YIELD MANAGEMENT IN HEALTH CARE CLINICS WITH WALK-IN TRAFFIC

Linda LaGanga
Stephen R. Lawrence

Mental Health Center of Denver, Denver, Colorado, USA
Leeds School of Business, University of Colorado, Boulder, Colorado, USA

ABSTRACT
An important problem in health care operations involves yield management in health clinics where patients often do not arrive for scheduled appointments (are "no-shows"). A second issue is the frequent presence of "walk-in" traffic where patients arrive for service unscheduled and unannounced, but must be served. In this paper, we investigate the management and scheduling of health care clinics in the presence of both walk-in traffic and appointment no-shows. We build on recent research that has developed methods to determine optimal and near-optimal appointment schedules when there are patient no-shows. Our analysis employs a clinic utility model that works to balance the benefits of seeing additional patients (both scheduled and walk-in) with the implicit costs of patient waiting time and the possibility of clinic overtime. Results demonstrate that compared to an all-walk-in clinic, even a modest number of scheduled appointments can significantly improve patient service and clinic performance.

Keywords: Service operations, yield management, appointment scheduling

INTRODUCTION
An important problem in health care operations involves yield management in health clinics where patients may not arrive for scheduled appointments (are "no-shows"). The complexity of appointment scheduling increases further in the presence of "walk-in" traffic where patients arrive for service unscheduled and unannounced. Walk-in traffic occurs when clinics permit or are required to allow patients to arrive for service at their own volition and according to their own schedule. Walk-ins can also occur in hospital settings where in-patients are brought to a clinic for treatment on an as-needed basis.

In this paper, we investigate yield management in health care clinics in the presence of scheduled appointments and walk-in traffic. We assume that one set of patients arrive to a clinic for previously scheduled appointments, and that a second set of patients arrive without appointments but expect to be serviced within a reasonable length of time. To address clinic scheduling with walk-ins, we build on recent research that has developed methods to determine optimal and near-optimal appointment schedules where some fraction of patients do not show for their appointments (LaGanga and Lawrence 2008). We extend this research to allow for the stochastic arrival of walk-ins that are either seen after scheduled patients, or are given priority ahead of waiting patients. We employ a clinic utility model that works to balance the benefits of serving additional patients (whether scheduled and walk-in) with the implicit costs of patient waiting time and the possibility of clinic overtime.

We assume that some fraction of scheduled patients are no-shows and that the arrival pattern of walk-ins is stochastic. Both the frequency of no-shows and the arrival pattern of walk-ins can vary by time of day. Computational results demonstrate that even when compared to an all-walk-in clinic, even a modest number of scheduled appointments can significantly improve patient service and clinic performance. Our results will prove useful to both health-care practitioners and researchers alike. From a practical perspective, our model allows clinician administrators to better manage a mix of walk-in and scheduled patients, and provides a means to evaluate the impact of walk-in traffic on the performance of a particular clinic. From a research perspective, our model provides a robust and effective means of evaluating yield management in health care clinics in the presence of no-shows and walk-ins. It can readily be extended to other clinic problems such as open access scheduling where patients call-in and are scheduled for the first open appointment slot on the same day or in the near future.

This line of research makes several contributions. To our knowledge, we are the first to formally examine appointment scheduling in the presence of walk-in traffic in any context. We demonstrate the trade-offs between variance reduction of appointment scheduling with the flexibility of walk-ins and provide a means to measure the costs and benefits of each. We show that the costs of introducing some walk-in traffic into an all-appointment clinic are modest, and that the benefits of introducing some appointments into an all-walk-in clinic are large. And finally, we demonstrate how our results can easily be integrated into and support an open access clinic where same-day appointments prevail.

A CLINIC SCHEDULING MODEL WITH APPOINTMENTS AND WALK-INS

We model a health care clinic that cares for patients both by appointment and by walk-in. The goal of the model is to represent the scheduling problem faced such a clinic as it works to maximize its access to patients while minimizing patient waiting time and clinic overtime. We first describe assumptions of the model, then define a clinic utility objective function, and finally outline a near-optimal solution procedure for the objective.

Model assumptions

We model a health care clinic that sees patients during a clinic session of duration D time units (notation is summarized in Table 1). A clinic session is a period of time (*e.g.,* a morning, a day) during which the clinic is in continuous operation, followed by a break (*e.g.,* a noontime break, end of work day, etc.). Each appointment "slot" j during a session is of fixed duration d and the length of a session is designed to be an integer multiple of d so that $D=Nd$, where N is the number of appointment slots scheduled in a session. We set the duration of each appointment to $d=1$ time units so that the length of a clinic session is $D = N$, without loss of generality. Our assumption of fixed appointment duration is appropriate in many clinics (*e.g.*, mental health clinics, routine care clinics, and many diagnostic testing clinics) where there is intrinsically little or no time-variability in the services provided. For clinics with more variable service durations, an appointment schedule that assumes fixed service times can nonetheless provide a useful heuristic solution (LaGanga and Lawrence 2007). The number of providers P in the clinic is fixed during a clinic session so that P patients may receive service during any appointment slot, $P \times N$ patients can be seen in a clinic session without overtime. Patients may be assigned to any provider, and are seen in the order that they arrive, whether scheduled or walk-in. Finally, if patients remain waiting for service at the end of a clinic session, the clinic will work overtime until all patients have been seen – patients are not sent away without service.

We assume that demand is exogenous to a clinic and is beyond its control, at least in the short run. Patients of the clinic either arrive for scheduled appointments or are walk-ins who arrive unscheduled and unannounced. Patients who arrive for appointments are assumed to be punctual (Soriano 1966), but some do not appear for their appointments (are no-shows) so that the patient show rate σ is often significantly less than 100% ($0 \leq \sigma \leq 1$). The number of patients arriving for an appointment slot j is therefore binomially distributed (Ross 1985, 26-28) with parameters a_j (the

Table 1 – Notation

$\hat{A}$	Expected number of arriving patients
$D = Nd$	Duration of a session
d	Duration of an appointment (deterministic)
N	Number of appointment slots in a session
P	Number of providers ($P \geq 1$)
S	Number of appointments scheduled for a session ($S \geq N$)
σ	Average show rate of for all scheduled appointments ($0 \geq \sigma \geq 1$)
π	Marginal net benefit of one additional patient
τ	Marginal cost or penalty of clinic overtime ($F > C$)
ω	Marginal cost or penalty of patient wait time (per patient)
s_j	Number of patients scheduled for service in slot j
a_j	Number of scheduled patients that actually arrive in slot j
λ	Arrival rate of walk-in patients (patients per appointment slot)
α_{jk}	Probability that k patients arrive in appointment slot j
θ_{jk}	Probability that k patients are ready for service at the start of slot j after new patient arrivals
S	Vector of number of patients s_j scheduled for each slot j
$\Pi(\cdot)$	Net benefit function
$\Omega(\cdot)$	Patient waiting cost function
$\mathrm{T}(\cdot)$	Clinic overtime function
$\mathrm{U}(\cdot)$	Utility function, where $\mathrm{U}(\cdot) = \Pi(\cdot) - \Omega(\cdot) - \mathrm{T}(\cdot)$

number of patients scheduled for slot j) and show rate σ. In this paper, we hold the appointment show rate constant across a clinic session, but our model can easily be adapted to to accommodate show rates that vary by appointment slot (LaGanga and Lawrence 2008). For walk-ins, we assume that walk-in patients arrive according to a Poisson distribution with mean λ. This assumption of Poisson arrivals is standard in queuing theory and has been empirically validated in numerous studies (Ross 1985).

Given show rate σ and walk-in rate λ, the clinic may decide to either overbook or underbook appointment slots so that the total number of appointments scheduled S varies from the total number of provider slots $P \times N$. The problem of the clinic is to create a schedule vector **S** with elements s_j, the number of patients scheduled to arrive at the start of each appointment slot $j\left(1 \leq j \leq N\right)$ so that the *utility* of the clinic is maximized (defined below).

Clinic utility objective function

We adopt a general objective function for the appointment overbooking problem that works to balance the interests of the clinic with those of patients and of service providers (LaGanga and Lawrence 2007). This objective trades off the benefit of servicing additional patients with the costs or penalties for keeping some patients waiting for service plus the expected cost of clinic overtime incurred when all scheduled patients cannot be seen during a clinic session. We address each of these elements below.

For every patient serviced, some net benefit π is generated for the clinic, or for its goals and objectives. This benefit can represent net financial profit (after variable costs), community service delivered (in the case of not-for-profit organizations), accrued goodwill or some combination of these and other benefits. The gross benefit of servicing a patient is reduced by variable service costs such as the costs of materials, external services, and other costs that vary directly with the number of patients served. Note that the costs of service providers are generally not included in benefit calculations since these costs are usually sunk and do not vary directly with the scheduling policy employed.

Depending on the construction of appointment schedule **S**, some patients may be required to wait for service if more patients have arrived than there is open capacity and thus incur a waiting cost. Patient waiting costs ω are incurred for each patient that must wait one appointment slot for service, and include lost patient goodwill and reduced satisfaction, lost future business, increased future no-show rates (Dyer 2005, Lowes 2005), and other similar expenses.

Uncertainty about patient no-show and walk-in rates may result in some patients waiting for service at the end of a clinic session (LaGanga and Lawrence 2007), so that the clinic may be required to work overtime. Clinic overtime costs τ are incurred every time unit that the clinic must work overtime, and can include lost patient and staff goodwill, overtime wages paid to clinic staff, and increased staff turnover.

Combining the expected benefits of seeing additional patients $\hat{\Pi}(\mathbf{S})$ less the expected costs of patient waiting $\hat{\Omega}(\mathbf{S})$ and expected clinic overtime $\hat{\mathrm{T}}(\mathbf{S})$ provides a representation of expected clinic utility $\hat{U}(\mathbf{S})$:

$$\hat{U}(\mathbf{S}) = \hat{\Pi}(\mathbf{S}) - \hat{\Omega}(\mathbf{S}) - \hat{\mathrm{T}}(\mathbf{S}) \tag{1}$$

$$\hat{U}(\mathbf{S}) = \pi\hat{A} - \frac{\omega}{\hat{A}}\left(\sum_{j=1}^{N}\sum_{k}(2k-1)\theta_{jk} + \sum_{k}\sum_{i=1}^{k}(i-1)^2\theta_{N+1,k}\right) - \tau\sum_{k}k^2\theta_{N+1,k} \tag{2}$$

The first term represents the net benefit of seeing an expected $\hat{A}$ patients with schedule **S**. The second term represents the expected waiting costs incurred, and the final term is expected overtime costs. Note that this representation assumes quadratic waiting and overtime costs, since long waits are disproportionately more annoying, disruptive, and costly than short waits (Maister 1984). A complete derivation of this utility function is available in LaGanga and Lawrence (2008).

In the utility equation above, variable θ_{jk} represents the probability of the number of patients waiting for service at the start of appointment slot j, including new arrivals. Let α_{jk} be the probability that k new patients arrive in slot j, either for scheduled appointments or as a walk-ins. Define $\alpha(\cdot)$ as the discrete distribution of the number of new arrivals in slot j, which is the convolution of binomially distributed appointments $b(s_j,\sigma)$ and Poisson-distributed walk-ins $p(\lambda)$:

$$\alpha(s_j,\sigma,\lambda) = b(s_j,\sigma) * p(\lambda) \tag{3}$$

the elements of which can be calculated as

$$\alpha_{jk} = \sum_{i=0}^{k} b(s_{ij},\sigma)\, p_{k-i}(\lambda) \tag{4}$$

These new arrivals will join any remaining patients waiting from prior appointment slots. The joint probability for the number of patients waiting for service at the start of an appointment slot after all new arrivals is:

$$\theta_{j,k} = \theta_{j-1,0}\alpha_{j,k} + \sum_{i=1}^{M}\theta_{j-1,i}\alpha_{j,k} + \sum_{i=M}^{k}\theta_{j-1,i+1}\alpha_{j,k-i} \tag{5}$$

The first term of this expression is the joint probability that there are no patients waiting at the start of slot j-1 and that k new patients then arrive. The second term is the sum of joint probabilities that there are no more than P patients waiting patients at the start of slot j-1 and that k new patients then arrive for slot j. These P (or fewer) waiting patients will all have been serviced by the P service providers in the prior appointment slot. The third term represents the summation of joint probabilities that some patients remain from the previous slot, such that the number remaining and arriving patients sum to k.

Problem solution technique

Utility function $\hat{U}(\mathbf{S})$ is quasi-concave in the number of appointments scheduled a_{ij} in each appointment slot, so good solutions can be found using a simple heuristic solution methodology. In brief summary, the iterative solution procedure combines gradient search followed by pairwise interchange. In an extensive computational analysis, this procedure provided solutions that were optimal for all problems for which provably optimal solutions could be found. Further details are available in LaGanga and Lawrence (2008).

COMPUTATIONAL RESULTS

To investigate the impact of walk-ins on appointment scheduling, we examined a problem instance that is representative of appointment scheduling problems that we have observed in practice. The clinic scheduling problem studied has 12 appointment slots (N = 12). We assumed that the marginal benefit of serving an additional customer is π = 1.0; the quadratic cost of a patient waiting time d for service is ω = 1.0; and the quadratic cost of clinic overtime is τ = 1.5 per period d, the duration of an appointment. The show rate for scheduled appointments is taken to be 70% (σ = 0.7) Expected demand for clinic services is 36 patients per clinic session, or 3 patients per appointment slot. Given this expected demand, the clinic is staffed with 3 providers (M = 3).

We examined four different clinic arrival scenarios: (1) all patients are walk-ins (λ=3); (2) two-thirds of patients walk-in (λ=2) and one-third are seen by appointment; (3) one-third are walk-ins (λ=1) and two-thirds are seen by appointment; and finally, (4) all patients are seen by appointment. An implicit assumption in this construction is that patients can be shifted between walk-in and appointment traffic, so that clinic demand remains unchanged across scenarios. This assumption allows a fair comparison between the four walk-in and appointment policies, but we acknowledge that it may be difficult to easily exchange walk-in and appointment demand, especially over a short time interval.

Using the solution methodology described in the previous section, appointment schedules and performance statistics were obtained for each of the four scenarios using a common desktop computer. Solutions times ranged from 23 to 75 seconds, using code that had not been optimized for speed. Results are summarized in Figure 1 and Table 2.

Figure 1 illustrates the arrival pattern for each of the four clinic scenarios. Figure 1a (all walk-in patients) shows an expected 3 walk-in patients arriving for each appointment slot. The number of walk-ins arriving in actuality will of course vary since walk-in arrivals are Poisson-distributed and may be more or less than the expected value of 3 patients. Table 2 shows that an all walk-in clinic will see more patients than the other 3 clinics (34.1 patients per clinic session), but does so at the expense of clinic overtime, with expected overtime equivalent to 3.8 extra appointment slots each day. This extensive clinic overtime is costly (−41.6 utility loss) and results in *negative* net clinic utility (−12.2).

Figures 1b and 1c show a mix of walk-in and appointment traffic. In Figure 1b, walk-ins constitute 67% of clinic demand (λ = 2), with appointments constituting the remainder. The optimal appointment schedule significantly overbooks the first appointment slot with 3 appointments, but does not book any appointments for slots 7-12. Such front-loading of an

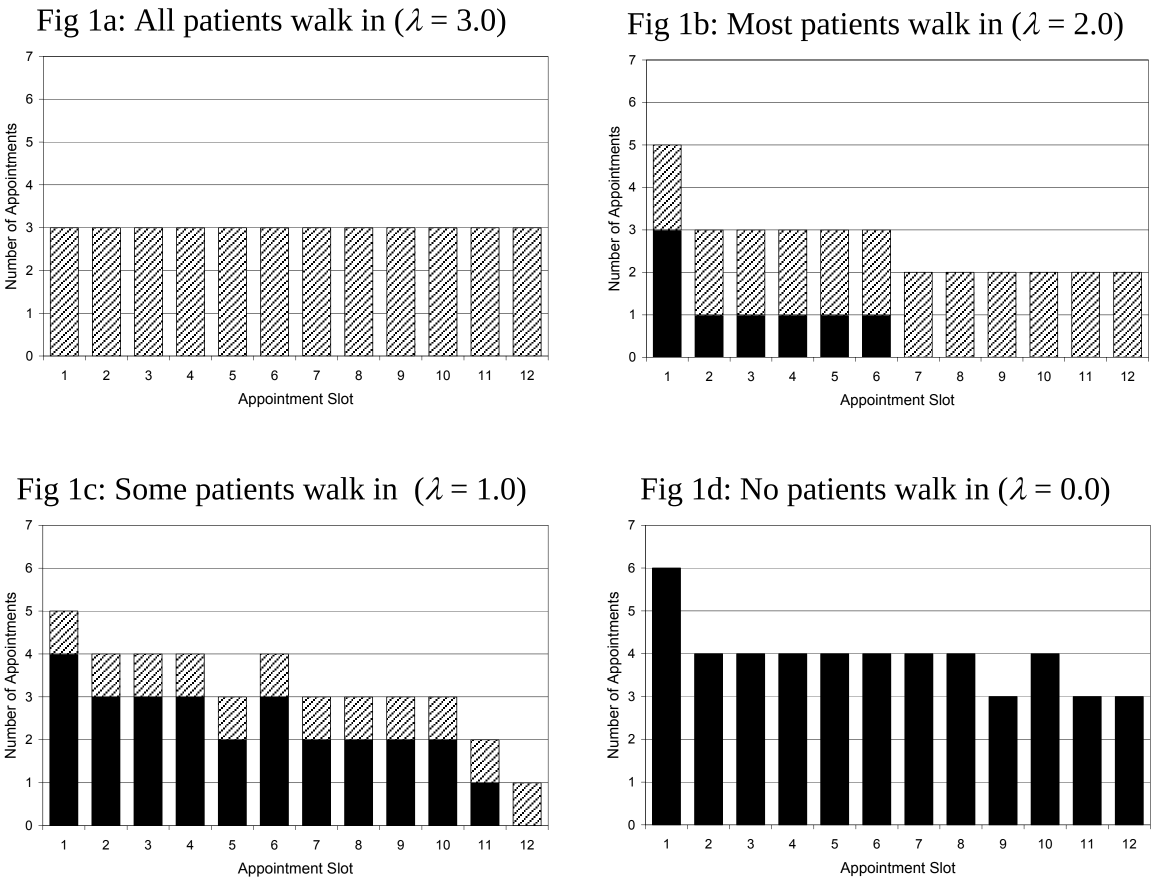

Figure 1 – Scheduled appointments for various patient walk-in rates λ
N=12, P =3, σ=0.7, π=1.0, ω=1.0, τ=1.5
Hatched bars represent expected walk-in patients; solid bars are scheduled appointments

appointment schedule serves to "prime the pump" with waiting patients who are then available if walk-in traffic fails to materialize and if there are appointment no-shows. In this second scenario, Table 2 shows that a total of 29.5 patients are expected to be serviced, provider utilization is a low 80.4%, but patient waiting times and clinic overtime are moderate, so net clinic utility is a positive 24.9.

Table 2 – Clinic performance for various patient walk-in rates λ
N=12, P=3, σ=0.7, π=1.0, ω=1.0, τ=1.5

	Arrival Rate for Walk-in Patients			
Performance Measures	**λ = 3.0**	**λ = 2.0**	**λ = 1.0**	**λ = 0.0**
Patients seen	34.09	29.48	30.84	32.88
Wait time per patient	1.82	1.11	1.14	1.00
Clinic overtime	3.77	0.64	0.32	0.29
Provider utilization	89.3%	80.4%	84.9%	90.6%
Utility gained from patients seen	34.09	29.48	30.84	32.88
Per-patient utility lost from waiting	–4.75	–1.94	–1.96	–1.66
Utility lost from clinic overtime	–41.58	–3.56	–1.75	–1.09
Net Clinic Utility	**–12.24**	**23.98**	**27.13**	**30.12**

Figure 1c shows the optimal appointment schedule when the majority of patients are seen by appointment, but some walk-in traffic is allowed ($\lambda = 1$). Again, the optimal appointment schedule is front-loaded to create a reservoir of waiting patients for the rest of the clinic session, and again, the number of scheduled appointments declines during the day with no appointments scheduled for the last appointment slot. Performance data in Table 2 shows that an expected 30.8 patients will be serviced with provider utilization of 84.9%. Patient waiting times are comparable to the second scenario, but clinic is cut by one-half, resulting in improved net clinic utility (27.1).

Finally, Figure 1d illustrates the optimal appointment schedule for the fourth scenario in which there is no walk-in traffic and all appointments are scheduled. This schedule is again front-loaded with the number of appointments generally declining during the clinic session. An expected 32.9 patients will be seen, expected provider utilization is 90.6%, expected patient waiting times and clinic overtime are the lowest of the four scenarios, which results the largest expected clinic utility (30.1).

DISCUSSION

Several patterns immediately are apparent from the results of this example. First, patient wait times and clinic overtime decline as the fraction of scheduled appointments increases. Second, the number of patients served generally increases with the fraction of scheduled appointments except for the case with all walk-in traffic, which has the most patients-seen of any scenario. Finally, net clinic utility increases with the fraction of scheduled appointments. It appears, in general, that the greater is the fraction of scheduled appointments, the greater is clinic performance.

Theoretical explanations

To better understand this pattern, Figure 2 illustrates the distributions of the expected number of patients waiting at several points (slots 1, 4, 8, and 12) in the course of a clinic session. The tables within Figure 2 show that throughout a clinic session, an all-walk-in schedule ($\lambda = 3$) has the highest variance in the distribution of the number of waiting patients, while an all-appointment schedule ($\lambda = 0$) provides the smallest variance, even though the no-show rate is a significant 70% ($\sigma = 0.70$).

This can be explained by comparing the shapes of the Poisson and binomial distributions that determine walk-in and appointment variability, respectively. The Poisson distribution is unbounded and skewed to the right, allowing the possibility that a large number of patients may arrive for service in any given appointment slot. When this happens, the clinic is overwhelmed with demand, leading to the need for extensive overtime. In contrast, the binomial distribution is bounded on the right to be no more than the number of appointments scheduled, and is skewed to the left. This means that the number of patients arriving can never be more than the number scheduled (and with no-shows, is usually less), which limits the potential need for clinic overtime.

From the perspective of queuing theory, setting capacity equal to expected demand in an all-walk-in clinic is equivalent to setting long-run server utilization to 100%. Well-known queuing results inform us that this will lead to infinitely long queues in the presence of system variability. While a health-care clinic does not operate for sufficient time to reach steady state, queuing results help to explain why the walk-in scenario leads to so much clinic overtime.

Practical implications

From a practical perspective, our results indicate that an all-appointment clinic will provide better clinic performance and better patient service than will any combination of appointment and walk-in patient flow. An all-appointment system will generally see more patients, limit patient waiting times, and reduce clinic overtime better than a system that allows walk-in patients. This result helps to explain why many health care clinics operate using an appointment-only model of demand flow management.

First Appointment Slot
(j=1)

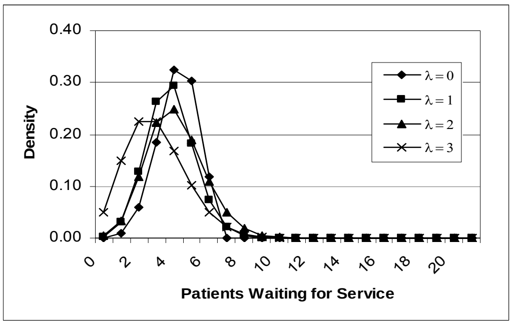

	$\lambda = 0.0$	$\lambda = 1.0$	$\lambda = 2.0$	$\lambda = 3.0$
Mean	4.20	3.80	4.10	3.00
Std Dev	1.12	1.36	1.62	1.73
Skew	-0.36	0.27	0.41	0.58

Fourth Appointment Slot
(j=4)

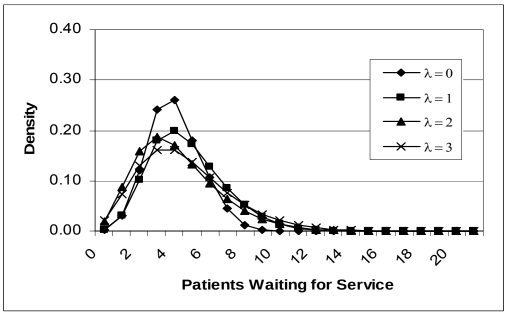

	$\lambda = 0.0$	$\lambda = 1.0$	$\lambda = 2.0$	$\lambda = 3.0$
Mean	4.00	4.70	4.13	4.55
Std Dev	1.52	2.10	2.34	2.59
Skew	0.30	0.63	0.82	0.78

Eighth Appointment Slot
(j=8)

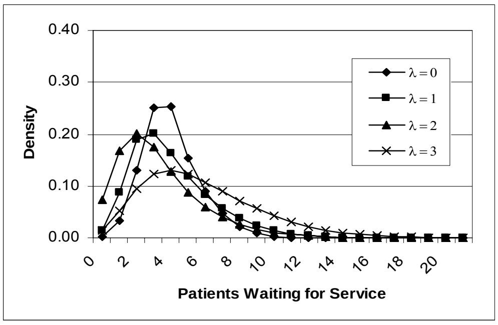

	$\lambda = 0.0$	$\lambda = 1.0$	$\lambda = 2.0$	$\lambda = 3.0$
Mean	4.00	4.02	3.32	5.73
Std Dev	1.67	2.35	2.47	3.37
Skew	0.65	1.02	1.21	0.84

Last Appointment Slot
(j=12)

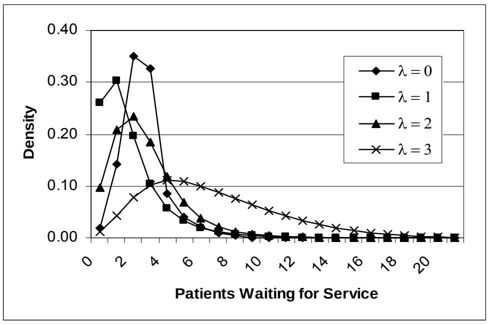

	$\lambda = 0.0$	$\lambda = 1.0$	$\lambda = 2.0$	$\lambda = 3.0$
Mean	2.60	1.74	2.71	6.57
Std Dev	1.29	1.82	2.05	3.91
Skew	1.14	1.79	1.40	0.81

Figure 2 – Comparison of Queue Lengths by Appointment Slot and Mean Walk-ins
(N=12, P =3, σ=0.7, π=1.0, ω=1.0, τ=1.5, quadratic costs)

However, there are many circumstances where walk-in traffic cannot or should not be avoided. Persons who are injured, sick, or in distress cannot wait for treatment and need immediate attention. A clinic that wishes to best serve its client population may therefore find it necessary to see at least some patients on a walk-in basis. Our results suggest how best to do this and how it will effect patient service and the clinic performance.

Figure 2 summarizes the net utility results of the four scenarios examined. For an all-appointment clinic, the figure shows that leaving some capacity open for walk-ins has only a relatively modest impact on clinic performance and net utility. From Table 2, allocating one-third of clinic capacity to walk-in traffic causes patient waiting times to increase by 14%, clinic overtime to increase by 10%, provider productivity to decline by 7%, and net 11% utility to decline by 11%. These outcomes may be far preferable to significantly increasing clinic capacity to buffer against the increased variability that walk-in traffic brings, or to turning away walk-in patients with emergencies, large and small.

For an all-walk-in clinic, Figure 3 is also informative, since it suggests that if a relatively small fraction of clinic capacity can be converted to appointment demand, extraordinarily large

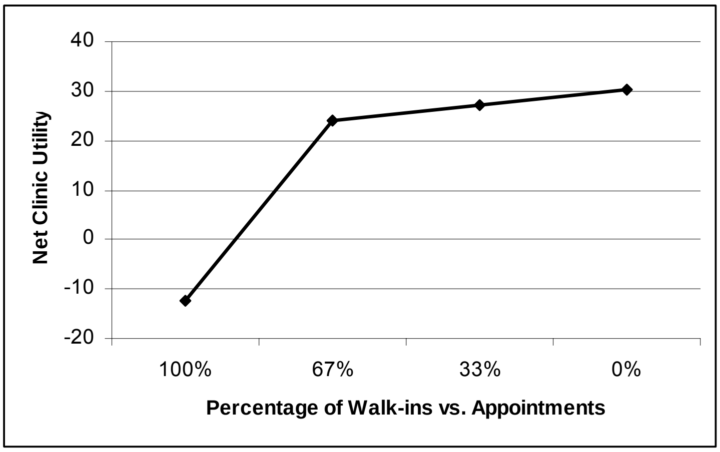

Figure 3 – Clinic performance as a function of walk-in percentage
N=12, P=3, σ=0.7, π=1.0, ω=1.0, τ=1.5

improvements in patient service and clinic performance can be achieved. From our example scenarios, Table 2 shows that allocating one-third of clinic capacity to walk-in traffic would decrease patient waiting times by 39%, decrease clinic overtime by 83%, and would increase net utility from –12.2 to 24.0. However, doing so would cause provider productivity to decline by 11%. These outcomes would probably be preferred over other alternatives such as increasing provider capacity (probably expensive) or reducing demand by turning patients away or raising prices.

Finally, we note that our results can easily be adapted to and offer support for open access scheduling, where patients call in for same-day scheduling of an appointment (Murray and Tantau 2000). If a clinic allows both scheduled appointments and walk-ins (the middle two scenarios above), the day's allotted appointments can be distributed to patients as they call in. If no available appointments meet the schedule of the patient, the patient can be invited to "walk-in" that day and wait for the next available provider, or the patient can make an appointment for the following day. By allowing both walk-ins and an appropriate number of scheduled appointments as demonstrated in this paper, a clinic can balance the needs of patients, the concerns of providers, and interests of the clinic in an open access setting.

CONCLUSION

In this paper we, we have studied the interaction of appointment scheduling and walk-in traffic in the context of yield management in health-care clinics. Using a clinic utility model that balances the benefit of seeing additional patients with the implicit costs of increased patient waiting time and clinic overtime, we developed and computationally analyzed the efficacy of appointments versus walk-ins as a method of demand flow management. Principle conclusions of our study include the following:

- All else equal, managing clinic demand by scheduling appointments provides better patient service and clinic performance than does any model that permits walk-in traffic.
- In circumstances where a degree of walk-in traffic is desired or necessary (for example, to provide service for emergency cases), allowing some walk-in traffic only causes modest declines in patient service and clinic utility. A clinic does not need to sacrifice much performance or endure large costs to provide flexibility for patients by allowing some walk-in traffic.
- For clinics that rely solely on walk-in traffic, converting a relatively small amount of demand from walk-ins to appointments can provide dramatic improvements in patient service and clinic performance. Even a clinic that relies solely on a walk-in model to provide flexibility

and convenience for its patients may find that allowing some appointment scheduling will produce beneficial results.

This paper makes several contributions to our understanding of clinic yield management, and appointment scheduling with walk-in traffic. First, to our knowledge, we are the first to formally examine the implications of mixing of walk-in traffic with appointment scheduling. Second, we have developed a solution methodology that optimizes the number of scheduled appointments in the presence of stochastic walk-in traffic. Third, we have demonstrated the general benefit of appointment scheduling over an all-walk-in model and have characterized the costs of allowing walk-ins. Fourth, we have shown that in cases where some walk-in traffic is required by a clinic; the costs of allowing walk-ins need not be high. Fifth, we have demonstrated that for an all-walk-in clinic, the introduction of even a modest degree of appointment scheduling can provide significant improvements in patient service and clinic performance. And finally, we have shown that our results are easily adapted to and supportive of open-access clinic operations where appointments are made for the same day that a patient calls in.

This paper suggests a number of avenues for further inquiry and future research. To date, we have only investigated a single problem instance of the appointment scheduling problem with walk-ins. While we believe that our results can be generalized, more extensive testing across a wide range of problem characteristics is warranted in order to prove this belief. Another important avenue of research will be to extend our analysis to circumstances where appointment durations are more highly variable. While many clinic settings can be adequately described as having fixed appointment durations, others cannot. For these latter clinics, it will be important to extend our model to accommodate variable service times. Finally, we are working with practitioners in the field to validate the practical efficacy of our results and to further refine our model.

REFERENCES

Dyer, O. (2005), "Sick of getting stood up? No-shows say it's because they need a little respect. *National Review of Medicine*, 2(1), 1/15.

Murray, M., and Tantau, C. (2000), "Same-day appointments: Exploding the access paradigm," *Family Practice Management*, September, 45-47.

LaGanga, L. R.and Lawrence, S. R. (2007), "Clinic overbooking to improve patient access and increase provider productivity," *Decision Sciences*, *38*(2).

LaGanga, L.R., and Lawrence, S.R. (2008), *Yield Management in Health Care Clinics to Improve Patient Service and Clinic Performance*, working paper, Leeds School of Business, University of Colorado, Boulder CO. Available online at http://Leeds.colorado.edu/overbooking.

Lowes, R. (2005), Practice pointers: How to handle no-shows. *Medical Economics*, 82(8), 62-65.

Maister, D. (1984), *The Psychology of Waiting Lines.* HBS Teaching Note #9-684-064. Harvard Business School Press.

Ross, S.M. (1985), *Introduction to Probability Models.* Orlando FL: Academic Press.

Soriano, A. (1966), Comparison of two scheduling systems. *Operations Research*, 14, 388-397.