Waiting time information services: what are the implications of waiting list behaviour for their design?

DAVID CROMWELL AND DAVID GRIFFITHS

David A. Cromwell is Senior Research Fellow at the Centre for Health Service Development, University of Wollongong, Wollongong, New South Wales. David A. Griffiths is Professor at the School of Mathematics and Applied Statistics, University of Wollongong.

Abstract

In some countries, patients requiring elective surgery can access comparative waiting time information for various surgical units. What someone can deduce from this information will depend upon how the statistics are derived, and how waiting lists behave. However, empirical analyses of waiting list behaviour are scarce. This study analysed three years of waiting list data collected at one hospital in Sydney, Australia. The results highlight various issues that raise questions about using particular waiting time statistics to make inferences about patient waiting times. In particular, the results highlight the considerable variation in behaviour that can exist between surgeons in the same specialty, and that can occur over time.

Assumptions underpinning the dissemination of waiting time statistics

In some countries with a publicly funded hospital sector, governments have begun posting waiting time statistics for elective surgery on their web-sites. The general aim of these information services is to assist patients and general practitioners (GPs) to decide on where the patient should be referred. By so doing, it is believed that patients can avoid waiting longer than necessary, and that imbalances in referrals will be reduced, although there is currently little published evidence that supports this view.

A review of six waiting time information services highlighted some statistical issues that might affect the accuracy with which patients can make inferences about their likely waiting time (Cromwell et al. 2002). First, the web-sites derived their statistics from different types of data, using either data on admitted patients (throughput data) or data on patients on the waiting list (census data). In many circumstances, throughput data statistics will provide the more accurate information about expected waiting times (Don et al. 1987), but these statistics may still be inaccurate, such as when levels of surgical activity change over time.

Second, the web-sites presented statistics based on different levels of aggregation (e.g. by specialty or surgeon, over 3 or 12 months). Low levels of aggregation may produce unreliable estimates due to small sample sizes. Yet high levels of aggregation can hide problems of particular units, making the statistics unresponsive. Users also risk making an ecologically biased inference (Morganstern 1998) about a patient's likely waiting time. For example, this will arise from specialty-level statistics if the waiting times of patients on the lists of individual surgeons vary substantially within a specialty. In summary, assumptions are made about waiting list behaviour that may not reflect what happens in reality.

Although various claims have been made about waiting list statistics (Don et al. 1987; Nicholl 1988; Mordue & Kirkup 1989), there are currently no empirical studies that examine how waiting list behaviour might affect the interpretation of different waiting time statistics. Consequently, a study was undertaken which examined waiting list behaviour and patterns of elective surgical activity at one hospital over a three year period, with a view to determining what statistical issues should concern users of waiting time information services. The analysis focussed on waiting list behaviour among patients assigned to the least urgent waiting list category because patients with non-urgent conditions are likely to wait many months for admission, and are likely to make most use of the waiting time information services.

Data collection

Waiting list data were collected from a teaching hospital in Sydney, Australia. Data could be extracted in two forms; one gave information about all patients removed from the list, while the other provided data about patients on the waiting list at a specified date. To provide information on activity between 1 July 1995 and 30 June 1998, de-identified data were extracted on all patients admitted or removed from the list during this period, together with the patients still waiting on 30/6/98. The period coincided partially with the State-wide waiting list reduction program that ran from May 1995 until 31 December 1995. During the program, the number of patients waiting in NSW dropped 56% from 44707 to 19589 (Shiraev & McGarry 1996).

Each patient record included the identifier of the treating doctor, a code for the intended procedure, the dates of listing and removal, and a code indicating the type of removal (ie, elective admission, emergency admission, or the reason for removal without admission). Other fields gave the final urgency category assigned to the patient, the number of re-categorisations, and the date of the last re-categorisation. Each record also included the final listing status assigned to the patient (ie, whether they were "ready for care", deferred or staged), the number of days the patient was listed as not ready for care, the number of times the listing status had changed, and the date of the last change.

The database contained the records for 46 surgeons, each of whom had been active for the full three years. There were 27,827 records in total, and there were at least two surgeons in each of the 10 surgical specialties covered. Each surgeon appeared to operate only a single waiting list. The urgency categories used by surgeons changed on 1 July 1997 (see Table 1). The new NSW category U7, required categories U3 and U4 to be renamed U8 and U9 respectively. The urgency data were recoded to be consistent with the current classification.

Table 1: Waiting list urgency categories and maximum desirable waiting time

Category code*	Used in years 1995/6 + 1996/7	Used in 1997/98
1	Admit, preferably within 7 days	Admit, preferably within 7 days
2	Admit, preferably within 30 days	Admit, preferably within 30 days
3	Admit as soon as possible	not used
4	Patient not ready for care	not used
7	not used	Admit, preferably within 90 days
8	not used	Admit as soon as possible
9	not used	Patient not ready for care

* The letter U has been added as a prefix to the category labels in the text of this article so that they are not misread as a derived statistic

While the database provided information on activity levels and the waiting times of admitted patients, the number waiting within a given urgency category on a specific date could not be directly derived. This was because the records did not include all the data required to reconstruct the sequence of categorisation for the 2608 patients whose urgency category changed during their time on the list. A change in urgency category can mean a change in listing status. Consequently, it was necessary to impute values for the missing fields. Correct values were derived using a variety of techniques. First, additional patient data were collected for 10 other census dates, which enabled some missing category values and dates of recategorisation to be derived. Second, the relationship between the urgency categories and listing status meant that, in specific circumstances, some missing urgency values and dates could be deduced logically. Others could be ignored as the period of unknown urgency occurred before July 1995.

When the missing urgency category could not be derived, values were inferred using heuristic rules that were based on known recategorisation sequences, and the length of total waiting time. If the date of recategorisation was missing, and one of the census dates occurred during the two periods, the missing date was set to be the census date. When a census date did not occur during the two periods, or a sequence of dates were missing, the first unknown phase of recategorisation was defined to last the whole period of unknown urgency. All other phases in the sequence affected by missing date(s) had zero duration. The outcome of this process is summarised in Table 2. Actual values were found for 1972 (63%) of the missing urgency categories.

Table 2: Outcome of process to find urgency categories and dates of recategorisation not contained in the initial database

	Number of recategorisations			
	0	1	2	3 or more
Number of patients	25219	2249	297	62
Phases missing urgency categories				
Number affected	0	2249	594	221
Number derived from data	0	1139	466	132
Number heuristically inferred	0	910	128	89
Phases missing dates of recategorisation				
Number affected	0	0	297	159
Number derived from data	0	0	275	105
Number heuristically inferred	0	0	22	54

Cross-sectional analysis of waiting list behaviour and surgeon activity

The basic assumption behind the waiting time information services is that waiting times for different surgeons are different. The cross-sectional distribution of waiting times of admitted patients (who did not change urgency category) within each urgency category confirms that this is true, at least for patients admitted from category U8. Twenty-nine surgeons had a median wait of less than 90 days, with 12 having a median less than 30 days. Nine surgeons had a median wait of over 180 days, the maximum being 324 days. For the high urgency categories, variation was clearly limited by their maximum desirable waiting time and this restricted the degree of difference between surgeons. For category U1, all surgeons except one had a median waiting time of under 4 days. For category U2, all but six surgeons had a median waiting time below 30 days.

The range of surgeons' median waiting times within each specialty is shown in Table 3. It shows that there can be considerable diversity among surgeons, and this extended to other aspects of their practice. The size of the waiting list, and differences in the rates of activity, can be seen from the cross-sectional figures in Tables 3 and 4. That there can be large variations between surgeons in the same specialty highlights the danger of making an invalid inference about surgeon-level behaviour from specialty-level statistics. In addition, the proportions of patients allocated to each urgency category varied between surgeons, although there were some similarities among surgeons in the same specialty, and the proportions changed over time. This change over time was not simply due to the introduction of category U7. The proportion of patients allocated on addition to the list to category U2 (admit within 30 days) increased for 36 of the 46 surgeons over the 3-year period.

Table 3: Range of median waiting times and average size of waiting lists for surgeons within specialties (minimum, maximum)

Specialty	Median waiting time (days)		Average census	
	Min	Max	Min	Max
Cardiac Surgery	16	19	0.5	5.4
ENT	90	189	30.6	92.0
General surgery	7	120	1.6	104.2
Gynaecology	18	105	1.1	44.3
Neuro-surgery	20	22	3.5	4.0
Ophthalmology	243	324	52.3	66.4
Orthopaedics	77	253	16.3	102.9
Plastic surgery	8	58	1.2	20.4
Urology	35	271	4.8	77.6
Vascular surgery	48	92	48.0	66.6

Table 4: Range of surgeon activity within specialties (minimum, maximum)

Specialty	Average monthly rate of admission for U8 patients		Average proportion of U8 patients of total admissions		Average monthly rate of removal for U8 patients	
	Min	Max	Min	Max	Min	Max
Cardiac Surgery	0.5	5.4	4%	35%	13%	18%
ENT	6.1	9.4	54%	84%	10%	32%
General surgery	1.3	26.3	21%	70%	4%	24%
Gynaecology	1.2	7.1	17%	83%	4%	39%
Neuro-surgery	4.1	4.3	37%	62%	2%	5%
Ophthalmology	3.4	4.2	66%	77%	30%	40%
Orthopaedics	1.9	9.1	49%	72%	23%	41%
Plastic surgery	1.7	4.8	13%	43%	3%	22%
Urology	2.0	6.5	16%	16%	10%	41%
Vascular surgery	4.1	5.8	28%	29%	28%	30%

The cross-sectional figures highlight other issues. First, a low admission rate will mean throughput statistics may be based on small samples, which could affect the smoothness of a sequence of statistics. Smoothness is an important characteristic in the context of waiting time information services because an erratic sequence would not lead anyone to have much confidence in the accuracy of the waiting time predictions. The likelihood of this being a problem is increased by the fact that patients are not generally admitted on a first come, first served basis. This will generally increase the variation in waiting times among patients.

An analysis was performed that examined the size of throughput data samples. The analysis assumed disseminated statistics were calculated for each surgeon, and were updated every month. When the statistics were based on data aggregated over three consecutive months, the minimum size of the samples for 28 surgeons was either 1 or 2 patients. Overall, 9% of all statistics were based on samples of this size. Sample sizes could be increased by aggregating data over the last three months in which at least one person was admitted. Using this rule, 15 surgeons had a minimum sample size of 3 patients, while the minimum for 25 surgeons was at least five. Moreover, this simple rule had a noticeable effect on the smoothness of the time series of monthly average (mean) waiting times. For 13 surgeons, the maximum difference between consecutive averages was over 90 days when derived using data from three consecutive months. When derived using data aggregated by the "at least one admission" rule, the difference decreased for 10 of these surgeons, in 8 cases by over 50 days. It made no difference in the other three cases. Overall, the maximum difference decreased in 22 cases, did not change in 16 cases and increased in 8 cases.

The issue of sample size also arises if data are aggregated by procedure as well as surgeon and urgency category. Indeed, analysis of the data on category U8 patients found that many procedures were performed infrequently, and their incidence within separate urgency categories was reduced further when category U7 was introduced. Only 13 surgeons performed any procedure on category U8 patients 30 or more times per year prior to the introduction of category U7, and only eight surgeons did so afterwards. This casts doubt on whether it would be possible to get sufficiently large samples for surgeon level procedure-based statistics to be reliable; aggregating data over periods of more that a year is unrealistic due to the potential for change in behaviour.

The final statistic included in Table 4 shows the differences between surgeons in the number of patients assigned to category U8 who were removed without admission. For some surgeons, over 25% patients in category U8 that joined the list were removed, although a significant fraction of these received treatment (either as an emergency admission or elsewhere). Such differences again have implications for the interpretation of waiting time statistics, particularly those based on census data. Census data containing a high number of patients who are eventually removed may bias waiting time statistics, and because such patients cannot usually be identified ahead of time, it is not clear how this effect could be counteracted.

Longitudinal analysis

Changes over time in the waiting list statistics of U8 patients were analysed using two time series, both being derived on a monthly basis. The first time series was the number of patients on the waiting list (the census), the number being counted at midnight on the last day of each month. The second series was the average waiting time of U8 patients admitted during the month and who had not changed urgency category. Time series were created for all surgeons, and were grouped according the broad types of behaviour.

Distinctions were made between seven types of census time series, based on the size of the waiting list and, for the larger waiting lists, commonly appearing patterns. These types classify surgeons whose census:

(C1) was typically less than 5 patients, and was often zero;

(C2) was typically less than 10 patients, but rarely zero;

(C3) fluctuated between 10 and 40 patients;

(C4) was originally high (>40 cases) but dropped greatly over the last year (as patients were reassigned to category U7);

(C5) was typically high and stable;

(C6) showed an increasing trend over the data collection period;

(C7) showed a significant dip and rise over the first 18 months of the data collection period, due to the waiting list reduction initiative.

The patterns exhibited by the time series of average waiting time were grouped into five classes, using as classification criteria the level of wait and, for longer average waiting times, type of pattern. These groups contained surgeons whose monthly average waiting time:

(W1) was typically low, and fairly stable;

(W2) appeared stationary but fluctuated;

(W3) increased from a low initial level;

(W4) increased over time, but that began at a higher level, and fluctuated more, than the times series of surgeons in group W3;

(W5) was high, and moved erratically.

Examples of each type of behaviour are shown for the time series of census and average waiting times in Figures 1 and 2. The gaps in the graphs of the average waiting times correspond to months in which no U8 patients were admitted.

Figure 1: Examples of the grouped census time series by surgeon

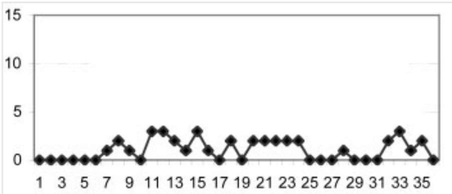

C1: Surgeons whose census was very low (typically <5)
Dr001, Dr012, **Dr026, Dr038, Dr040

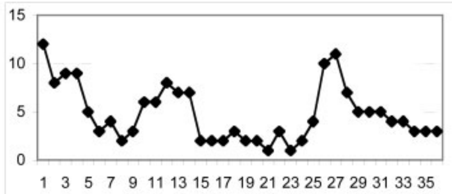

C2: Surgeons whose census was low (typically <10)
Dr002, Dr005, Dr010, Dr011, Dr016,Dr023, Dr027
Dr028, **Dr042, Dr043

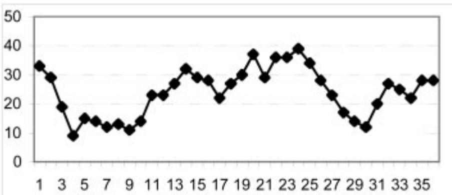

C3: Surgeons whose census ranged between 10 and 40
Dr004, **Dr009, Dr014, Dr018, Dr021, Dr022, Dr025
Dr036, Dr039, Dr041

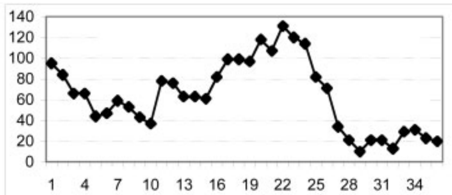

C4: Surgeons whose census dropped greatly in last year
**Dr008, Dr013, Dr020

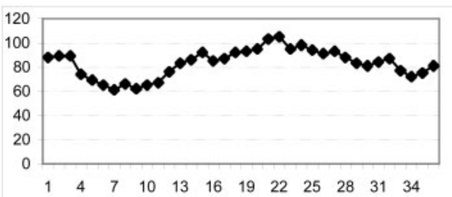

C5: Surgeons whose census was high and stable
Dr029, **Dr033, Dr037, Dr044, Dr046

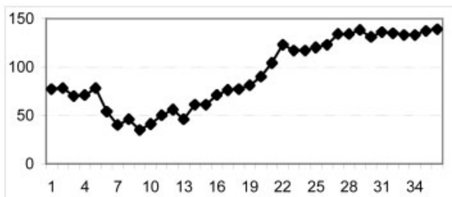

C6: Surgeons whose census showed an positive trend
**Dr003, Dr006, Dr007, Dr015, Dr017, Dr024, Dr030
Dr032, Dr034, Dr035

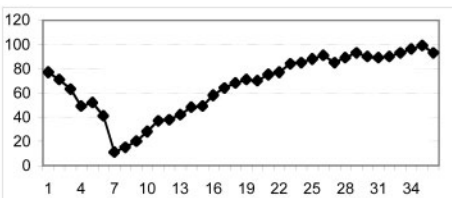

C7: Surgeons whose census had a large dip in first year
Dr019, **Dr031, Dr045

Note:
Data in a graph relates to the surgeon marked with (**).
x-axis: Time in months
y-axis: Number of patients waiting
Census data collected on a monthly basis,
starting on midnight 30 June 95.

Figure 2: Examples of time series of average waiting times by surgeon

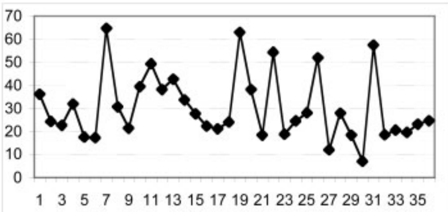

W1: Surgeons where waiting time was low and stationary
Dr001, Dr002, Dr005, Dr010, Dr013, **Dr023, Dr026
Dr027, Dr028, Dr038, Dr040

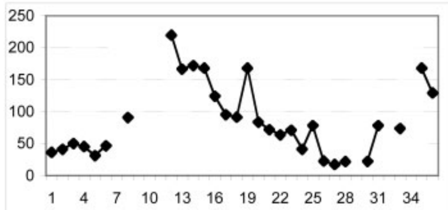

W2: Surgeons whose waiting time fluctuated
Dr004, Dr006, Dr008, Dr011, Dr016, Dr018, Dr022
**Dr041, Dr043

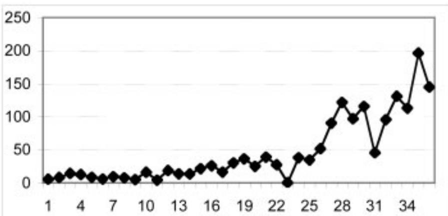

W3: Surgeons whose waiting times rose from low levels
Dr007, Dr009, Dr012, **Dr014, Dr015, Dr017, Dr019
Dr021, Dr025, Dr036, Dr039, Dr042

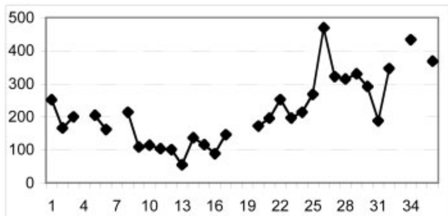

W4: Surgeons whose waiting times became excessive
**Dr003, Dr020, Dr024, Dr029, Dr030, Dr031, Dr032
Dr037, Dr045

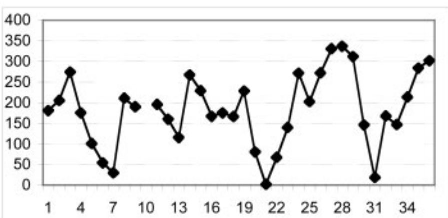

W5: Surgeons whose average waiting time fluctuated greatly
Dr033, Dr034, **Dr035, Dr044, Dr046

Note:
Data in a graph relates to the surgeon marked with (**).
x-axis: time in months
y-axis: waiting time in days

Several of the longitudinal features are important with respect to using statistics presented by waiting time information services to make inferences about likely waiting times. First, the groups confirm the extent to which surgeons in the same specialty can differ. The four specialties containing four or more surgeons (general surgery, gynaecology, orthopaedics and urology) had surgeons in at least three of the census groups and at least three of the average waiting time groups. This is further evidence that aggregating data at a specialty level may produce statistics that will lead to invalid inferences about how long someone might expect to wait. Second, Figure 2 shows how admitting patients in a way that greatly deviates from a first come, first served policy can affect waiting time statistics. This is the principal reason for the erratic behaviour in the time series of average waiting time, which is especially noticeable in group W5. Third, it shows how quickly the underlying level of waiting time can change due to an imbalance between the rates at which patients join and leave the list. That some series are non-stationary implies that both census and throughput waiting time statistics will be biased as the simple ways in which these are derived do not take this non-stationarity into account. This bias may be large in situations where the waiting time (and so the forecast lead time) is many months.

The second and third features also highlight another dilemma facing designers of waiting time information services. Aggregating data over long periods of time will increase sample size and this will generally reduce the amount by which the average waiting time is influenced by patients with unusual waiting times. However, doing so increases the risk of the statistic being unresponsive to change in waiting time behaviour. Finally, it is worth noting that, at a surgeon level, the relationship between average waiting times and the size of the waiting list is not simple. High average waiting times do not always correspond to long waiting lists. Moreover, the relative movements of the series over time can differ. An example of this is surgeon Dr020. This surgeon's census dropped greatly during the last year, while the average waiting time increased.

The effect of changing urgency category or list status on waiting time

It is possible to define the waiting time of patients who change urgency category and/or listing status in various ways. The definition is not consistent between countries and, in some cases, has changed over time. Although the number of patients who change category is a small proportion of all U8 admissions, their effect on disseminated statistics may not be negligible if statistics are derived from small samples and the waiting times of patients who change category are systematically different.

An analysis was undertaken comparing the waiting times of admitted U8 patients who changed and did not change urgency category or listing status. A complete description of this is contained in Cromwell (2002) and only a brief description is given here. The initial step was to examine the waiting times of patients who were listed in category U8. This used the cohort of patients who joined in the same month and who had all left the waiting list during the data collection period; patients from months with incomplete data were excluded. Of the 8043 admitted patients within this cohort, 7141 had not changed category, 786 had changed once, while the remaining 116 had changed category on two or more occasions. Each patient who had changed category was matched to a patient who had joined the same surgeon's list at the same time, and the difference in their waiting times was derived. For patients who changed category, waiting time was defined as the time spent listed as "ready for care". The average differences between the observed sequences of categories are shown in Table 5. Although the differences were over 20 days for three types of sequence, the difference was only statistically significant in one case. This was because of the large variation in differences, due mainly to patients not being admitted on a first come, first-served basis.

Table 5: Summary of differences in waiting time between patients who changed urgency (cases) and the patients who did not (controls)

Category sequence	Average difference (days)	Standard error (days)	Case wait < control	Case wait = control	Case wait > control	Sign Test statistic
U8-U9	-21.6	4.1	487	14	243	8.47
U8-other	-79.0	27.8	26	0	16	1.70
U8-U9-U8	33.0	22.3	39	0	50	1.06
Other U8	1.6	41.9	15	0	12	0.77

The second part of the analysis compared the waiting times of individuals who changed category with those who did not when patients were grouped by their final urgency category. The waiting times of patients who had changed category were compared with the median waiting time of patients who had not changed category but who had been admitted by the same surgeon. The median wait was derived from those patients who had been admitted either in the same month, or the preceding or following months. If waiting times for each type of patient were equivalent, then the waiting time of the patients who changed category should be evenly distributed above and below the median. The analysis used the same definition of waiting time as above.

The analysis revealed substantial and statistically significant differences in waiting times between U8 patients who had and had not changed urgency (see Table 6). On average, the waiting times of those patients who changed category once were less than the reference median, while the waiting times of patients who changed category two or more times exceeded the reference median by an average of 19 days. The proportion of patients who waited less than the median was statistically different from 50% for the "two categories" group (Sign Test, 26/41, p=0.030). Similarly, the proportion of patients who waited more than the median was statistically different from 50% for the "three or more categories" group (Sign Test, 97/169, p=0.032).

Although not conclusive, the analysis suggests that patients who change category have a different waiting experience to those who do not. This raises questions about whether patients who change category should be included the samples from which statistics presented by waiting time information services are derived, and whether such services need to warn patients that the statistics will be most accurate for those patients who do not change urgency or listing status.

Table 6: Difference between the reference median and the waiting times for patients who changed urgency

Sequence group	Patients who changed category					
	No. of cases	Average wait (days)	Average of reference medians (days)	Average difference (days)	Std error of difference (days)	No. of cases who waited less than the median
Two categories	42	81	168	-86	25.6	26
Three or more categories	170	192	173	-19	11.0	72

Conclusion

The statistics presented by waiting time information services reveal how patient waiting times differ at various surgical units. There is nothing statistically incorrect in this if an inference is not drawn from these statistics. Assuming the data contain minimal errors, they will accurately describe the experience of patients over a period of time. But this is a managerial perspective, and it is not how patients or GPs will use the services. They will (and are encouraged to) make inferences about what the differences will mean for someone who is about to be referred.

This study provides some insight into the dangers of using waiting time statistics in this way. The first issue it raises concerns the level at which data are aggregated. Some services have presented data for non-urgent patients aggregated by specialty (Cromwell et al. 2002), yet the observed differences between surgeons in relation to activity, the waiting list census, and average waiting times suggest that inferences about surgeon level behaviour from such statistics are likely to be biased. Other services have presented statistics based on throughput data aggregated by surgeon and procedure. Low levels of activity suggest that, if data are collected over short intervals of time (three consecutive months), the statistics may be unreliable due to small sample sizes, but if data are collected over longer intervals (12 months), the statistics may be biased due to changes in waiting times over time. Changes in behaviour over time will also bias statistics derived from census data.

Several other practical issues were identified. First, patients who change urgency category or listing status appear to have different waiting times than those patients who do not. This suggests that these "unusual" patients should be excluded from statistics presented by the information services. Second, the high rate of removal without admission has implications for statistics derived from census data. This also has implications for another type of waiting list statistic, namely, the clearance time. This was initially defined by Cottrell (1980) as the census divided by the average admission rate. However, these results suggest that, to produce an accurate estimate of how long it will take to clear all patients currently on a waiting list, it is necessary to include a term for removals.

The study also provides some general insight into waiting list behaviour at the level of a surgeon, at the level that waiting lists are often managed. Two aspects of behaviour are worth emphasising. First, the impact of patients not being admitted on a first-come, first-served basis was clearly visible in some time series of average waiting time statistics. That the series can fluctuate so violently suggests that either some adjustment process or a more sophisticated sampling process might be warranted. This applies to the use of such statistics for managerial purposes as well as to using them to assist patient decisions. Second, the analysis casts doubt on whether relationships between factors observed at high levels of aggregation hold at lower levels. In particular, it suggests the positive association between waiting list length and waiting times, strongly asserted by Yates (1987), and supported by an analysis of regional level data (Harvey 1993) is not a general relationship.

The study does suffer from several limitations that may make the results atypical. Chief amongst these were the waiting list reduction program that ran in 1995 between May and December, and the introduction of the new urgency categorisation in 1997. These both contributed to changes over time in waiting list behaviour and activity levels, and consequently the observed longitudinal changes may be greater than would be expected. However, another effect of the waiting list reduction program would have been to reduce surgeon differences, at least initially. Consequently, it is believed that such differences would be a general feature at most hospitals.

The other main limitation was the need to infer urgency categories and dates for some patients who changed urgency category. Nonetheless, the inaccuracies in the derived waiting list census are not thought to affect the conclusions of the study. A large proportion of sequences with missing categories matched a frequently occurring sequence (ie, U8 then U9 then U8). Moreover, analysis of how the census might change if all inferred values were wrong found that the calculated census would differ from its real value by at most 10%. Twenty-nine of the 46 surgeons had a maximum error of three or fewer patients. The surgeons that were most affected (Dr007 and Dr008) had maximum potential errors of 15 and 16 patients respectively. The next highest were four surgeons that shared a maximum potential error of seven patients.

References

Cottrell KM 1980, 'Waiting lists: Some problems of definition and a relative measure of waiting time', *Hospital and Health Services Review*, vol 76, pp 265-269.

Cromwell DA 2002, *The use of information in the management of elective surgery waiting lists. PhD thesis*. University of Wollongong, Wollongong.

Cromwell DA, Griffiths DA, Kreis IA 2002, 'Surgery dot.com: the quality of information disseminated by web-based waiting time information services,, *Medical Journal of Australia*, 177, in press.

Don B, Goldacre MJ, Lee A 1987, 'Waiting list statistics 3 : comparison of two measures of waiting times', *BMJ*, vol 295, pp 1247-1248.

Harvey I 1993, 'And so to bed: access to inpatient services', in S Frankel, R West (eds) *Rationing and Rationality in the National Health Service: the persistence of waiting lists*, MacMillan Press, Basingstoke.

Mordue A, Kirkup B 1989, 'An appraisal of waiting list problems', *Health Trends*, vol 21, pp 110-113.

Morganstern H 1998, 'Ecologic studies', in KJ Rothman, S Greenland (eds) *Modern Epidemiology. 2nd ed*. Lippincott-Raven, Philadelphia.

Nicholl J 1988, 'Comparison of two measures of waiting times', *BMJ* vol 296, pp 65.

Shiraev N, McGarry J. 1996, 'Waiting list reduction program: results to November 1995', *NSW Public Health Bulletin*, vol 6(12), pp 147-150.

Yates J 1987, *Why are we waiting? an analysis of hospital waiting lists*, Oxford University Press, Oxford.