct, the happiness of their customers.

2 Experiments

2.1 Methods

We started with a simple model: a random number of customers would come at the beginning of each hour, and stay for exactly one hour. This way, the maximum number of people in our cafe at any time would just be the number of people who showed up for that hour, and each hour would be independent of the next. Our second model involved a constant number of people arriving at the beginning of each hour and staying for a random amount of time[2]. In all cases our objective function was as follows: (N is the number of tables, P is the number of customers)

$$\min N\alpha(N-P)^+ + \beta(P-N)^+$$

subject to the constraints that $N > 0$, $N * C <= B$ (where C is the cost per table and B represents the budget allocated for tables).

α and β are the parameters that control the relative importance of saving money on tables and hoping customers without a table will still buy something to go. A higher α value implies more stinginess, since it penalizes having more tables than customers. Note: our model assumes that a table will seat one person (1-to-1 correspondence and no table will seat two people) and that if there is a free table someone will sit there.[3] Furthermore the α and β weights were related as $\alpha + \beta = 1$.

Every model had two forms: one in which the expectation was pushed through the objective function and a second in which the expectation was empirically calculated by taking an average over objective functions, that is:

$$\min_N \frac{1}{L}\sum_L \sum_h \alpha(N-P_{l,h})^+ + \beta(P_{l,h}-N)^+$$

where L is the number of samples and h is the number of hours the cafe is open (for both models this was set as 8am-12am). Part of our analysis was comparing the two models.

[1] In our models we assume customers come and go independent of one another, and that cafe goers murmur at a volume that does not adversely affect the ambiance

[2] This seems a better fit to reality, since classes end on the hour, but students have varying endurance levels for cafe studying

[3] Thus, if a Stanford students braves a Berkeley campus cafe he will always take an empty table next to a Berkeley student, even if he taunted by Berkeley students for the duration of his stay. If you can't be 'em, at least be near 'em.

3 Analysis and Discussion of results

Our first analysis of our stochastic optimization was to check that Jensen's inequality was satisfied[4] and to make sure our results made sense. Next, to explore the relationships among the components of our models, we varied parameters and examined their effects on the optimal number of tables the cafe shop owner should buy. The parameters we investigated included: the distributions that represented the student stay time (model 2), α/β weights (both models 1 and 2), and the arrival/stay rates used (Poisson in model 1 and exponential in model 2 respectively).

3.1 Jensen's inequality

Both model 1 and model 2 satisfied the requirement that $f(N, E[P]) \leq E(f(N, P))$, meaning that our implementation, at least on this count, did not violate any rules. For model 1, pushing expectation through yielded 10 as did just calculating the expectation. For model 2, pushing expectation through yielded 10, whereas calculating the expectation yielded $12 \geq 10$.

3.2 α and β weight values

Figure 1 shows that as α gets larger, the optimal number of tables to buy decreases because a high α means we are adamantly against wasting money on tables that are not used. With a lower budget, the budget constraint dominates over the effect of α; Conversely, for moderate and high budgets, α quickly dominates as the deciding factor. This pattern is observed no matter what the customer arrival and staying times are; the graphs from model 1 and model 2 are effectively identical, meaning that even a simple model (model 1) performs reasonably well in describing the desired behavior, compared to a more complicated model (model 2).

3.3 Rates

Figure 2 shows the changes that come about when the average amount of time customers stay is varied (for model 2, Figure 2b) and when the amount of people who arrive on average is varied (for model 1 - Figure 2a).

3.4 Distributions

Another parameter we changed was the distribution from which the random variable representing how long each customer stays was sampled from. We tried exponential (original), Gaussian, and uniform distributions. For 10 samples, the results for each distribution (plotted against varying alphas for different budget values) showed hardly any difference between the distributions (Figure 3). So even though an exponential distribution would be the most logical choice for modeling the customer stay time, the distribution from which the stay times are sampled from do not seem to have a marked difference between them.

3.5 Limitations

These models were all oversimplifications of how a cafe actually works. However we hope that the simplifying assumptions made were legitimate in the context of campus cafes such that the calculated number of optimal tables could act as a worst case scenario (i.e. an upper bound) on the true number of tables needed. More realistic models may yield a better estimate of the true number, but would have been more difficult to implement. Future directions might include exploring slightly more complicated models, for example, modeling the number of customers as a continuous time markov chain, or including a model of their spending.

4 References

a. Lab5 part 2 was a source of inspiration for this project.

b. Kabir, and his example problem on piazza, helped us formulate our objective function.

c. `http://articles.bplans.com/13-tips-open-successful-coffee-shop/` [5]

[4] `https://en.wikipedia.org/wiki/Jensen%27s_inequality`

[5] Not actually a reference. But perhaps a reference for our future business venture.

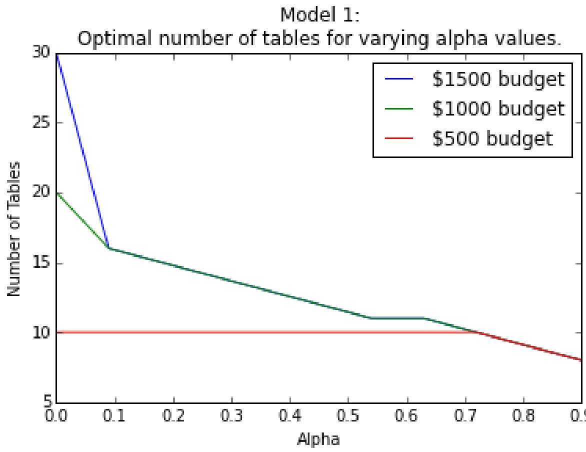

(a) Model 1

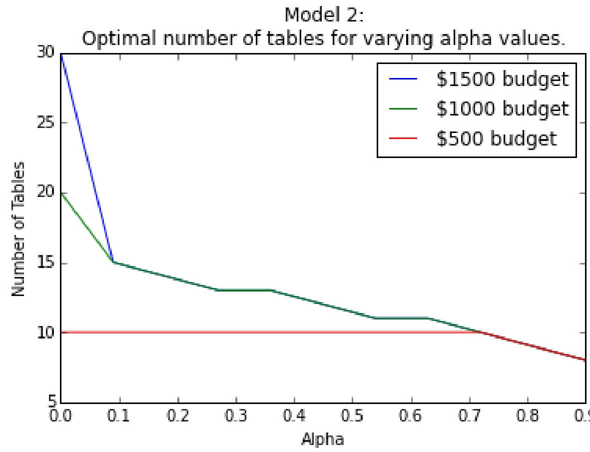

(b) Model 2

Figure 1

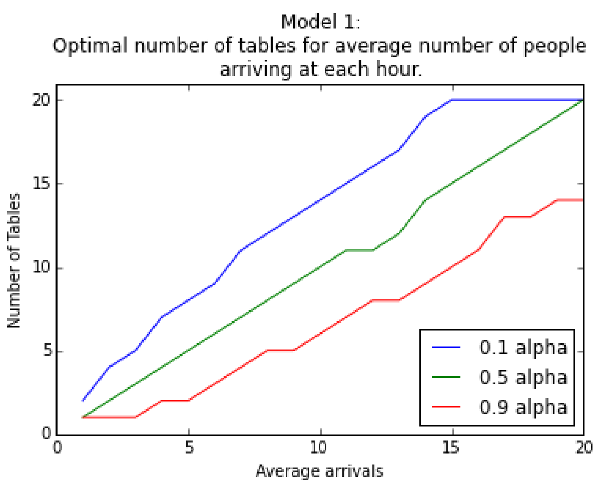

(a)

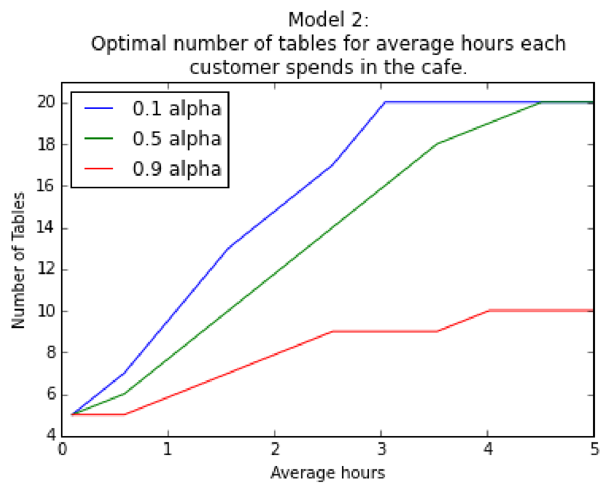

(b)

Figure 2

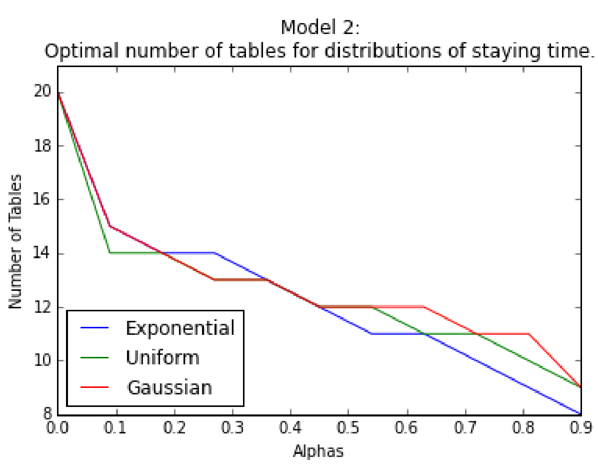

Figure 3