INVESTIGATION OF QUEUEING PARAMETERS

FOR

EVALUATING

BOSTON HOUSING AUTHORITY MAINTNENANCE OPERATIONS

by

Theodore James Langton

B.A. Hampshire College
Amherst, Massachusetts
(1980)

Submitted to the Department of
Urban Studies and Planning
in Partial Fulfillment of the
Requirements of the
Degree of

MASTER OF CITY PLANNING

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 1984

Signature of Author ______________________________
Department of Urban Studies and Planning, 23 May 1984

Certified by ______________________________
Aaron Fleisher
Thesis Supervisor

Accepted by ______________________________
Ralph Gakenheimer, Chairman
Departmental Graduate Committee

"Investigation of Queueing Parameters
for
Evaluating
Boston Housing Authority Maintenance Operations"

Masters Thesis in City Planning
Theodore Langton
Masachusetts Institute of Technology
June 1984

ABSTRACT

This paper analyses current problems in evaluating BHA maintenance operations using recorded data, offers suggestions for designing monthly performance measures, and discusses criteria for evaluating proposed operations policies. On the assumption that accurate projections of demand are preferable to strict prescheduling as a basis for designing maintenance systems that are responsive to tenants needs, I have investigated the conditions necessary for using stochastic queueing models to project the consequences of alternative operating policies.

The analysis uses Consistent System statistical programs developed by the Laboratory of Architecture & Planning at MIT and run on the Multics operating system at MIT's Information Processing Services. In Chapter II the relationship of data structure to maintenance operations is described and variables are chosen from October 1983 work order data provided by the BHA. I then use techniques from linear regression, analysis of variance, goodness of fit tests and queueing theory in Chapter III to define the behavior of work order arrival processes. A similar analysis is presented for service times in Chapter IV.

The results suggest that calls for service are not Poisson distributed, although the limited sample size makes it difficult to draw definitive conclusions. I also test the sensitivity of various methods for comparing observed and hypothesized probability distributions at different arrival rates and sample sizes. Because it is likely that systematic maintenance problems are causing work orders to be generated non-randomly, a method is outlined for identifying building systems failures from task code data to be recorded by a modified work order processing system. The extent to which work orders are generated in a Poisson manner can then be used as one measure of how well buildings are being maintained.

Chapter IV provides reasonable evidence that service times are not exponentially distributed and suggests that queue interdependency may explain observed service time distributions. Poisson-based queueing models therefore would not currently provide acceptable accuracy for use in evaluating proposed operating policies.

I then use more generally applicable relationships from queueing theory in Chapter V to analyze turn-around times and queue lengths, and to compare priority policies. First, a regression model indicates that service priority is given to recent work orders rather than to emergencies per se. The tendency to delay the service of older work orders creates backlogs which are not fully reflected in mean

turn-around times. In addition, since several work orders may be generated for a single repair job, it is difficult to estimate the number of tenants in queue. The resulting ambiguities are not primarily due to data structure, however, but to operating problems. Although the number of servers appears adequate, inefficient priority-of-service policies and interdependent queues seriously hinder the system's responsiveness to demand. Therefore, suggestions are made for reducing interdependency and a method is described for comparing priority policies with respect to total expected waiting times.

Acknowledgements

I would like to warmly thank Aaron Fleisher, Professor of Urban Studies & Planning, and Statistics Instructor Ed Kaplan at MIT for their time, teaching, suggestions, criticisms, encouragement and inspiration. I am also grateful to Gwen Friend, Ricardo Duran and the Boston Housing Authority for the time and effort they generously donated, and for the use of their information. Then there are my parents, whom I can never thank enough, and Cynthia Lagier - who did the graphics, who understands what the Poisson distribution is all about, and to whom I was a phantom for several months.

"The first intimation that things were getting out of hand came one early-fall evening in the late nineteen forties. What happened, simply, was that between seven and nine o'clock on that evening the Triborough Bridge had the heaviest concentration of outbound traffic in its entire history...

The bridge personnel, at any rate, was caught entirely unprepared. A main artery of traffic, like the Triborough, operates under fairly predictable conditions. Motor travel, like most other large-scale human activities, obeys the Law of Averages - that great, ancient rule that states that the actions of people in the mass will always follow consistent patterns - and on the basis of past experience it had always been possible to foretell, almost to the last digit, the number of cars that would cross the bridge at any given hour of the day or night. In this case, though, all rules were broken...

The incident was unusual enough to make all the front pages next morning, and because of this many similar events, which might otherwise have gone unnoticed, received attention... It was apparent at last that something decidedly strange was happening. Lunchroom owners noted that increasingly their patrons were developing a habit of making runs on specific items; one day it would be the roast shoulder of veal with pan gravy that was ordered almost exclusively, while the next everyone would be taking the Vienna loaf and the roast veal went begging. A man who ran a small notions store in Bayside revealed that over a period of 4 days, 274 successive customers had entered his shop and asked for a spool of pink thread...

At this juncture it was inevitable that Congress should be called on for action... In the course of the committee's investigations it had been discovered, to everyone's dismay, that the Law of Averages had never been incorporated into the body of federal jurisprudence, and though the upholders of States' Rights rebelled violently, the oversight was at once corrected, both by Constitutional amendment and by a law - the Hills-Slooper Act - implementing it. According to the act, people were required to be average, and, as the simplest way of assuring it, they were divided alphabetically and their permissible activities catalogued accordingly. Thus, by the plan, a person whose name began with "G," "N," or "U," for example, could attend the theater only on Tuesdays, and he could go to baseball games only on Thursdays, whereas his visits to a haberdashery were confined to the hours between ten o'clock and noon on Mondays."

- Robert M. Coates, "The Law", 1947

CONTENTS

CHAPTER I

INTRODUCTION

1) Background & Purpose

Since 1980, when a long history of severe funding and management problems culminated in a 33% vacancy rate in Boston public housing, the Boston Housing Authority (BHA) has been in receivership. For decades, one of the Authority's major problems has been the management of maintenance operations. Consequently, much of the city's hopes for getting the BHA out of receivership rest on the extent to which maintenance operations can be improved.

A system based on tenant-initiated work order requests has been in use for some time, and recently work order data have been recorded on computer tapes so that summary profiles of maintenance operations can be generated. Largely because of the heavy demands already placed upon managers and operations staff, however, no systematic attempts have been made to analyze work order data, despite the considerable monthly effort required to record and store them. Although performance measures developed from these summaries might provide relatively unambiguous yardsticks for evaluation and scheduling, the form of the data and their reliability have impeded the creation of such performance measures.

It would be useful then to investigate tools and procedures that might help feedback from past maintenance operations to inform current practice and to project some likely waiting time consequences of alternative operating policies and scheduling methods. First, we would like to determine the feasibility of projectively evaluating changes in priority of service policies. Before implementing new policies, development managers and operations staff should be able to make informed decisions based on the improvements expected from these changes under a variety of conditions. Such changes might be projected by a set of performance measures estimated from queueing models. The degree to which improvements are expected from a given policy would be assessed by comparing these projections to a similar set of "observed" performance measures calculated statistically from the previous month's work order data. Over time, this would allow projections to be tested and refined, and day to day scheduling operations better anticipated.

Before such models can be constructed, however, probability distributions for the number of arrivals and service time completions in a given time period must be estimated from recorded data, and the reliability and structure of these data must be assessed. Another purpose, therefore, is to suggest any changes in data structure that could lead to more useful performance measures being culled from work order data on a monthly basis. In addition to providing information which can be compared with queueing models, these measures must also be used to update inputs for such models. Observed performance measures should be designed to be quickly and easily extracted, and the structure of the information should help us to clearly interpret changes in actual operations. In this sense, we are concerned both with "projective" and "reflective" forms of evaluation.

2) Projective Evaluation and Queueing Analysis

From the data, simple statistics on the arrival rates of calls for maintenance service and on service times can be isolated by priority class (emergency, routine), craft (licensed, skilled) and development. This information can be used to determine the extent to which the consequences of alternative operations policies can be projected. These consequences include the size of backlogs and the costs and waiting times associated with projected levels of congestion. If the system follows one of several well-known behavior patterns, we will be able to make quite detailed queueing estimates. The importance of probabilistic models lies in the fact that waiting times often increase exponentially with only incremental increases in calls for service. Such models can help suggest a policy which could avert congestion by projecting the conditions under which it is likely to occur.

If properly structured and carefully implemented, such information could be usable by and useful to managers, supervisors, craftspeople, tenants and operations staff, and could provide a method of "planning for" work orders which are about to "happen" rather than requiring they be prescheduled long in advance. Strict advance prescheduling can limit the system's ability to respond to new information and therefore increase costs and waiting times for many calls. Conversely, a greater ability to dynamically respond to, create and communicate information would provide an opportunity to choose from among a wider range of operating policies, scheduling methods and crafts roles than at present.

While the testing of more complex systems and policies would be more easily done on a mainframe computer, enormous centralization of information and operations is not a technical requirement of a good evaluation system. The present analysis was intentionally undertaken with little prior knowledge of on-site maintenance operations in order to test the ability of the data themselves to provide information useful to central operations staff. In practice, however, the value cannot be overemphasized of having dedicated people at each development capable of linking simple but well structured data analysis to day-to-day operations. Indeed, solutions to a host of system design, implementation and policy questions need to continuously adapt. At any time, most of these issues should be able to be resolved outside of central operations - in the developments, where maintenance operations take place. One performance measure of the BHA work order processing & evaluation system's design might even be the degree to which decision making power is enhanced and the range of choices increased for those working (and living) at any place in the system.

A dynamic scheduling, processing and evaluation system as outlined here is based on the idea that noone knows exactly when and where the next need for maintenance will "happen", or what crafts will be required to service it, but that we can make reasonable projections of many calls that are likely to occur. We have called this process "projective evaluation" because it is concerned with how anticipated decisions result from and help create information which continuously evaluates the system. The extent to which this process operates continuously and maintains its adaptability over time may be another measure of its success.

3) Performance Measures for Reflective Evaluation

If observed arrivals and service times do not behave in ways that allow more powerful probabilistic models to be constructed, other general relationships in queueing theory should still enable us to estimate several types of congestion. Either way, to undertake "reflective evaluation" the system first needs the ability to draw meaningful summary statistics (performance measures) from recorded data. These measures can then be used to observe (rather than test) how general policies have worked in practice and how demands change over time. They also help specify more complex policies and queueing models we might want to test.

Averaged performance measures produced by reflective evaluations could already increase the system's ability to respond to demand, since they provide somewhat the same type of information as queueing models. One difference is that these performance measures may be used as inputs to a probabilistic model. By themselves, however, many reflective measures are simply percents and averages taken from monthly data when these data are in a form that permits clear interpretation. As currently recorded, it is difficult to make use of observed maintenance data. One problem is that observed measures may be affected by other variables than those we wish to measure, and it may be difficult to attribute performance changes to specific changes in policy. Therefore we also want to distinguish between operating problems and data structure problems, and suggest what questions performance measures may be designed to answer.

Samples of these performance measures for reflective evaluation have been included in subsequent chapters. Like the projective evaluation system which could grow from them, such measures should also provide information useful to those throughout the maintenance system. Using sample data from October 1983, steps have been demonstrated for extracting several monthly performance measures, mostly making use of simple database management routines. Finally, several operating policies are listed which might improve maintenance operations as profiled by these performance measures, and suggestions are made for reducing ambiguities in retrospectively evaluating trial policies.

CHAPTER II

HOW MAINTENANCE OPERATIONS ARE REFLECTED IN THE STRUCTURE OF WORK ORDER DATA

1) Profile of Past Operations

In any year, the BHA maintenance system generates nearly 60,000 work orders. To make this information useful to the system, we need to have a clear image of how maintenance has happened operationally. A primary investigation should help us choose a sample for observation and suggest how the data might be "sliced" to provide useful feedback. A balance must be struck between an overly fine grained analysis that prevents us from drawing general conclusions and one so global that it is useless for suggesting specific operating solutions.

To get a sense of the relationship between maintenance operations and the structure of work order data, several interviews were held at the BHA with Gwen Friend, who also provided recent literature. These enabled the following general profile of operations to be drawn.

The process generating work orders can be seen as one in which calls for service arrive at a processing facility. Each large development at the BHA typically has an office that generates these work orders. In most cases, a tenant discovers a need for performing some type of repair, whether in his or her own apartment, or somewhere in the development. The tenant then calls the maintenance office, where a work order clerk asks a series of questions to determine whether repair is needed, and if so, how it shall be described on the work order form. In the case of repairs requiring service by a variety of craftsmen with different skills, a work order is generated for each component of the repair job. Thus, repairing a hole in a wall may require separate work orders for carpentry, plastering and painting.

Throughout the day, these work orders are collected by the development's maintenance supervisor, sorted by priority according to emergency or routine status, and scheduled for service later that day or on following days, together with those orders outstanding from previous days. As the supervisor sorts work orders, he also estimates the service time and the cost for each job.

Each supervisor is in charge of a maintenance crew specific to that development. With a few exceptions, crews at large developments do not service work orders from other developments, but operate semi-autonomously at their own locations. These crews are of different sizes and receive calls for service at different rates. Furthermore, maintenance crews are composed of craftspersons from several specialized craft or skill types, and workers in each category only service work orders corresponding to their particular craft.

Therefore, as the supervisor sorts work orders by priority class, he also sorts them by craft type. Within each priority class and craft category, work orders are to some extent serviced in a first-come, first-served (or FIFO - first-in, first-out) manner. Although supervisors maintain different scheduling styles, a set of limited guidelines for priority scheduling were centrally adopted three years ago. There has been a tendency throughout the BHA, however, to backlog work orders which are either more diffcult to service (those involving heavy budgetary demands, hard-to-get parts and supplies, seasonal work, etc), or are considered trivial because they will be serviced by other long-run maintenance operations, or both. The type of orders backlogged may also vary from one development to another, depending on any other development-wide projects that may have priority, such as landscaping or general infrastructure repair.

This profile roughly outlines the system for several years prior to February 1984. It also indicates that the data available describe a maintenance system which for some time had not undergone major changes, and that sample data may be used to create a more detailed profile of system operations for this period. October 1983 was chosen as a sample for the analysis, because it is one of the more recent months for which information seems to be representative of year-round maintenance operations.

2) Recent and Proposed Operations Changes

As this continuous but uneasy Pax Romana may be both too complex and expensive to maintain, however, several major changes have recently been made, and others are in store. In February 1984, a policy was announced by which all outstanding work orders more than a month old were to be purged, other than emergencies and those involving energy conservation, cost savings and inspections. In addition, no new work orders are to be accepted outside of these categories. This move followed the expansion of the Living Unit Inspection (LUI) program, designed to eliminate tenant-initiated routine work orders by servicing them once yearly for each apartment. Under the LUI program, each apartment is prescheduled for inspection on a particular day of the year. Apartments are inspected in succession, and one round of inspections (a visit to all apartments) is supposed to take one year.

The idea is that most minor or routine tasks will thereby be discovered, and that routine work orders for each apartment will be generated once yearly by the inspector. A tenant calling the maintenance office with a request for routine service is then told to wait until the date on which his or her apartment is scheduled for inspection, even if that date is eight months away. Although the strictness with which development managers and supervisors enforce these rules may vary, the policy assumption is that a unit of waiting time for routine service is far less costly to the tenant than a unit of waiting time for emergency service. Reductions in turn-around times for tenant-initiated emergency service are expected to follow from the increased efficiency of servicing all routine work orders simultaneously for a given apartment. Further reductions in both turn-around and total service times are expected as living unit inspections discover more general signs of decay and take care of these problems before they degenerate into emergencies.

Comparison of this system with the earlier, backlogged version of October 1983 would be a fascinating exercise, but data for the new system will not be reliable until the program has been in operation for several months. Such a comparison would be difficult for three reasons, however. First, one cannot know how long a tenant has effectively been waiting for routine service up to the time his or her apartment is inspected. Under the assumption that the tenant-borne costs associated with routine waiting times are negligible, the new system would probably compare favorably with that operating in October. But the implications of other assumptions are more difficult to measure. Second, the decision to purge backlogged work orders will exaggerate the reductions in turn-around times which are a consequence of the upgraded LUI program by mixing them with reductions resulting from purges. This difficulty can partly be overcome by looking at the October turn-around times on a restricted interval of 1 to 20 working days - as if the purges had also been in effect in October. But here, the effect of congestion would have been to also raise turn-around times for those work orders served relatively quickly, and there may be no way to eliminate the interaction effect. Third, a new maintenance contract has recently been negotiated under which many tasks formerly coded as "specialized" (and thus requiring service from a craftsperson only of that skill type) have been reclassified as "neutral". This permits a wider range of craftspeople to service many routine tasks, reducing the total response time required for multicraft routine maintenance requests.

There are also further changes ahead. These concern the process by which and the form in which information generated by work orders will be recorded and used for evaluation. In September of this year, the BHA is scheduled to adopt a modified version of the Dallas Housing Authority's work order processing system, together with the computer hardware and software which is an integral part of it. While the final form the system will take is still unclear, each maintenance request is expected to be specifically coded on the work order form, resulting in about 400 distinct task codes organized into 9 basic crafts. It is hoped that the increased detail will enable materials and labor costs to be standardized by job code.

In addition, the creation of a centralized work order processing office is foreseen in order to automatically generate work order forms and to reduce recording errors, which have plagued the decentralized, development-based offices. These information processing changes will not necessarily alter the way in which work orders are actually serviced, but tenants will now call a central maintenance facility instead of a neighborhood office. Rather than centralize scheduling, however, the new system is first of all intended to better monitor and assist the actual servicing changes that have already taken place through purges, a new union contract and an expanded LUI program.

High hopes are placed on the ability of the new work order form to provide more detailed evaluative information. The form in present use, however, also has a detailed task categorization scheme. Although only slightly less elegant than the one proposed, the present scheme is used only intermittently. It is possible that the amount of usable detail in the proposed system will be increased by the improved form this information takes, but there are certainly no guarantees. Given the underutilization of data provided by current work order forms, we need to ask how much is due to the lack of a more elegant data structure. We then ask what performance measures the future work order form and processing system may be better designed to inform. A major objective of this design should be to eliminate the errors and overlapping variables that characterize the work order form used up to the present. Let us then structure our data analysis by looking more closely at the form this information takes.

3) <u>The Work Order Form</u>

Over the past several years, data have been collected for all of the approximately 5,000 work orders serviced each month. Roughly 1,000 of these come from the smaller and generally newer elderly developments, and the rest from family developments, which usually have more severe maintenance problems. A typical monthly tape thus contains about 5,000 records of raw data, - each record or line summarizing information from a single work order. The data entries represent categories of information found on any work order form, and include:

- the work order number, development, apartment, supervisor and employee numbers;
- class, craft, cause and task code;
- call-in date, assignment, completion and inspection dates;
- estimated labor time, labor costs and materials costs;
- actual labor time, labor costs, materials and total costs.

A copy of the work order form used during this period can be found in table 1 of Appendix II. Several of these categories will not be used for the analysis, and several more require some explanation.
(1)

The work order, apartment, supervisor and employee numbers do not concern us here, although they are obviously useful for tracing the history of a particular work order, for identifying building systems problems and comparing employee performance. The cause category has

---

(1)
A sample of recorded data from the October 83 tape can be found in table 2 of the appendix. Only obvious errors have been deleted. Some variables have also been transformed to facilitate calculations. All weekends and holidays have been eliminated to create a working-day profile of operations.

also been dropped due to the ambiguities (and therefore errors) associated with judging why a specific problem occurred. Such judgements tend to vary from one worker to another. Task codes have not been considered because the information is both too fine grained for our analysis and too badly organized to specify standard costs per task. Estimated and actual costs have also been left out of the analysis, although in the future it will be interesting to study how these costs change with new operating policies.

4) Variables Chosen for Analysis

We assume that turn-around times are the statistic of greatest importance to tenants in need of service, and that they are also a measure of how efficiently the system responds to demand. Tenants bear implicit costs based on the lengths of these turn-around times, and it is these "hidden" costs we want to know more about. Without giving them explicit dollar values, we want to understand the variables which affect turn-around times, using information on arrival rates of calls for service and on service times associated with different class and craft types. These parameter estimates, together with information on the number of servers (or craftspersons) by development, and the specification of the queueing discipline (rules for the order in which calls are to be serviced), will usually be all we need to specify a variety of models.

A) Development

If each large development has its own maintenance crew, then each also has its own unique set of queues to be modelled. Since the larger family developments interest us primarily, we have selected 12 developments which generated over 200 work orders in the month of October. These are:

Dev No.	Name	Units	Work Orders	WOs per 100 Units
101	Charlestown	1149	201	17.5
103	Mission Hill	1023	435	42.5
105	Orchard Park	774	190	24.5
108	Maverick	414	292	70.5
109	Franklin Hill	375	222	59.2
114	Mission Hill Ext	588	197	33.5
120	Columbia Point	1504	233	15.5
123	M. E. McCormack	1016	213	21.0
124	Old Colony	873	475	54.4
501	West Broadway	676	229	33.9
508	Orient Heights	354	172	48.6
510	Gallivan Blvd	251	173	68.9
	TOTAL	8997	3032	

At times we will look at information aggregated across these developments, and at other times focus in on the operations of a specific one.

B) Class

More detailed information on variations in arrival rates and service times can be derived by measuring priority class and craft type against call-in dates, completion dates and labor times. The 12 class codes are meant to describe the overall type of work order. In one sense, orders are coded by priority (emergency, routine); in another by special source, if any (Housing Inspection Dept, Emergency Response Service, Court Order, LUI), and in yet another sense by special reasons for which service might be required (lead paint, vacancy, extraordinary, modernize, safety, security). Any work order falls into at least one of these subcategories, but may fall into all three. Since only one of these overlapping class codes is specified on a work order form, however, a great deal of information can be lost or misclassified. In the future, this information might be organized into two unique and exhaustive categories covering priority and source, with perhaps a third for any special characteristics not detailed elsewhere.

For this analysis, we are primarily concerned with the priority in which calls are serviced. Fortunately, central BHA policy has attempted to sort each of the 12 class items by the order in which they should be handled. Assuming that central policy has had some effect, this allows us to retrospectively analyze all work orders by priority. Hereafter, we will use the term "class" to refer to service priority. The variable "class" takes on two values - emergency and routine - with respect to which all work orders are defined. The terms "emergency" and "routine" - used in a more comprehensive sense than that in which they appear on a work order form - are defined to include work orders coded as

Emergency	Routine
1 emergency	4 vacancy
2 lead paint	5 routine
3 Housing Inspection Department	6 extraordinary
8 safety	7 modernize
9 security	12 Living Unit Inspection
10 Emergency Response Service	
11 court order	

This information could be used to model either the existence of one priority-based queue for each development, or two independent queues - one for each priority class.

C) Craft

The craft category refers to the type of craftsperson needed for the maintenance service. According to union contracts no longer in effect, each work order was required to be serviced by a craftsperson of the corresponding type. This policy substantially increased the costs and turn-around times associated with some jobs, since a painter, for example, could not begin painting until an expensive, licensed electrician had arrived to remove the light switch cover plates. These rules apparently went much further than state law, which only required that more complex tasks be performed by workers of the corresponding specialized skill type. The fact that licensed craftspersons are both more expensive and sometimes legally required for certains tasks suggests that we differentiate licensed work orders from those of other skill types. The 19 craft codes might then be aggregated into two groups. Similar to our definition of class, the variable "craft" takes on two values - licensed and skilled - by which all work orders (except rare manager actions) can be defined.

Licensed	Skilled
4 electrician	1 appliance
10 plumber	2 auto mechanic
	3 carpenter
	5 fireman
	6 glazier
	7 laborer
	8 painter
	9 plasterer
	11 roof
	12 site,structure
	13 steamfitter
	14 welder
	16 exterminator
	17 tile setter
	18 bricklayer
	19 cement finisher

Although it can be argued that much information is lost by grouping so many skills together, licensed work orders account for 42% of all work orders generated, whereas any one skill type accounts for very little. Since, for the period under study, craftspersons of a particular type were required to service each order, one might assume that a different queue exists for each craft. Workers in some queues may therefore have a great deal more idle time than others. We have ignored these distinctions here for two reasons: one is that workorders are backlogged, usually by the hundreds, for each development. This suggests that most workers always have calls in queue. The other reason is that average service times do not appear to differ dramatically by skill type. Although idle times by craft have not been investigated here, such retrospective information could help judge the accuracy of idle times projected by queueing models. Further studies should certainly attempt more fine grained analyses, once the new work order processing system is securely in place.

D) Class/Craft Combination

Each work order has both a class value and a craft value. We may wish to construct a model having two craft queues for each development. Within each of these queues, work orders may be assigned nonpreemptive priority based on their class values. On the other hand, another simple model might be based on four independently operating queues in which either licensed or skilled servers perform only emergency or routine work. Each work order could then be assigned to one of these queues, based on unique combinations of class and craft.

- emergency licensed (emlic)
- emergency skilled (emski)
- routine licensed (roulic)
- routine skilled (rouski)

5) Conclusions

Overall, we have defined four variables by which the parameters can be measured.

A) development
B) class
C) craft
D) class/craft

The strategy we follow is to use linear regression, analysis of variance, and goodness of fit tests to estimate the probability distributions of the number of arrivals and of service times, and to see how these estimates differ with respect to the variables outlined above. Turn-around times have been analyzed in a similar fashion to provide a base for comparison with those which could eventually be estimated by queueing models. Once reasonable assumptions can be made regarding the forms taken by these distributions, parameter estimates used as inputs to queueing models can then be easily calculated using the same database management techniques which provide other reflective performance measures.

CHAPTER III

ESTIMATING ARRIVAL RATES

1) Objectives

By observing arrival processes for the queues we have defined, we can determine whether simple queueing models can be fit to these arrivals. Here, a variety of methods are used to test whether arrivals correspond to a Poisson process. Due to its pleasant mathematical properties, this process has become the basis for some of the simplest and most powerful queueing models in common use. Another reason it is often used is that the Poisson process describes events which occur randomly in time. Therefore, if systematic maintenance problems are not occuring, we would expect work orders to arrive approximately in a Poisson manner.

In the case of Poisson arrivals, scheduling operations may clearly benefit from queueing models capable of projecting congestion. We would therefore like to point out any Poisson arrival processes and suggest models that may help reduce turn-around times. But since Poisson processes describe random events, we can also use our test results as criteria for determining whether maintenance problems are ordinary or systematic. If arrivals are non-Poisson, then the maintenance system would benefit from an ability to diagnose problems by separating independent occurrences from systematic ones. We begin with a short profile of Poisson processes, followed by an explanation of test methods and results. An interpretation of these results can be found at the end of this chapter.

2) The Poisson Arrival Process

The exact time at which a call for service will arrive is not known in advance. If one were to visit a maintenance office, the time interval between any two arrivals would be different each time. But over a large number of observations, these interarrival times might follow one of several classic probability distributions. If we take one day as our time unit of analysis, we can observe a daily demand function showing the aggregate number of work orders that arrived on a particular day throughout the 12 high-demand developments we have chosen. A plot of daily demand in October (fig 1), shows work orders as a function, $\lambda(t)$, of the working day on which they arrived. Weekend days have been eliminated.

figure 1

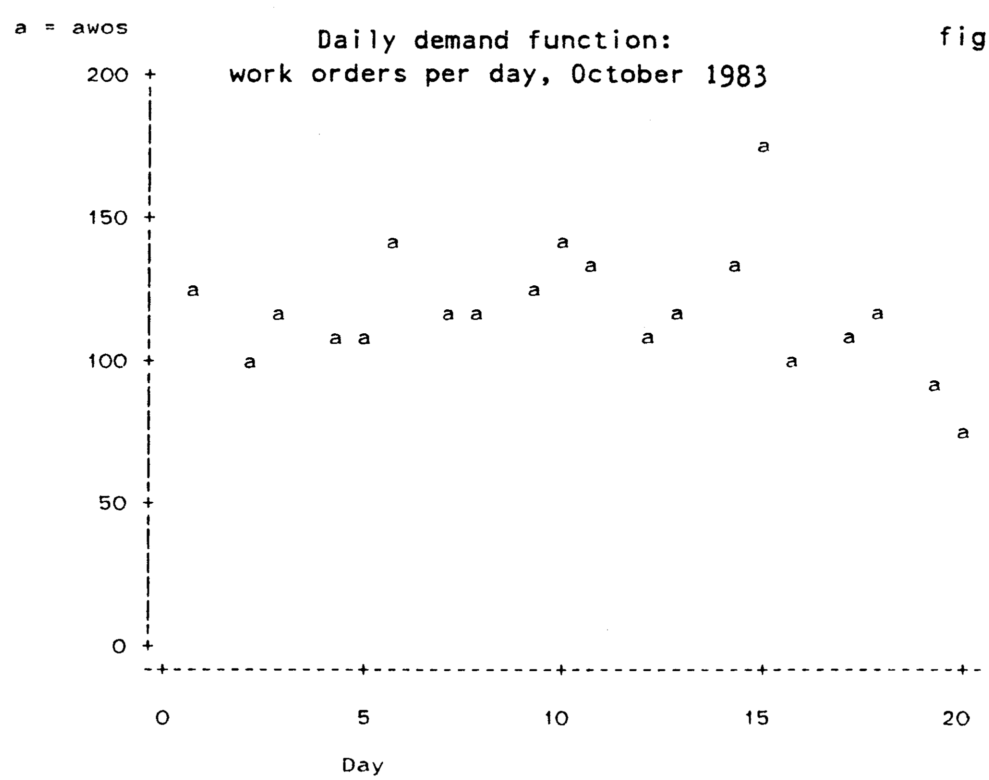

Toward the end of the month, the number of arrivals appears to drop because the data contain call-in dates only for those orders serviced in October. Especially near the end of the month, calls are arriving that will not be immediately serviced, and not appear in the data until November. The mean arrival rate of these calls is 117 per day with a standard deviation of 21. Further identification of errors would probably raise the mean and reduce the standard deviation somewhat.

We begin with the null hypothesis that arrivals can be modelled as a time-homogenous Poisson process. This process has been shown especially useful in approximating many aspects of urban service systems. It typically applies when customers (arrivals) are drawn from a large population, any one of which has a very small probability of "arriving" on a given day. For general Poisson processes, the numbers of arrivals occurring during non-overlapping time intervals are statistically independent. If in addition the interarrival times follow an exponential distribution, then the actual arrival pattern is called a time-homogenous Poisson process. In any time interval, these Poisson arrivals are said to occur "randomly" in time, and the probability that there are exactly n work orders generated in a particular day is given by

$$Pr(n/day) = Pr_n = \frac{(\lambda t)^n e^{-\lambda t}}{n!} \qquad \begin{array}{l} t = 1 \text{ day} \\ n = \text{arrivals/day} \end{array}$$

If this equation fits the distribution of the number of arrivals as we observe them in the data, then other well-known properties of homogenous Poisson arrival processes will enable us to define the system's behavior particularly well.

There are a variety of methods for testing whether this hypothesis is true for the observed arrivals. First, we plot a histogram showing the number of days in October on which n arrivals occurred (fig 2).

figure 2

wos/day		frequencies
90	xx	2
100	xx	2
110	xxxxxxx	7
120	xxx	3
130	xx	2
140	xx	2
150	x	1
160		0
170		0
180	x	1

With the exception of the outlier to the right, the histogram resembles that of a Poisson probability mass function (pmf), and suggests we test the actual pmfs more closely for specific developments. For each development, we then calculate means and variances of daily arrival rates (λ) by class, craft, and class/craft combination (table 3).

Choosing development 108, a relatively busy project, we next create table 1, showing the daily demand function. From column 1, we can make a separate matrix (table 2) to show the number of arrivals per day and the fraction of days on which they were observed. This is the observed probability mass function ("frac", fig 3a). Using the observed mean arrival rate λ, we also generate a theoretical pmf (Pr_n_arr) showing how arrivals per day would be distributed if they corresponded perfectly to a time-homogenous Poisson process. This is plotted in fig 3b.

table 1

Daily demand function,
all 12 developments

#	ci_day	wos	em	rou	lic	ski	emlic	emski	roulic	rouski
1	277	16	6	10	8	8	3	3	5	5
2	278	5	1	4	3	2	0	1	3	1
3	279	6	1	5	5	1	1	0	4	1
4	280	7	4	3	3	4	2	2	1	2
5	281	5	2	3	1	4	1	1	0	3
6	285	17	12	5	9	7	7	5	2	3
7	286	12	3	9	4	8	1	2	3	6
8	287	15	4	11	8	6	3	1	5	6
9	288	9	3	6	3	6	1	2	2	4
10	291	17	4	13	4	12	1	3	3	10
11	292	17	7	9	4	13	2	5	2	8
12	293	7	4	3	4	2	2	1	2	2
13	294	11	8	3	3	8	3	5	0	3
14	295	13	2	11	6	7	2	0	4	7
15	298	16	3	12	5	11	0	3	5	8
16	299	12	4	8	1	11	0	4	1	7
17	300	18	10	7	7	11	1	9	5	3
18	301	3	1	2	1	2	0	1	1	1
19	302	5	2	3	0	5	0	2	0	3
20	305	11	4	6	4	7	1	3	3	4

table 2

Probability mass function
all work orders
development 108

#	wos	freq	frac	Pr_n_arr
1	3	1	0.050	0.003
2	5	3	0.150	0.021
3	6	1	0.050	0.039
4	7	2	0.100	0.062
5	9	1	0.050	0.107
6	11	2	0.100	0.119
7	12	2	0.100	0.110
8	13	1	0.050	0.094
9	15	1	0.050	0.055
10	16	2	0.100	0.038
11	17	3	0.150	0.025
12	18	1	0.050	0.015

3) A Summary of Tests for Poisson Arrival Processes

A) Chi-Square and Linear Regression Tests

A Chi-Square goodness of fit test was first used to compare frequencies based on these two distributions, but low counts in the expected cells made it necessary to aggregate the frequencies into classes so general that the test became useless.

Then we attempted a simple linear regression analysis using one probability distribution to predict the other. The details of this analysis are presented in the appendix to Chapter III. Ambiguities in the results led us to test the sensitivity of regressions to known figure 3a

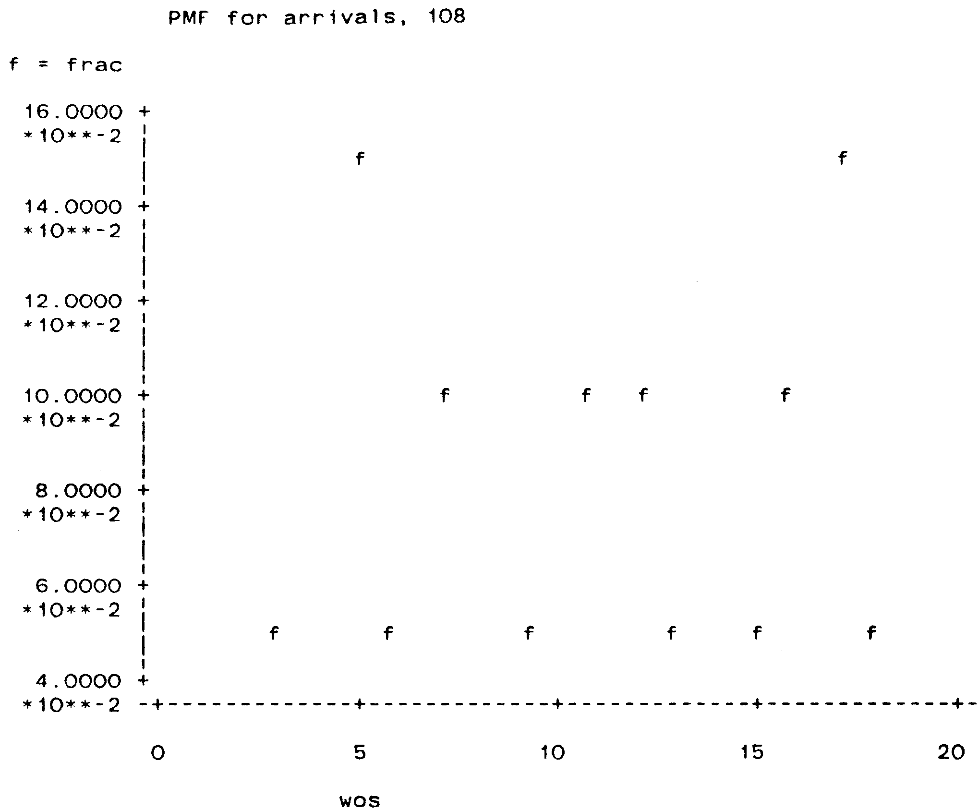

figure 3b

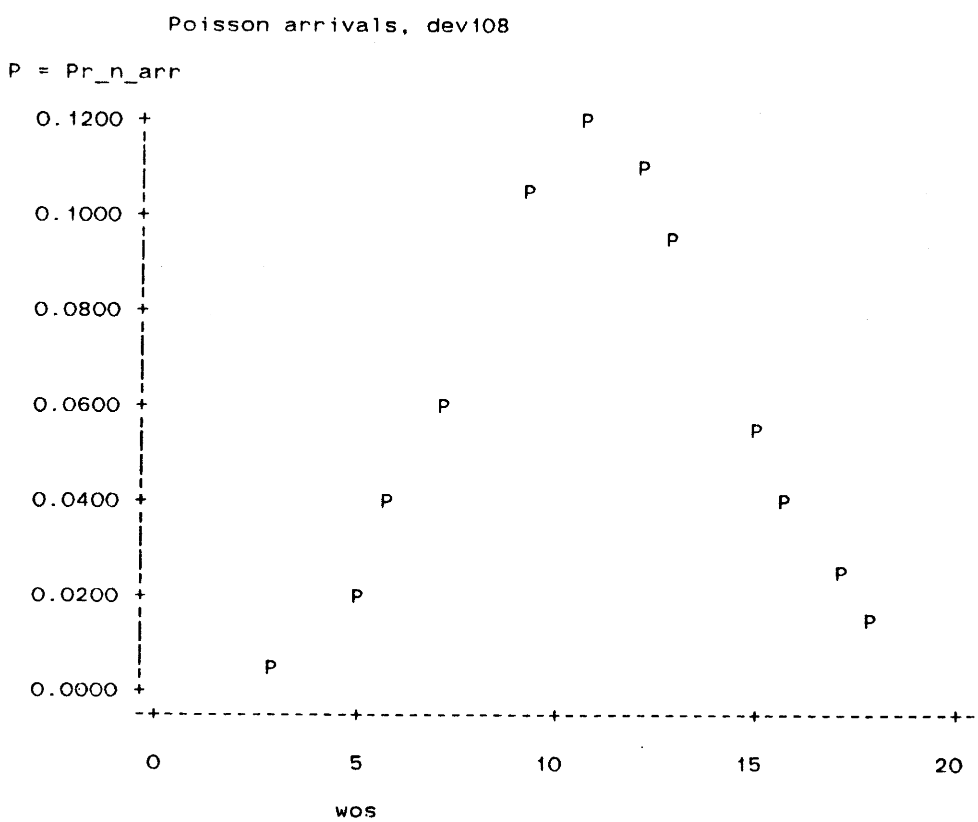

differences between probability distributions. For low arrival rates and small sample sizes, regression results appear to underestimate these differences. The tests are especially insensitive to differences between cumulative distributions.

Finally, we found that non-constant variances require that a weighted least squares approach be used. Such a regression test would account for inhomogenous variances (just as the Chi-Square test does), but it would also indicate how consistently and in what direction the distributions differ. No time was available to execute a weighted least squares regression test, however.

B) The Kolmogorov-Smirnov Test

Another method for testing the degree of agreement between a cumulative function of observed probabilities and an hypothesized cumulative distribution is the Kolmogorov-Smirnov (K-S) goodness of fit test. This avoids the problem of glossing over differences between the two distributions because it subtracts each cell in the observed function from the corresponding cell of the expected or hypothesized distribution, and uses the maximum resulting difference as a test statistic, D, whose distribution in repeated sampling is known.

$$D = \text{maximum}|\text{expected} - \text{observed}|$$

Here, D is small under the null hypothesis that the observed probabilities are equal to the Poisson. A large D, however, leads us to reject the Poisson model at some chosen significance level, whose critical value depends on the sample size, n. Like many tests, the K-S statistic tends to favor the null hypothesis for small sample sizes. In our case, we want to be careful because the finer our class and craft categories are sliced, the more the arrival processes may appear Poisson.

We can give the non-Poisson alternative the benefit of the doubt by choosing a significance level of $a=0.20$. By this decision rule, we are willing to reject the Poisson model if our observations appear non-Poisson while allowing a 20% probability of chance error. Critical values at the 0.05 level have also been provided for reference. But neither decision rule accounts for recording errors or provides us with a definitive answer.

Table 4 in Appendix III summarizes the outcomes of the K-S test for development 108. The distributions used to calculate the test statistic D are given in table 5 and figure 2 of the appendix. Like the regression on cumulative probability distributions, the K-S test judges all arrival processes to be Poisson at any reasonable level of significance. An analysis using the K-S test to compare a series of randomly generated probabilities to a Poisson distribution based on different values of the mean arrival rate shows the test to be highly sensitive to changes in this rate. The problem faced with cumulative regressions has to some extent been overcome (here we used random rather than observed distributions to test for changes in the arrival rate).

The interesting conclusion, however, is that the K-S statistic also appears highly sensitive to changes in sample size. Our tests for development 108 used a sample of 20 days and found arrivals to be

Poisson. But for the same sample size, the K-S test also judges a cumulative distribution of random uniform probabilities to be Poisson, provided we use the mean arrival rate implied by those random numbers as our estimate of lambda in generating the expected Poisson distribution. The same is generally true for all sample sizes less than 60. At 60, the test discriminates well between hypothesized and random distributions for an arrival rate of 18 work orders per day, but not for 12. It is only for sample sizes of at least 100, however, that it appears clearly useful for testing hypotheses throughout the range of the arrival rates in our observed distributions. The test is therefore inadequate for estimating arrival distributions from monthly samples. The results of these investigations are presented in tables 6 and 7 of Appendix III.

C) <u>Measures of Central Tendency Across Developments</u>

Finally, Poisson arrival processes also have the property that the mean number of Poisson events occurring in a given time period is precisely equal to the variance. The means and variances of daily arrival rates are given in tables 3a and 3b by class, craft, class/craft and development. If the means and variances for any development are equal, this does not prove the arrivals are Poisson, but very different means and variances would suggest that they are not.

Using data from all 12 developments, linear regressions were run to test how well the observed means predict the variances. While this masks the differences between any two developments, it permits us to view tendencies across them. In making such generalizations, however, one must also be careful to avoid the "ecological fallacy" of interpreting results from regressions on sample means just as one would interpret them from tests run on actual data. Regressions on means naturally tend to account for more of the observed variation simply because the means of several samples vary less than do the data in those samples.

Means & variances of arrival rates
by class, craft & development

table 3a

dev	wos_m	em_m	rou_m	lic_m	ski_m	emlic_m	emski_m	roulic_m	rouski_m
101.000	8.150	2.050	6.050	3.100	5.000	0.900	1.150	2.200	3.800
103.000	14.400	1.050	12.950	6.400	8.000	0.600	0.450	5.600	7.350
105.000	7.350	5.250	2.100	2.700	4.650	2.400	2.850	0.300	1.800
108.000	11.100	4.250	6.650	4.150	6.750	1.550	2.650	2.550	4.350
109.000	8.150	6.600	1.550	3.650	4.500	3.100	3.500	0.550	1.000
114.000	7.750	3.900	3.850	2.750	5.000	2.000	1.900	0.750	3.100
120.000	10.150	8.800	1.350	5.200	4.950	4.600	4.200	0.600	0.750
123.000	7.650	1.000	6.650	3.150	4.500	0.600	0.400	2.550	4.100
124.000	20.000	1.450	18.500	8.350	11.650	0.450	1.000	7.900	10.600
501.000	8.900	0.350	8.500	3.800	5.050	0.150	0.150	3.650	4.850
508.000	6.950	0.500	6.450	3.150	3.800	0.350	0.150	2.800	3.650
510.000	6.750	6.100	0.650	2.400	4.350	2.200	3.900	0.200	0.450

table 3b

dev	wos_s^2	em_s^2	rou_s^2	lic_s^2	ski_s^2	emlic_s^2	emski_s^2	roulic_s^2	rouski_s^2
101.000	24.344	4.260	14.364	3.675	11.472	1.147	1.608	3.222	6.381
103.000	31.618	1.313	32.262	6.462	16.524	0.569	0.682	6.991	15.610
105.000	15.920	12.831	3.881	4.012	9.400	3.830	5.607	0.221	3.640
108.000	24.305	8.934	12.236	6.240	12.931	2.683	4.661	2.996	6.260
109.000	12.766	16.459	3.629	4.661	6.579	5.462	7.840	0.787	3.262
114.000	21.987	6.305	9.923	3.881	12.946	3.580	3.356	0.828	7.885
120.000	36.024	26.378	2.660	9.536	15.312	7.513	11.540	0.780	1.355
123.000	13.184	0.947	10.452	3.397	7.840	0.780	0.358	1.946	6.938
124.000	64.626	3.629	65.740	12.766	62.758	0.576	2.631	11.465	62.457
501.000	20.730	0.239	20.794	8.486	6.472	0.134	0.134	8.556	6.240
508.000	14.577	0.473	14.258	5.818	6.802	0.450	0.134	5.009	6.975
510.000	14.304	10.824	1.293	2.462	9.505	2.062	6.938	0.274	0.787

Four rounds of regressions were run testing the means of each class, craft, and class/craft category in table 3a against the corresponding variances in the columns of table 3b. Results are given in table 8 of the appendix. The high slopes of the fitted regression lines suggest that typical arrival processes for most classes, crafts, and class/craft combinations may conform more closely to hyperexponential or uniform distributions than to the Poisson. This appears especially true for skilled and routine/skilled arrivals. Another possibility is that many arrivals correspond to time-varying or inhomogenous Poisson processes - in other words that arrivals in non-overlapping intervals are independent, but that "rush hours" and peak periods also occur.

Failures in all classes and crafts do seem to be associated, since the only categories which appear somewhat Poisson are also the most finely sliced. These include emergency and emergency/skilled problems, whose variances were smaller in development 108. On this basis, other developments exhibit a greater overall tendency toward related and systematic failures of the emergency/skilled type. Development 120 seems to have had special problems in this respect. Further studies should isolate and detail such problems for each development.

4) Interpreting Test Results

Based on all of the tests presented above, we can neither accept nor reject the hypothesis that the number of arrivals is Poisson distributed. For our sample sizes, the only tests which led us to accept the hypothesis proved highly insensitive to differences between distributions. There are also major problems with the tests which led us to reject the Poisson model. It is certainly true, however, that there is considerable variability in daily arrivals at any development. This may be due to systematic maintenance problems which cause work orders to be generated in groups.

In continuously modernized buildings without systematic problems, one would expect work orders to be generated in an entirely random manner. By definition, however, buildings are constructed at one point in time. Because their components are standardized, those components also tend to need repair or replacement at about the same time, and this causes work orders to be generated non-randomly. The existence of group arrivals is partly due to the fact that several work orders are often generated for a single repair job. One job may appear to be many separate jobs. But groups of non-Poisson routine arrivals also indicate that minor but associated problems are simultaneously occurring due to worn out building components. These problems are associated because they happen concurrently (such as falling plaster), not because one minor problem necessarily causes another. The analysis indicates that such routine problems make up the major portion of non-Poisson work orders for development 108, and that the repairs required to service such problems tend to vary greatly by skill type. One would expect to find such a state of affairs in old developments (Maverick was first occupied in 1942) having maintenance crews that respond to minor problems one at a time.

On the other hand, non-Poisson arrivals can also be systematic in the sense that one problem contributes to or causes another. Extreme examples include emergency roof, site or structure failures leading to a variety of smaller "associated" problems without apparent causal connections between these minor failures. Similarly, licensed problems such as the failure of outdated plumbing or electrical systems may generate a host of associated routine work orders which appear unrelated. Such bunches of routine work orders could well indicate that more serious problems are about to happen, and that there are major problems underlying these minor failures.

The fact that we are unable to assume arrivals to be Poisson distributed does not help us to construct simple and accurate queueing models. But it does suggest that the more serious (and interesting) maintenance problems involve precisely those work orders which do not arrive in a Poisson manner. We can use this fact to help us identify possible systematic failures.

It is important that we not confuse service priorities with the degree to which a problem may result from systematic failures. Emergency work orders are not necessarily the main indicators of serious systems problems. While they require priority service, these emergencies can be the result of either related or isolated occurrences. If most arrival categories are non-Poisson, we should view work orders as symptoms of a maintenance process that systematically generates certain types of problems. By definition, these problems are the cumulative effects of systematic neglect. Just as systematically, therefore, we must define the underlying failures which are likely to be causing different "bunches" of associated problems, whether they be emergency or routine, licensed or skilled.

5) Diagnosing Systems Failures

A) The Bottom-Up Approach

A method for identifying system failures can be briefly outlined as follows. One first identifies a specific bunch of problems that are regularly occurring. Again, these problems are "associated" because they are of the same type, rather than because they necessarily cause each other. Only rarely will bunches will be observable in the space of one day, since failures may not create symptoms all at once. It is more likely that neglected systems will be identified by isolating "bunches" that form over periods of several weeks or even months.

Each bunch can then be viewed as one element in a hypothetical set of such symptoms. At the disaggregated level, these bunches may appear unrelated to other elements of the same set. Indeed, workers may fail to notice (or to record) the existence of a "set" of problems, since they tend to service work orders of a particular type. But these diverse elements may in fact be related through systems failures. The bottom-up approach asks what possible systems failures are implied by the diverse non-random problems we observe in work order data.

B) The Top-Down Approach

The same method for identifying systems failures can be seen from another perspective. First, a list of possible systems failures is created. From these, one can generate hypothetical sets of problems which could follow from such failures. Each set implies a corresponding systems failure, and each element in a set constitutes a symptom or bunch of associated work orders which would appear in the data if the related systems problem may be occurring. Any systems problem may therefore imply a variety of possible symptoms, not all of which would necessarily manifest themselves in the event such failures exist. From the appearance of a few symptoms or bunches, one could then diagnose possible failures. Furthermore, if a given underlying problem is suspected, one should be able to anticipate other bunches of work orders that may "happen", since these are simply the remaining elements of the system's set. The diagnosis is complicated somewhat, however, by the fact that a given symptom may imply a variety of systems failures, and therefore belong to several sets at once. This should be clear from the diagram (figure 4) on the following page.

The term "systems failures" should not conjure images of exploding boilers, flooded corridors and collapsing ruins, however. The reason we need a method for diagnosing them is precisely that they may otherwise go unrecognized. The method outlined here is essentially meant to be used for analyzing neglected or undermaintained building systems, since they may be hidden behind work orders which appear to workers as isolated events. Rather than cataloguing every imaginable system, the bottom-up approach should be used initially to identify those "most neglected" systems for which system sets should be constructed. The detailed task code scheme provided by the new work order processing system should be quite useful for this.

Figure 4. **DIAGNOSING SYSTEMS FAILURES FROM TASK CODE DATA.**

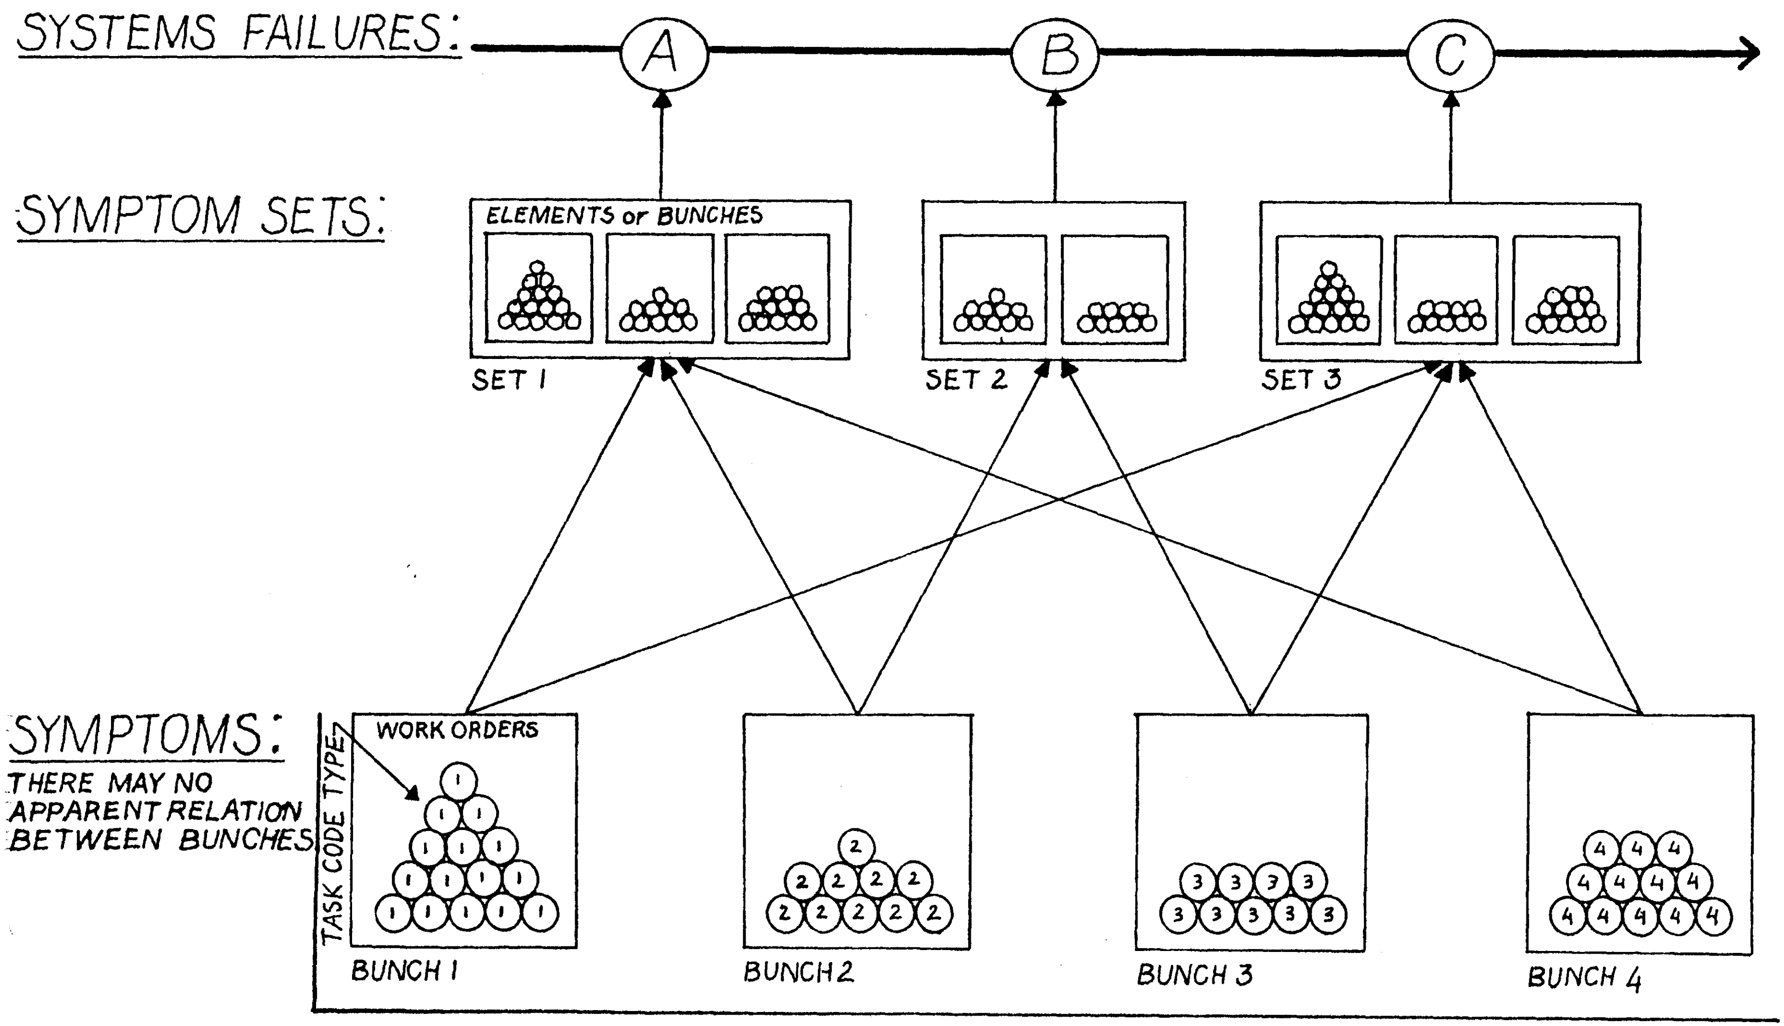

6) Summary & Conclusions

Our analysis of arrival rates therefore indicates that

1) A variety of non-random problems are occurring;
2) These failures are associated but are not necessarily causing other failures to occur;
3) Underlying systems failures which may be causing diverse types of non-random problems are not always recognized in the field or by monthly reviews of work order data, but
4) they can be identified by constructing sets of symptoms which appear in the data as bunches of associated work orders. These sets can then be used to structure monthly reviews.
5) Data should be organized by the new work order processing system so that work orders can be sorted into bunches by task code, and
6) Living Unit Inspection data should be structured to further identify such problems by classifying the condition of components common to many apartments.

Finally, waiting time consequences for most arrival processes might not be approximated with acceptable accuracy by simple Poisson-based queueing models. It appears quite likely that work order arrivals are uniformly distributed, although other models may apply. Conversely, however, the degree to which arrival distributions change over time into Poisson processes may be used as a measure for evaluating improvements in building conditions. As systematic and associated failures are identified and serviced, one would expect arrivals of subsequent work orders to more closely follow Poisson distributions. In turn, this would enable more accurate projections to be made from simple queueing models. From this perspective, the modified and expanded LUI program and the task code scheme of the proposed work order processing system seem to be oriented in the right general direction.

CHAPTER IV

SERVICE TIMES

1) The Poisson Service Process

Service times measure the time actually spent performing maintenance tasks. These are usually very small in comparison with the time customers spend in queue (by far the largest component of turn-around times), but they play an important role in determining the lengths of those queues.

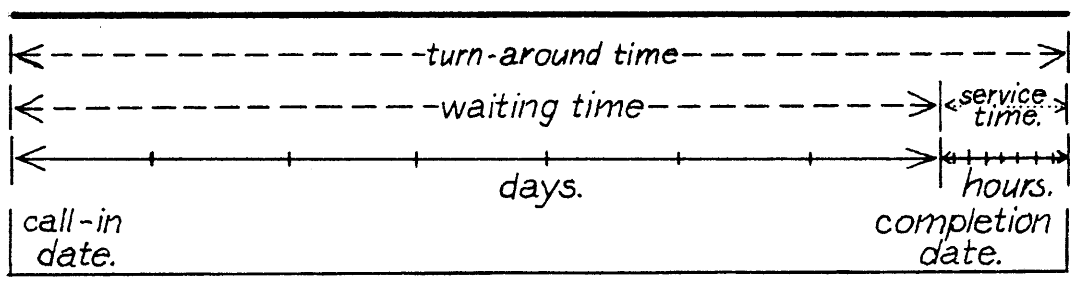

The strategy for estimating service time distributions is similar to the one we used for analyzing arrivals. There, we tested to see if the interarrival times follow an exponential distribution, which is equivalent to saying that the arrival rates could be modelled by a Poisson probability mass function. We assumed the interarrival times to be independent, and tested this assumption by measuring events per unit time with the aid of the pmf.

For a Poisson service process, however, it is the time between service completions (the service times) that are independent and which follow an exponential distribution. Instead of measuring probability distributions of events in a fixed time interval, we now measure probabilities that the time interval needed for an event to occur will assume some value. Therefore, we approximate a Poisson service process by calculating expected probabilities using the probability density function (pdf) for a negative exponential distribution, which is given by

$$\Pr\{st \in (t, t+dt)\} = ue^{-ut}\,dt \qquad t > 0$$

This is the set of probabilities that a service completion occurs in any given time interval, or, equivalently, the probabilities that service times are of a given duration.

As λ stood for the mean arrival rate of calls per day, here u equals the mean service rate per hour and t is the independent random variable describing the number of hours. Mean service times calculated from the data, however, are given by 1/u. Tables 4a and 4b present means and standard deviations for these service times by development, class, craft, and class-craft combination.[1]

Taking the inverse of the observed mean service time (1/u) for development 108 transforms it to a mean service rate (u), which can be substituted into the formula for the negative exponential distribution. Integrating the areas under the curve for each service rate then gives the expected probability density function for our hypothesized Poisson service process.

In table 1, development 108 work orders have been sorted by service times for all class and craft categories. The aggregate column

---

(1) For further reference, table 1 in Appendix IV shows the actual number of workorders in each service time category by class, craft, and class-craft. These figures are aggregated across all 12 developments.

table 1

Service times, by class & craft

#	hrs	wos	em	rou	lic	ski	emski	emlic	roulic	rouski
1	0.500	24	8	16	7	17	7	1	6	10
2	1.000	69	25	42	16	52	18	7	9	35
3	1.500	40	11	28	12	28	8	3	9	20
4	2.000	47	13	33	30	17	7	6	23	11
5	2.500	9	3	6	5	4	2	1	4	2
6	3.000	32	14	18	20	12	5	9	11	7
7	3.500	1	0	1	1	0	0	0	1	0
8	4.000	16	7	9	7	8	3	3	4	6
9	4.500	4	1	3	2	2	1	0	2	1
10	5.000	9	2	7	3	5	1	1	2	5
11	5.500	1	1	0	0	1	1	0	0	0
12	6.000	10	4	6	6	4	1	3	3	3
13	6.500	2	0	2	1	1	0	0	1	1
14	7.000	2	0	2	0	1	0	0	0	2
15	8.000	16	7	9	2	14	5	2	0	9
16	9.000	1	1	0	0	1	1	0	0	0
17	10.000	1	0	1	0	1	0	0	0	1
18	11.000	1	0	1	0	1	0	0	0	1
19	12.000	1	1	0	0	1	1	0	0	0
20	16.000	2	2	0	1	1	1	1	0	0

table 2

Service time
Probability density function
development 108

#	hrs	freq	frac	st_pdf
1	0.500	24	0.082	0.141
2	1.000	69	0.236	0.119
3	1.500	40	0.137	0.101
4	2.000	47	0.161	0.085
5	2.500	9	0.031	0.072
6	3.000	32	0.110	0.061
7	3.500	1	0.003	0.052
8	4.000	16	0.055	0.044
9	4.500	4	0.014	0.037
10	5.000	9	0.031	0.032
11	5.500	1	0.003	0.027
12	6.000	10	0.034	0.023
13	6.500	2	0.007	0.019
14	7.000	2	0.007	0.032
15	8.000	16	0.055	0.023
16	9.000	1	0.003	0.017
17	10.000	1	0.003	0.012
18	11.000	1	0.003	0.009
19	12.000	1	0.003	0.006
20	16.000	2	0.007	0.002
21	22.000	1	0.003	0.000
22	24.000	2	0.007	0.000
23	30.000	1	0.003	0.000

"wos" (all work orders) was then used to calculate the observed and expected probability distributions given in table 2.

2) <u>Methods for Testing Exponential Service Times</u>

A) <u>Linear Regression</u>

No matter how the arrivals are distributed, we would expect there to be a good chance that service times are exponential. Regression tests comparing the distributions in figures 1a and 1b are once more unreliable due to non-constant variances. A weighted least squares

figure 1a observed service time pdf
development 108

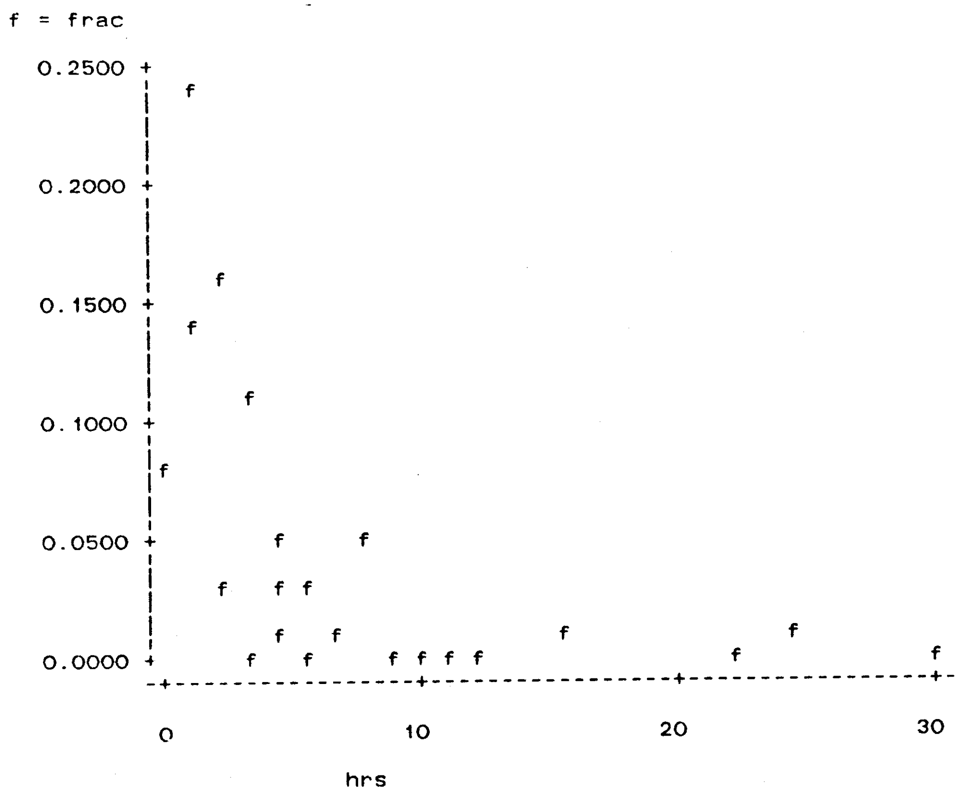

figure 1b expected exponential pdf

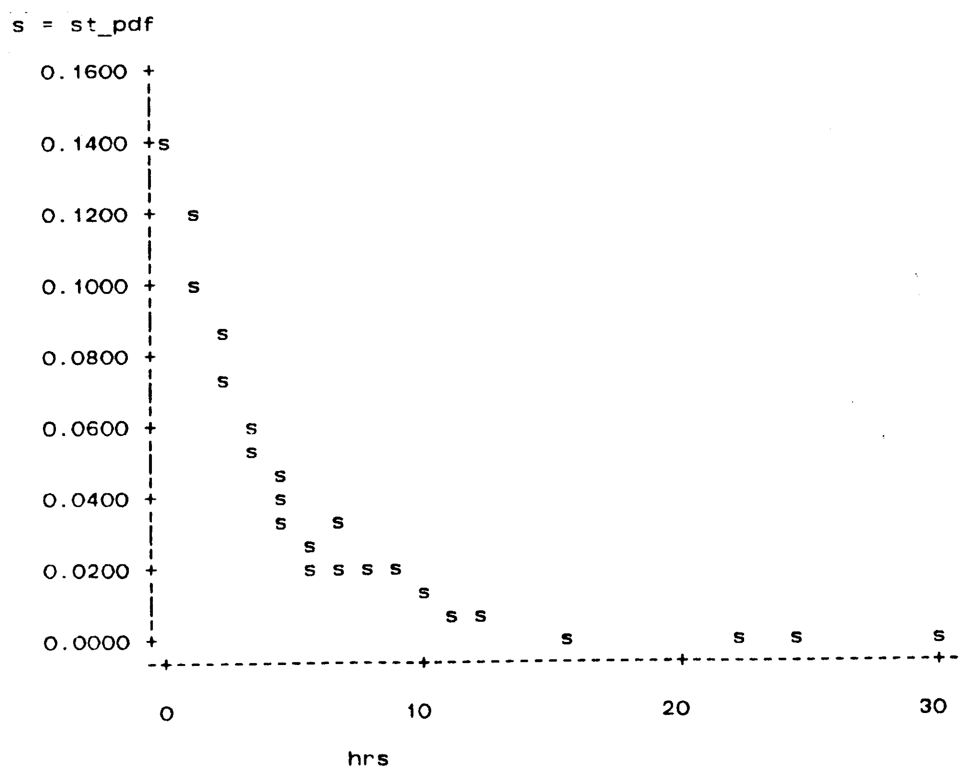

approach is again necessary if a regression test is to be used. Details of simple regression results are given in Appendix IV.

B) The Kolmogorov-Smirnov Test

Service time distributions can be reviewed, however, using the Kolmogorov-Smirnov goodness of fit test. We avoid the problem of small sample size here, because our sample is the number of work orders rather than the number of days. Results of the tests are given in table 3, and the calculations are in table 4 of the appendix. For all large sample sizes, the K-S statistic leads us to reject the exponential hypothesis.

C) Measures of Central Tendency

For arrival processes, we used the fact that Poisson means and variances are precisely equal as a final test of the degree to which the observed arrivals in development 108 are representative of those in the other 11. With the exception of emergency, emergency/licensed and routine/licensed jobs, we found the variances to be much higher than the means, - suggesting that uniform or hyperexponential distributions may be more accurate predictors of some arrival processes.

We can follow an analogous strategy as a final test for exponential service times. For development 108, a K-S test has provided reasonable evidence that service times are not exponentially distributed. A comparison of the means and variances may enable us to

Kolmogorov-Smirnov Goodness of Fit Test
for
Observed vs. Expected Poisson Service Time Distributions, by Class & Craft,
Development 108

Ho: observed = exponential
H1: observed $\neq$ exponential

Class/Craft Category	D stat.	Sample size	Critical Value		Decision
		n	a = .20	a = .05	
All work orders	.177	292	.063	.080	reject Ho
Emergency	.135	100	.107	.136	accept at .05
Routine	.189	188	.078	.099	reject
Licensed	.193	114	.100	.127	reject
Skilled	.216	174	.081	.103	reject
Emlic	.173	37	.175	.223	accept at .20
Emski	.185	62	.136	.173	reject
Roulic	.200	76	.123	.160	reject
Rouski	.228	109	.102	.130	reject

table 3

Kolmogorov-Smirnov Goodness of Fit Test
for Observed & Expected Service Time Distributions

table 4

Means & standard deviations of service times
by class, craft & development

dev	wos_m	em_m	rou_m	lic_m	ski_m	emlic_m	emski_m	roulic_m	rouski_m
101.000	2.512	2.353	2.574	1.969	2.891	2.227	2.448	1.873	3.051
103.000	3.397	3.083	3.456	2.890	3.856	3.015	3.235	2.822	4.034
105.000	3.095	3.165	2.930	3.467	2.922	3.613	2.867	2.357	3.010
108.000	3.009	3.020	3.037	2.842	3.089	3.257	2.863	2.651	3.271
109.000	4.027	3.148	8.263	3.100	4.629	2.704	3.450	5.107	10.104
114.000	3.429	3.868	3.020	4.034	3.176	4.262	3.557	3.438	2.942
120.000	2.470	1.784	5.885	1.821	3.125	1.825	1.739	1.794	9.045
123.000	4.185	5.826	3.987	3.122	4.906	4.077	8.100	2.952	4.632
124.000	2.220	2.216	2.224	1.790	2.533	2.500	2.080	1.745	2.586
501.000	3.753	6.944	3.637	2.478	4.620	14.250	1.250	1.924	4.752
508.000	3.965	1.917	4.119	2.134	5.252	1.778	2.333	2.185	5.342
510.000	2.173	2.083	2.795	2.636	1.934	2.716	1.760	2.125	3.179

dev	wos_s	em_s	rou_s	lic_s	ski_s	emlic_s	emski_s	roulic_s	rouski_s
101.000	3.907	3.042	4.181	1.386	4.921	1.817	3.745	1.191	5.282
103.000	4.663	4.204	4.803	2.905	5.739	4.078	4.378	2.617	6.104
105.000	3.972	4.434	2.623	5.731	2.840	6.063	2.921	1.651	2.734
108.000	3.575	2.978	3.890	2.843	4.006	2.876	3.072	2.841	4.496
109.000	6.843	5.204	11.146	2.599	8.455	2.019	6.478	4.063	13.451
114.000	5.516	5.803	5.229	3.543	6.148	3.768	7.033	2.892	5.565
120.000	5.281	1.964	11.674	2.184	7.112	2.267	1.591	1.668	14.848
123.000	5.322	5.680	5.258	2.562	6.475	2.216	7.884	2.596	6.306
124.000	3.208	2.219	3.282	2.102	3.789	2.374	2.178	2.083	3.919
501.000	7.757	11.690	7.579	4.009	9.374	15.354	0.645	1.228	9.524
508.000	7.201	1.165	7.439	2.304	8.995	1.121	1.443	2.430	9.115
510.000	2.018	1.814	3.062	2.191	1.888	2.232	1.469	1.959	3.555

tentatively extend these conclusions to other developments. The method is similar to the one employed for arrivals, but not identical, because service time distributions have been modelled using a pdf rather than a pmf. For a Poisson service process, the mean service time is $1/u$ and the variance is $1/u^2$. Therefore, the standard deviation (rather than the variance) is precisely equal to the mean, since

$$\sigma = \sqrt{\sigma^2} = \sqrt{\frac{1}{\mu^2}} = \frac{1}{\mu}$$

To test how closely the values in the columns of table 4a match those of table 4b, we use the same method as in comparing the means and variances of arrival rates. A first glance suggests that the means in each class and craft category are fairly close to their corresponding standard deviations. Results of the regressions are presented in table 5 of Appendix IV.

The r-square terms are uniformly higher than tests run for development 108 alone, as one would expect from regressions run on sample means. The slopes of lines fit by the regressions, however, are fairly close to one. Means and standard deviations are furthest from matching for skilled jobs, but we expect more variation here, due to the greater variety of job types in this category.

3) Interpretation & Conclusions

If we had found service times to be exponential, then the time remaining until the next service completion would be independent of the particular time at which we choose to view the service process. It would also be independent of the type of work order previously serviced. This is the Markov or "no memory" property of Poisson service processes. Exponential service times would therefore indicate that different types of repairs are being done within any class or craft category, since these repairs are taking different amounts of time to service. Such repairs would be unassociated and not consequences of systems failures. We might then expect any categories of Poisson arrivals to correspond to exponential service time categories.

But matters are more complicated than this. Two unrelated jobs may take the same time to be serviced, just as two associated jobs may have different service times. The fact that several workers may be needed to handle a given problem indicates that some jobs are interdependent and that the service times for these jobs are mixed together. Service interdependency means that queues can form on queues as some tasks are delayed in mid-service until help arrives from workers of another specialized skill needed to complete the task. Furthermore, this additional server may himself leave other jobs standing in a second queue while he is servicing minor tasks in the first. Given this interdependent operating method, Poisson arrival and service processes would not necessarily correspond.

Recording is another difficulty. "Bumps" appear in the service time distributions for jobs that take convenient time intervals to service (4 hours, 8 hours, 16 hours, 40 hours). These times may contain rounding errors and thus not correspond to theoretical durations of particular types of service. Rounding errors for the many short service times may be especially large relative to theoretical service times.

Despite these recording errors, systematic maintenance problems and job interdependency appear to be more important reasons we observe no mathematically neat distribution underlying service times.

There is a more important explanation, however, for the observed service time distributions than all those we have discussed so far. It is based on the fact that means and standard deviations do not merely correspond within class and craft categories; these measures of central tendency are also remarkably similar across those categories. Service times are on the whole much more regular or homogenous than the arrival rates we observed. This is an indication that maintenance crews are responding to calls by cosmetically treating diverse problems in similar ways.

Rather than undertaking projects aimed at the long term improvement of conditions, crews may be so used to operating in a "reactive mode" that the solutions they prescribe and apply merely alleviate symptoms in the short run where they should be diagnosing and treating underlying failures. Crews appear overworked and understaffed, yet this may indicate a need to reorganize service priorities and methods, rather than to expand personnel. The Living Unit Inspection program seems one way to begin to explore "initiatory" modes of operation. But it is important that these inspections not concentrate their efforts so heavily on apartment interiors that they are unable to devote attention to more global buildings systems.

Systematic problems often cannot be resolved through piecemeal reaction. A structure therefore must be found for the LUI program that avoids the tendency merely to react to scheduling routines, and encourages workers to initiate diagnoses. Because of their diverse nature, many systematic problems will not be isolated through data analysis, but can only be discovered in the field by craftspersons and supervisors who operate diagnostically.

The method for constructing system arrival sets outlined in the previous chapter can also be used to create corresponding service solution sets and to define system-oriented schedules that coordinate the diverse maintenance operations needed to upgrade undermaintained systems.

CHAPTER V

TURN-AROUND TIMES

1) Uses for the Data

For all jobs serviced in October, turn-around time (tat) data indicate how many working days each order spent in the queueing system from the day on which the request was called in through the day service was completed. Weekend days and holidays have been eliminated so that turn-around times more accurately reflect the actual working time required for the system to respond to demand. Of course, from tenants' point of view, weekend days are included in the time they must wait for service. But the inclusion of weekend days also exaggerates working day turn-around times by as much as 30 days for calls which are backlogged several months. Eliminating non-working days may therefore help us to better understand the relationship between actual system operations and turn-around times.

Had we been able to conclude in the previous chapter that arrival and service time distributions follow a Poisson process, then we could have modelled maintenance operations for development 108 as an M/M/n queueing system. This would have allowed us to compare hypothetical operating policies and to draw relatively specific conclusions, such as: "For policy x, 90% of emergency work orders will be serviced within 24 hours"; or "The expected average waiting time for routine calls is 5.5 days"; or "There is an 80% probability that w work orders or less will be in queue at time t".

(1)

Despite the fact that the observed arrival and service time distributions prevent us from making such specific projections, however, an analysis of observed turn-around times can reveal several important waiting time consequences which follow from the observed operating policy. First, we have fit a linear regression model showing the number of work orders that have a particular turn-around time. Second, we have used some general queueing equations applicable for nearly any arrival and service processes to estimate mean backlogs. Finally, a simple method has been illustrated for comparing hypothetical priority policies based on total waiting time calculations.

2) A Linear Regression Model for Observed Turn-Around Times

The regression model presented below indicates there is a linear relationship between the log of the turn-around time and the log of the number of work orders corresponding to a given time. (2)

$$\log(\text{wos}) = 6.5 - 1.3 \times \log(\text{tat})$$

$$\text{or,} \quad \text{wos} = e^{6.5}\,(\text{tat})^{-1.3}$$

---

(1) Any M/M/n queue is characterized by Poisson arrivals, exponential service times and n identical servers.

(2) See Appendix V for detailed results.

This equation defines the number of work orders which spent t days in the queueing system. Because both variables in the regression equation have been logged, the resulting linear relationship can be interpreted as a constant elasticity. In other words, a 1% change in the number of work orders is accompanied by a 1% change in the length of the turn-around time.

This equation can be used to compare October turn-around times with those following from recent operations changes. However, reductions in routine times following these changes will be exaggerated by the fact that routine work orders are now generated solely by the Living Unit Inspection program.

figure 1

Fitted and actual relationships
between
work orders and turn-around time

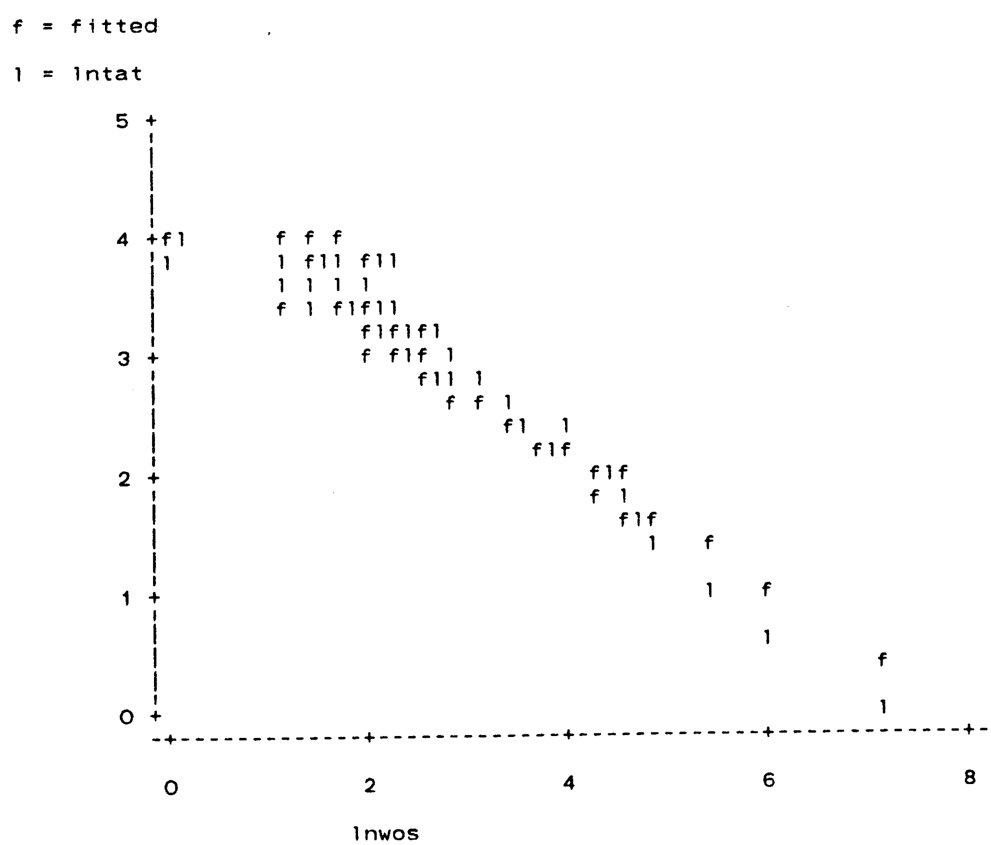

Figure 1 shows how closely the fitted regression line corresponds to the data. The equation also holds for all classes and crafts of work orders (see Appendix V) and suggests that over all developments there is no appreciable difference between emergency and routine turn-around times.

In a truly priority-based queueing system, newly arriving routine calls would not be serviced until all emergencies had been handled. The data, however, are more representative of a maintenance system in which separate queues exist for emergency and routine work orders. If for example one crew services only emergencies, while another handles only routine calls, then both queues may be equally congested.

In addition, BHA staff have noted the difficulty of comparing mean turn-around times in any one development or month with those of another. The difficulty arises because those crews operating quickly in any month (or receiving fewer calls) will begin to service backlogged orders toward the end of that month. Because these orders have been in the system a long time, the crew's mean turn-around time will be driven upward by the act of servicing them. Reliance on such statistics for comparing month-to-month performance creates incentives to "wish away" or purge backlogged orders.

The resulting confusion (table 1) is not due to the structure of the data, however, but rather to the operating method in use. By this method, work orders which have already spent a month in queue are sent back to the end of the line, where they are ignored until the system is uncongested enough to deal with them. Service priority is above all given to work orders arriving in the current month, rather than to emergency calls per se. Only if older work orders cease to be sent to

table 1
Means & standard deviations of turn-around times
by class, craft & development

#	dev	tat_m	tat_s	em_m	em_s	rou_m	rou_s	lic_m	lic_s	ski_m	ski_s
1	101	4.327	8.071	2.653	3.192	4.897	9.107	4.538	5.113	4.197	9.595
2	103	18.810	30.493	42.608	36.959	13.456	26.099	19.299	29.289	17.748	31.000
3	105	8.859	11.788	8.258	10.785	10.211	13.786	4.729	9.152	10.794	12.403
4	108	8.788	17.969	6.646	18.747	10.051	17.657	7.500	10.094	9.729	21.545
5	109	13.014	27.920	13.854	29.912	8.486	14.934	12.062	38.626	13.593	19.242
6	114	12.418	39.932	5.238	21.982	19.430	50.970	3.544	9.157	16.894	47.999
7	120	7.804	19.363	6.853	18.757	12.462	21.762	6.548	21.265	9.061	17.257
8	123	10.615	17.196	3.905	9.049	11.369	17.737	10.585	19.566	10.635	15.543
9	124	8.096	21.471	12.258	22.429	7.795	21.420	7.104	21.855	8.815	21.201
10	501	11.938	26.641	4.222	6.704	12.301	27.167	8.753	26.524	14.200	26.758
11	508	14.560	36.202	3.000	3.715	15.381	37.323	5.478	12.586	21.021	45.154
12	510	11.071	20.540	11.277	21.747	9.682	9.073	8.661	19.308	12.351	21.139

#	dev	emlic_m	emlic_s	emski_m	emski_s	roulic_m	roulic_s	rouski_m	rouski_s
1	101	2.333	2.436	2.893	3.685	5.351	5.598	4.625	10.854
2	103	31.242	33.449	49.045	37.492	17.269	28.125	9.931	23.614
3	105	4.615	9.232	10.750	11.114	5.571	9.181	10.860	14.261
4	108	4.758	7.604	7.726	22.662	8.886	10.914	11.048	21.164
5	109	13.239	42.077	14.255	19.503	6.000	6.069	9.833	18.014
6	114	2.262	3.507	8.214	30.785	7.133	16.754	22.028	55.310
7	120	5.602	21.609	8.172	15.193	12.000	18.808	12.818	24.230
8	123	1.182	0.603	6.900	12.749	12.042	20.662	10.957	15.765
9	124	11.700	19.345	12.524	24.207	6.852	22.004	8.502	20.990
10	501	1.000	0.000	8.250	8.995	9.118	27.093	14.469	27.198
11	508	3.000	4.036	3.000	3.464	5.803	13.293	21.596	45.755
12	510	8.451	20.488	12.763	22.341	10.000	9.381	9.500	9.247

the back of the line will meaningful reductions in mean turn-around times appear in the data after a period of several months.

3) Total Waiting Times & the Number of Tenants in Queue

The question "How quickly is maintenance service responding to demand?" can be posed in another way: "How many tenants are waiting for service, and how long have they been waiting?". When we review turn-around time statistics, it is as if we observe each work order as it is leaving service. We know how long these work orders have spent in the system, but not how many are currently waiting (figure 2).

figure 2

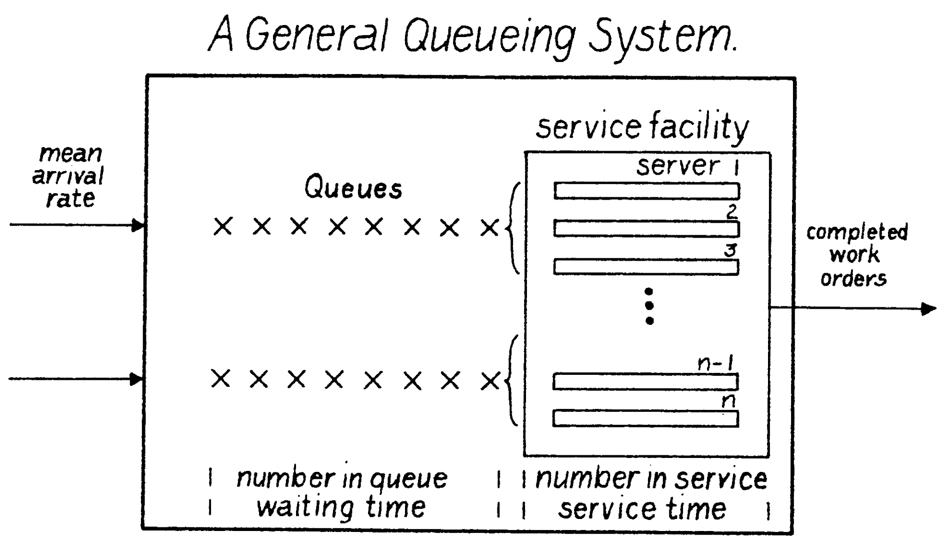

Fortunately, BHA Monthly Management Reports contain estimates of the number of outstanding work orders by development. The size of these backlogs does not change dramatically from one month to the next and suggests that the system is approximately in a steady state. This means that the number of tenants in the queueing system remains roughly constant over time. A general relationship in queueing theory known as "Little's formula" can be used to compare BHA backlog estimates with those expected from a steady state system with the same arrival and turn-around time statistics.

Little's formula $L = \lambda W$

Here L is the expected steady state number of work orders in the queueing system, λ is the mean arrival rate (Poisson or not), and W is the expected steady state time in the system for a given work order, or the mean turn-around time.

table 2 Estimated Tenants in the Queueing System

Dev	Arrival Rate λ	x	Mean Tat W	=	Estimate w/Little's Formula L	BHA Estimate
101	8.15		4.33		35	352
103	14.40		18.81		271	671
105	7.35		8.86		65	154
108	11.10		8.79		98	140
109	8.15		13.01		106	219
114	7.75		12.42		96	211
120	10.15		7.80		79	65
123	7.65		10.62		81	284
124	20.00		8.10		162	627
501	8.90		11.94		106	294
508	6.95		14.56		101	143
510	6.75		11.07		75	77

Calculations of L based on Little's formula underestimate the number of outstanding work orders as given in Monthly Management Reports (table 2). More tenants are actually waiting than we would expect to find. The problem is not that mean turn-around times have been skewed upward by the servicing of backlogged orders in any month, but rather that they fail to reflect the higher turn-around times associated with those orders still unserviced. Real mean turn-around times are higher than they appear in the data, and therefore actual backlogs are larger than those estimated by Little's formula.

One reason actual turn-around times are so long (and backlogs so large) is that some older work orders are repeatedly being sent to the back of the line. But this may not account for all of the difference between L and the BHA estimate. Another reason may be that many repair jobs involve several work orders (and servers), since more than one skill may be required to perform the repair. For such interdependent jobs, servers in one queue may be kept waiting for related or prerequisite orders to be serviced by those in another. There is also then a queueing process operating for these servers, many of whom may be kept "busy" waiting in the queue.

For a steady state queueing system,

$$\rho = \frac{\lambda}{ku} < 1 \qquad k = \text{number of servers}$$

must be true. If ρ (rho) is greater than one, there will be no steady state, since calls are arriving faster than servers are theoretically able to handle them. In "unstable" queues of this type, even minor increases in the arrival rate lead to major congestion. Because wild fluctuations in backlog are not occurring, we assume that a steady state exists and that the number of servers is theoretically adequate. A good deal of the actual backlog may therefore be due to the existence of tandem or interdependent queues. In order for turn-around times to be reduced, the interdependencies leading to "server queueing" also must be reduced. The recent reclassification of many tasks as "neutral" should certainly improve the situation, but the "neutral" job category may have to be expanded to include a wider range of tasks. It is important to know which tasks contribute most heavily to queue interdependency. These will appear as "bunches" of work orders in monthly data.

We have seen how the operating policy with respect to backlogs and the delays due to interdependent servicing account for the low estimate of L given by Little's formula. In addition, the fact that several work orders may be generated for one repair job means that BHA backlog figures provide too high an estimate of the steady state number of apartments (tenants) in the queueing system. If only one work order were generated per job, and if backlog policy were changed and queue interdependency eliminated, then Little's formula should provide an accurate estimate of those in the system. To find the steady state number in queue (L_q), we subtract the number in service on any day (λ/u) from L. Here u has been transformed to a daily service rate.

$$L_q = \lambda W_q = \lambda\left(W - \frac{1}{u}\right) = \lambda W - \frac{\lambda}{u} = L - \frac{\lambda}{u}$$

In a one-order-per-job operating system, L_q would also equal the steady state number of outstanding work orders. The waiting times for such a system can be summed to give the total time tenants spend in queue. Various queueing disciplines (priority policies) can then be compared with respect to total tenant waiting times.

4) <u>Comparing Priority-of-Service Policies</u>

A method for comparing queueing disciplines can be briefly illustrated using two hypothetical priority policies. One policy is

FIFO (First In, First Out). A variant of a FIFO system with nonpreemptive priority for emergency calls might be described as follows: There are two basic queues - licensed and skilled - each of which has an emergency and a routine component. Newly arriving emergency calls queue up behind those emergencies already waiting, while new routine arrivals take their places behind other routine calls. No routine jobs are serviced until all emergencies have been handled, but new emergencies do not interrupt the servicing of a routine order which has already begun (figure 3).

figure 3

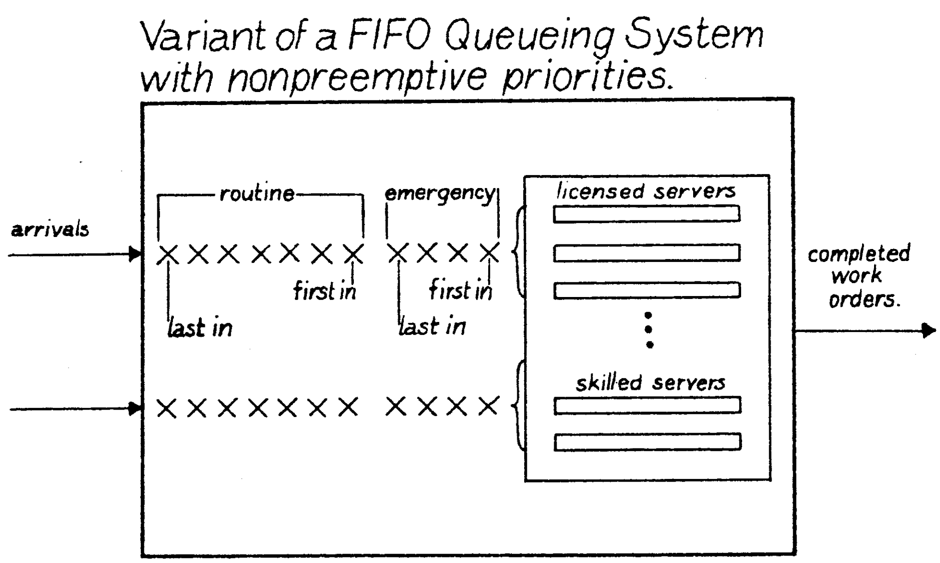

A second system, based on "Shortest Expected Processing Time" (SEPT), is similar in many ways to the one just outlined. Again, there might be 2 (or n) queues, each organized so that emergency orders continue to have nonpreemptive priority over routine calls. But within each emergency or routine sub-queue, those jobs are serviced first which have the shortest expected processing (service) time. On any day, this priority is assigned by the maintenance supervisor without considering how long a given call may have already waited for service (figure 4).

figure 4

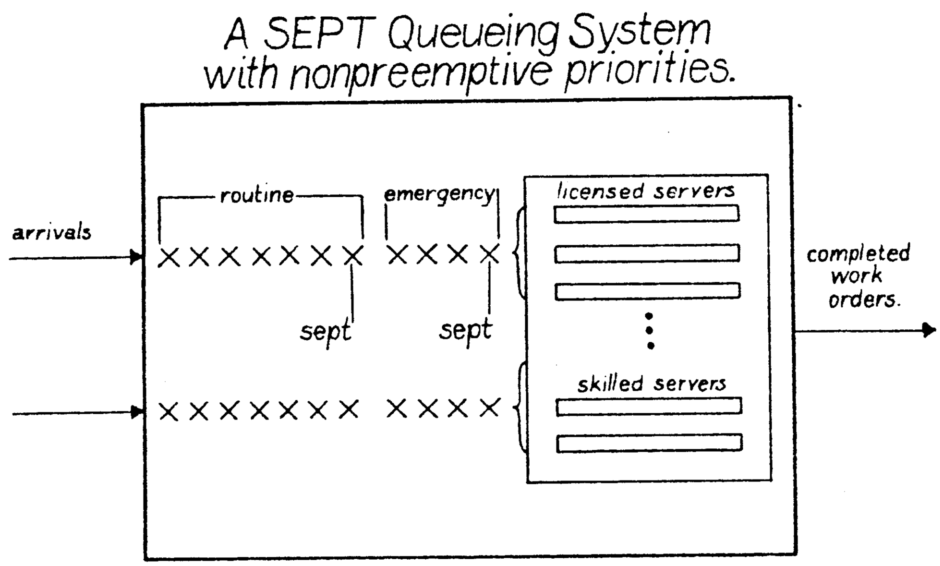

The FIFO and SEPT queueing systems can be compared with respect to the total time tenants spend waiting in each system. Using monthly data, a FIFO work order profile can be created by sorting calls at any development according to craft (licensed or skilled) and then in ascending order by arrival date. Note that "class" distinctions have been ignored, since knowing when the emergency queue is empty involves a more complex analysis, while the simplification used here does not affect a comparison of the two policies.

To make the comparison, each work order is assigned a waiting time, which is the sum of the service times for those work orders ahead of it in queue. The total time tenants spend in the system is simply the sum of these waiting times. In the following tables, imaginary data have been used to illustrate the method.

FIFO Waiting Times for Licensed Calls

WO No.	Arrival date	Service time (hours)	Waiting time (hours)
1	1	10	0
2	1	3	10
3	1	5	13
4	1	2	18
5	1	1	20
6	1	2	21
			23

Total Waiting Time = 105 hours (13.1 days)

For the SEPT system, work orders for each arrival date can be further sorted in ascending order of expected processing time.

SEPT Waiting Times for Licensed Calls

WO No.	Arrival date	Service time	Waiting time
5	1	1	0
4	1	2	1
6	1	2	3
2	1	3	5
3	1	5	8
1	1	10	13
			23

Total Waiting Time = 53 hours (6.6 days)

In practice, work orders are continuously arriving that have shorter service times than some which may have just been serviced. Therefore, the total waiting times calculated by this method are only approximations of those we would actually witness. Nevertheless, a comparison of these hypothetical policies provides reasonable evidence that a queueing discipline based on Shortest Expected Processing Time would be preferable to a simple FIFO system in terms of reducing the total time tenants spend waiting for service. For equity reasons, supervisors might want to put an upper limit on the time a work order may be delayed due to its long expected processing time.

Other queueing disciplines could be evaluated in a similar fashion by comparing the total waiting times implied by each policy.

APPENDIX I

Key to Tables in the Appendices

Appendix III

Tables 1,2 & Figure 1 :

- wos : work orders. This is the number of work orders arriving per day.
- em : emergency work orders
- rou : routine " " "
- lic : licensed " " "
- ski : skilled " " "
- emlic : emergency/licensed work orders arriving per day
- emski : emergency/skilled " " " " " " "
- roulic : routine/licensed " " " " " " "

- freq : frequencies or days on which n arrivals occurred
- frac : the observed probability mass function, or the fraction of days on which n arrivals occurred
- Pr_n : the "expected" Poisson pmf. The observed mean arrival rate used to calculate this distribution is the weighted mean of the orders arriving times the number of days on which n arrive. The expected pmf is therefore a likely profile of the observed distribution if the observed is Poisson.

Table 3 :

- The calculations used for this test are pmfs given in table 7.

Table 4,5 & Figure 2 :

- Symbols are as in table 1 above. Also :

- cumPr_n : the cumulative expected pmf, or cumulative distribution function
- cumfrac : the cumulative observed pmf

- allKS : For all arrivals, this is the difference between cumPr_n and cumfrac, whose maximum value used as the D estimate in the K-S test.
- eKS : same as allKS, but for emergency arrivals
- rKS : " " " " " routine " "
- lKS : " " " " " licensed " "

.
.
.

- rsKS : " " " " " routine/skilled arrivals

Tables 6,7 :

- wos : arrivals per day for all work orders
- un20, un40, 60, 100 : A set of 20, 40, 60, 100 uniformly distributed random frequencies (days) on which n arrivals would occur.
- frac20, 40, 60, 100 : Random uniform pmf based on un20, 40, ...100
- uncdf20, 40, ...100 : uniform cdfs based on frac20, ...100
- poi_pmf20, 40, ...100 : Expected poisson pmfs were generated using the mean arrival rate implied by the uniform istribution. Because they have the same mean, we are testing how sensitive the regression and K-S tests are to the forms the distributions take.
- poi_cdf20, ...100 : cdfs based on the poi_pmf20, ...100
- D20, ...100 : Differences between the uncdf20, ...100 and poi_cdf20, ...100. The maximum in this set is the D statistic used to test K-S sensitivity at sample sizes 20 to 100.

<u>Appendix IV</u>

Table 3 :

- hrs : service times in hours
- emfrac : observed pdf for emergency service times
- rfrac : " " " " routine " " "

  . .
  . .
  . .
- rsfrac : " " " " routine/skilled service times

- epdf : expected pdf for emergency service times
- rpdf : " " " " routine " " "

  . .
  . .
  . .
- rspdf : " " " " routine/skilled service times.

Table 4 :

- cepdf : cumulative expected pdf for emergency service times
- cefrac : " " observed " " " " " " "
- eks : differences between cepdf and cefrac

134817

BOSTON HOUSING AUTHORITY MAINTENANCE WORK ORDER

D E V | U N I T | M M D D Y Y | TENANT LAST NAME | TEL NO.

REQUESTER — IF DIFFERENT FROM ABOVE | TEL NO.

DEV. NAME: ______________________

ADDRESS: ______________________

FLOOR: ______________ UNIT SIZE: ______ BR.

(SEE REVERSE)

CLASS

CRAFT

ERS?✓

CLERK I.D. # | INITIALS

ASSIGNED:

DATE | TIME

EST. HRS. | EST. LABOR | EST. MATERIAL

DESCRIBE WORK INCLUDE AREA

STOCK NO.	ITEM	QUAN.	UNIT #	EXTENDED
			TOTAL	

TENANTS PERMISSION TO ENTER OR PERMISSION RECEIVED BY

DATE NOT HOME | DATE NOT HOME | DATE — VOID

CAUSE: (SEE REVERSE)

CRAFTSMAN'S WORK DESCRIPTION

EMPLOYEE NO.	NAME	HRS.	MINS.	CODES

TENANT ACCEPTANCE | M M D D

COMPLETED BY | M M D D

INSPECTED BY | M M D D

SUPER I.D. NO. | SUPERVISOR | M M D D

BLUE: WORKMAN/EDP | WHITE: WORKMAN/SUPER | YELLOW: TENANT | PINK: W.O. CLERK

Work Order Form, Boston Housing Authority, October 1983

table 1

APPENDIX II

table 2

Sample Work Order Data
Boston Housing Authority
October 1983

wo

#	dev	class	craft	mats	labor	total_cost	ci_day	a_day	c_day	hrs	tat
1	101	1	1	0	1948	1948	298	300	300	2.500	3
2	101	1	1	0	1169	1169	277	300	300	1.500	18
3	101	5	1	0	1169	1169	295	302	302	1.500	6
4	101	1	1	0	1169	1169	300	302	302	1.500	3
5	101	5	1	1962	12540	14502	271	274	278	12.000	6
6	101	1	1	0	779	779	299	300	300	1.000	2
7	101	1	1	0	1169	1169	299	300	300	1.500	2
8	101	1	1	0	1169	1169	286	300	300	1.500	11
9	101	99	1	0	1169	1169	300	302	302	1.500	3
10	101	1	1	0	1558	1558	277	279	279	2.000	3
11	101	1	1	0	1948	1948	279	277	279	2.500	1
12	101	1	1	0	779	779	294	302	302	1.000	7
13	101	5	1	0	3116	3116	279	279	279	4.000	1
14	101	5	2	0	8670	8670	305	305	305	7.500	1
15	101	1	3	250	490	740	278	278	278	0.500	1
16	101	1	3	1390	980	2370	298	298	298	1.000	1
17	101	5	3	125	490	615	271	281	281	0.500	9
18	101	5	3	125	490	615	271	283	281	0.500	9
19	101	5	3	388	980	1368	279	279	279	1.000	1
20	101	4	3	0	1256	1256	264	264	277	1.000	10
21	101	5	3	0	0	0	281	281	281	0.500	1
22	101	4	3	388	490	878	278	281	281	0.500	4
23	101	5	3	0	490	490	281	281	281	0.500	1
24	101	4	3	388	980	1368	302	302	302	1.000	1
25	101	4	3	388	490	878	278	278	278	0.500	1
26	101	1	3	250	1470	1720	298	286	299	1.500	2
27	101	1	3	388	490	878	302	302	302	0.500	1
28	101	5	3	388	490	878	300	300	300	0.500	1
29	101	5	3	388	490	878	302	302	302	0.500	1
30	101	5	3	0	980	980	264	278	278	1.000	11

APPENDIX III

1) Linear Regression as a Test of Probability Distributions

A simple regression test was devised to observe how closely the two pmfs correspond by using one distribution as the independent variable or predictor of the other. We assumed the two sets of probabilities would match perfectly if the test accounted for all of the variation by fitting a line of slope=1, passing directly through the origin.

Since we did not precisely know the extent of errors in the data, the tests were run in several ways. First, the observed pmf was used as the independent variable to explain variations in the expected probabilities.

$$\Pr\{\text{expected}\} = B0 + B1 \times \Pr\{\text{observed}\} \quad (1a)$$

This is a regression of the "expected on the observed" probabilities. The relationship was then turned around to test how well the expected Poisson pmf explained the observed variation.

$$\Pr\{\text{observed}\} = B0 + B1 \times \Pr\{\text{expected}\} \quad (1b)$$

Although the fraction of total variation explained (r-square) is the same in each case, the slopes (B1) and y-intercepts (B0) may differ. We controlled the y-intercepts by eliminating the constant terms from the equations and forcing the fitted lines to pass through the origin. This second round of tests gave both different slopes and different r-square values.

$$Pr\{expected\} = B1 \times Pr\{observed\} \qquad (2a)$$

$$Pr\{observed\} = B1 \times Pr\{expected\} \qquad (2b)$$

A comparison of the results from each of these four regressions was initially used to indicate whether all arrivals to development 108 taken as a group could be approximated by a Poisson model. The test results were

$$Pr\{exp\} = 0.07 - 0.11 \times Pr\{obs\} \quad r^2 = 0.01 \qquad (1a)$$

$$Pr\{obs\} = 0.09 - 0.10 \times Pr\{exp\} \quad r^2 = 0.01 \qquad (1b)$$

$$Pr\{exp\} = 0.56 \times Pr\{obs\} \quad r^2 = \text{error} \qquad (2a)$$

$$Pr\{obs\} = 0.96 \times Pr\{exp\} \quad r^2 = \text{error} \qquad (2b)$$

Development 108 arrivals certainly did not appear to be Poisson overall.[1]

We thought that arrivals for some classes and crafts may have been more closely approximated by the Poisson model, however. Table 1 summarizes regression results for each of the four equations by class, craft, and class/craft combination. The observed and expected probabilities for each category are given in table 2. The distributions are plotted in figures 1a through 1d.

All four series of regressions indicated that only emergency, licensed, emergency/licensed and emergency/skilled arrivals might

---

(1) In cases for which the correlation of the dependent and independent variables is extremely small, the "rgop" program in the Consistent System sometimes yields negative values for r-square when the constant term is forced out of the model.

follow nearly Poisson patterns. All other classes and crafts appeared to have non-Poisson arrivals, as indicated both by their low r-square values and slopes nearer to zero than to one. But how reliable were these results?

We then tested the regression test itself by generating a distribution of random uniform probabilities and using regression to compare them with a Poisson distribution. To calculate our Poisson values, we used the mean arrival rate implied by the corresponding uniform probabilities as our estimate of lambda. If the test showed no correlation between the distributions (as we would expect), then we would have had some evidence that the regression method used was accurate enough for our purposes.

The results of these tests are given in table 3. They indicate that regression is an inadequate tool for comparing noncumulative probability distributions based on small samples. This explains why our finely sliced categories appeared to be nearer to Poisson distributions. For larger samples, the test distinguishes between uniform and Poisson distributions, but not between the Poisson and the normal. In addition, regressions using cumulative distributions were shown to be extremely insensitive to any differences between distributions.

Most important, however, the non-constant variances in the distributions suggest the need for a weighted least squares approach if regression models are to be used. Unfortunately, no time was available to execute this test. On the basis of simple regression tests, we could therefore neither accept nor reject the hypothesis that observed arrivals are Poisson distributed.

table 1
Linear Regressions of Observed Arrivals on Expected Poisson Arrivals
by class & craft
development 108

Linear Regressions of Observed Arrivals on Expected Poisson Arrivals

Class/Craft Category	Expected on Observed			Observed on Expected			Exp. on Obs.		Obs. on Exp.	
	r^2	B0	B1	r^2	B0	B1	r^2	B1	r^2	B1
All work orders	.01	.07	.11	.01	.09	-.10	error	.56	error	.96
Emergency	.57	.02*	.63	.57	.03*	.91	.54	.74	.52	1.13
Routine	.00	.08	.03	.00	.08	.04	error	.69	error	.75
Licensed	.54	.02*	.71	.54	.04*	.76	.53	.81	.44	1.03
Skilled	.10	.04*	.40	.10	.08	.24	.03	.71	error	.92
Emlic	.90	-.03*	1.07	.90	.04*	.84	.89	.96	.85	1.01
Emski	.53	.02*	.83	.53	.06*	.64	.52	.92	.36	.94
Roulic	.07	.08*	.50	.07	.14	.15	error	.93	error	.90
Rouski	.10	.07*	.34	.10	.08	.29	error	.83	error	.85

* - insignificant B0 coefficient, indicating the fitted line may pass through the origin.

table 2 Observed & Expected Poisson Probability Mass Functions by Class & Craft, Development 108

wos

wos	freq	frac	Pr_n
3	1	0.050	0.003
5	3	0.150	0.021
6	1	0.050	0.039
7	2	0.100	0.062
9	1	0.050	0.107
11	2	0.100	0.119
12	2	0.100	0.110
13	1	0.050	0.094
15	1	0.050	0.055
16	2	0.100	0.038
17	3	0.150	0.025
18	1	0.050	0.015

em	freq	frac	Pr_n
1	3	0.150	0.061
2	3	0.150	0.129
3	3	0.150	0.183
4	6	0.300	0.194
6	1	0.050	0.117
7	1	0.050	0.071
8	1	0.050	0.038
10	1	0.050	0.008
12	1	0.050	0.001

rou	freq	frac	Pr_n
2	1	0.050	0.029
3	5	0.250	0.063
4	1	0.050	0.105
5	2	0.100	0.140
6	2	0.100	0.155
7	1	0.050	0.148
8	1	0.050	0.123
9	2	0.100	0.091
10	1	0.050	0.060
11	2	0.100	0.036
12	1	0.050	0.020
13	1	0.050	0.010

lic	freq	frac	Pr_n
1	3	0.150	0.065
3	4	0.200	0.188
4	5	0.250	0.195
5	2	0.100	0.162
6	1	0.050	0.112
7	1	0.050	0.066
8	2	0.100	0.034
9	1	0.050	0.016
0	1	0.050	0.016

ski	freq	frac	Pr_n
1	1	0.050	0.008
2	3	0.150	0.027
4	2	0.100	0.101
5	1	0.050	0.137
6	2	0.100	0.154
7	3	0.150	0.148
8	3	0.150	0.125
11	3	0.150	0.039
12	1	0.050	0.022
13	1	0.050	0.011

emlic	freq	frac	Pr_n
1	7	0.350	0.329
2	4	0.200	0.255
3	3	0.150	0.132
7	1	0.050	0.001
0	5	0.250	0.212

emski	freq	frac	Pr_n
1	5	0.250	0.187
2	4	0.200	0.248
3	4	0.200	0.219
4	1	0.050	0.145
5	3	0.150	0.077
9	1	0.050	0.001
0	2	0.100	0.071

roulic	freq	frac	Pr_n
1	3	0.150	0.199
2	4	0.200	0.254
3	4	0.200	0.216
4	2	0.100	0.138
5	4	0.200	0.070
0	3	0.150	0.078

rouski	freq	frac	Pr_n
1	3	0.150	0.056
2	2	0.100	0.122
3	5	0.250	0.177
4	2	0.100	0.193
5	1	0.050	0.168
6	2	0.100	0.121
7	2	0.100	0.075
8	2	0.100	0.041
10	1	0.050	0.009

Column 1 = number of arrivals
Column 2 = freqencies (days)
Column 3 = "frac" = observed probabilities
Column 4 = "Pr n" = expected Poisson pmf

figure 1a

#	em	freq	frac	Pr_n_arr
1	1	3	0.150	0.061
2	2	3	0.150	0.129
3	3	3	0.150	0.183
4	4	6	0.300	0.194
5	6	1	0.050	0.117
6	7	1	0.050	0.071
7	8	1	0.050	0.038
8	10	1	0.050	0.008
9	12	1	0.050	0.001

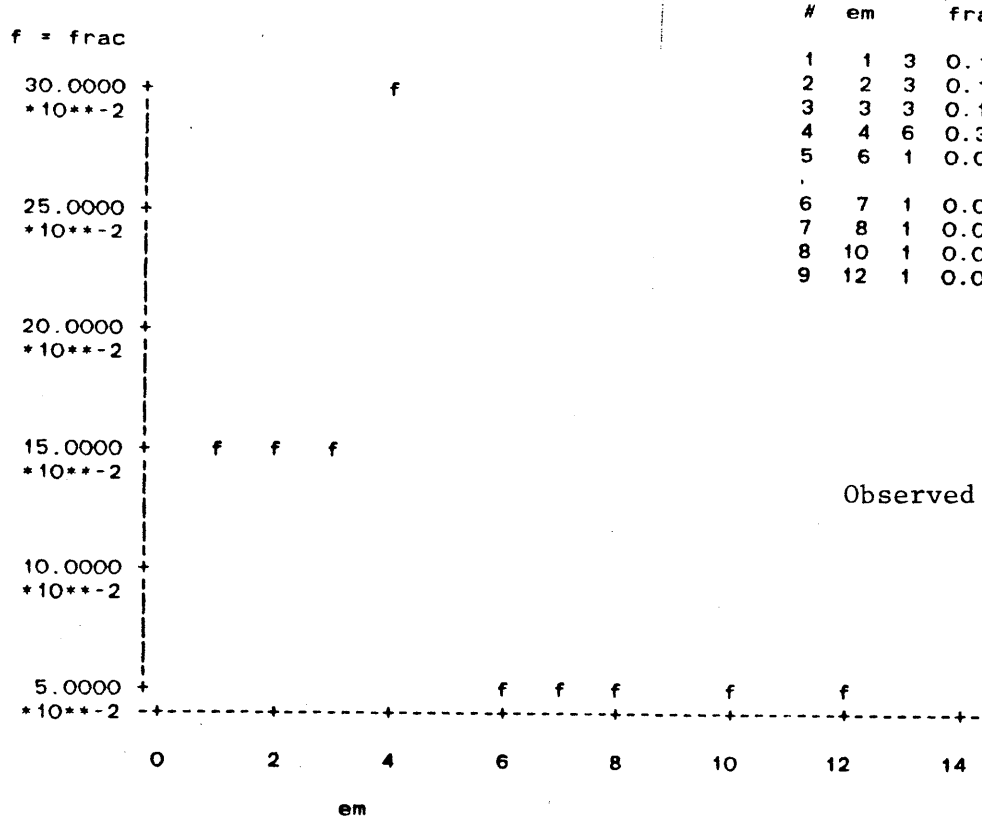

Observed pmf for emergencies

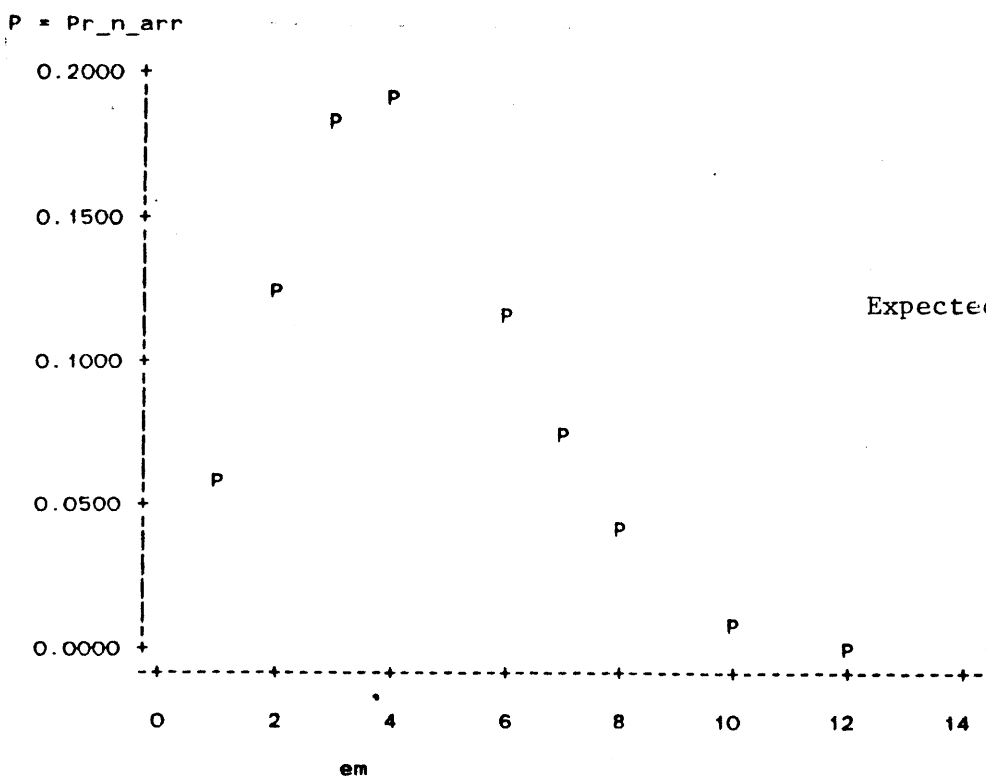

Expected Poisson pmf

figure 1b

#	rou	freq	frac	Pr_n
1	2	1	0.050	0.029
2	3	5	0.250	0.063
3	4	1	0.050	0.105
4	5	2	0.100	0.140
5	6	2	0.100	0.155
6	7	1	0.050	0.148
7	8	1	0.050	0.123
8	9	2	0.100	0.091
9	10	1	0.050	0.060
10	11	2	0.100	0.036
11	12	1	0.050	0.020
12	13	1	0.050	0.010

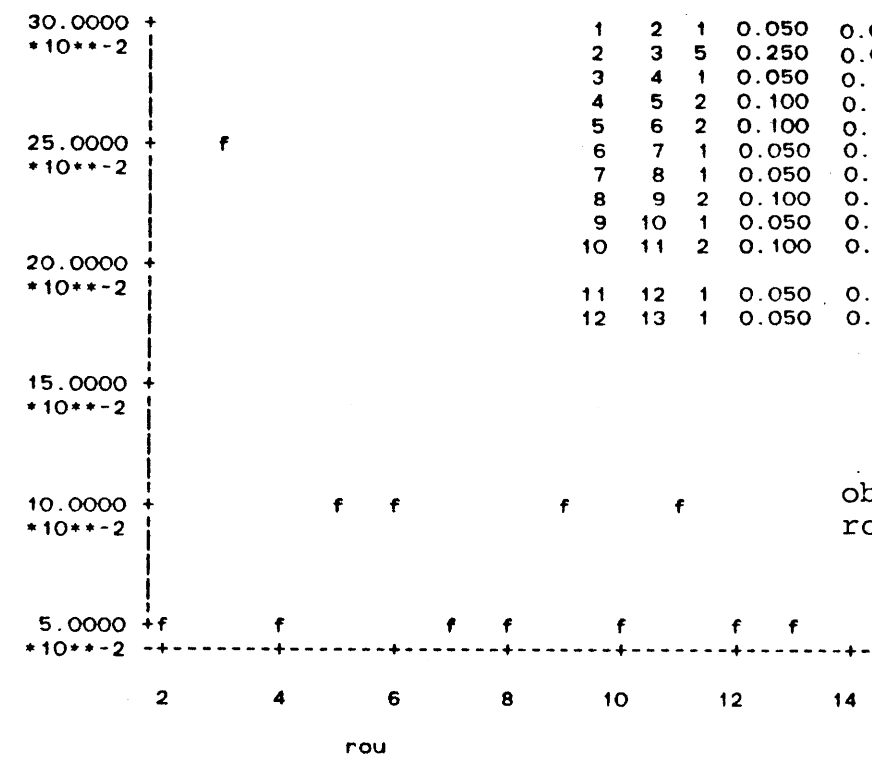

observed pmf for
routine work orders

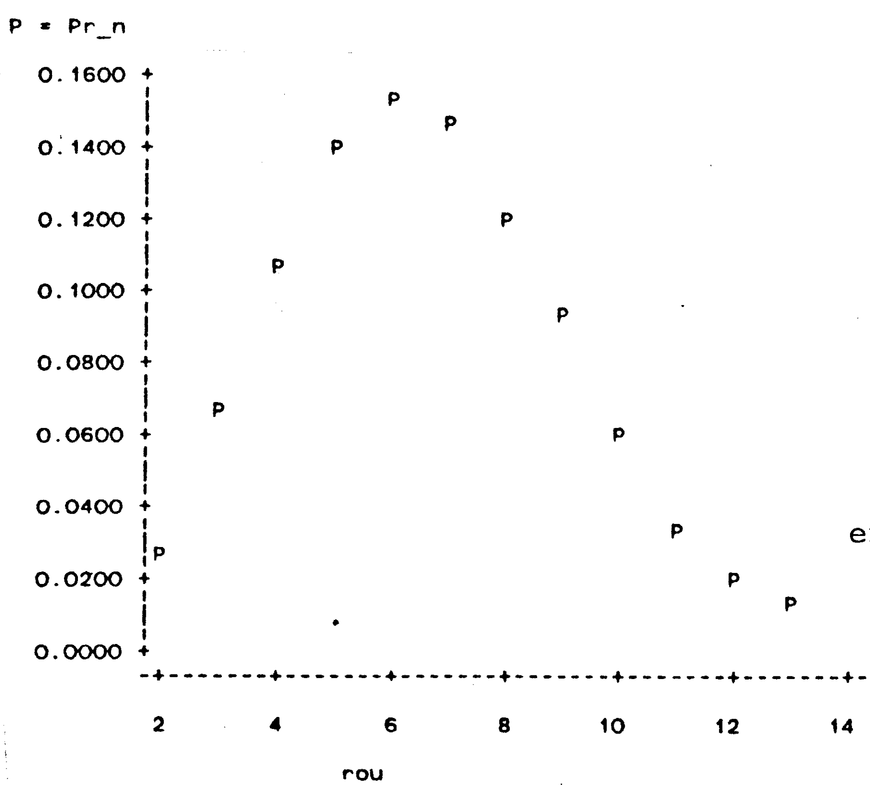

expected Poisson pmf

figure 1c

#	lic	freq	frac	Pr_n
1	1	3	0.150	0.065
2	3	4	0.200	0.188
3	4	5	0.250	0.195
4	5	2	0.100	0.162
5	6	1	0.050	0.112
6	7	1	0.050	0.066
7	8	2	0.100	0.034
8	9	1	0.050	0.016
9	0	1	0.050	0.016

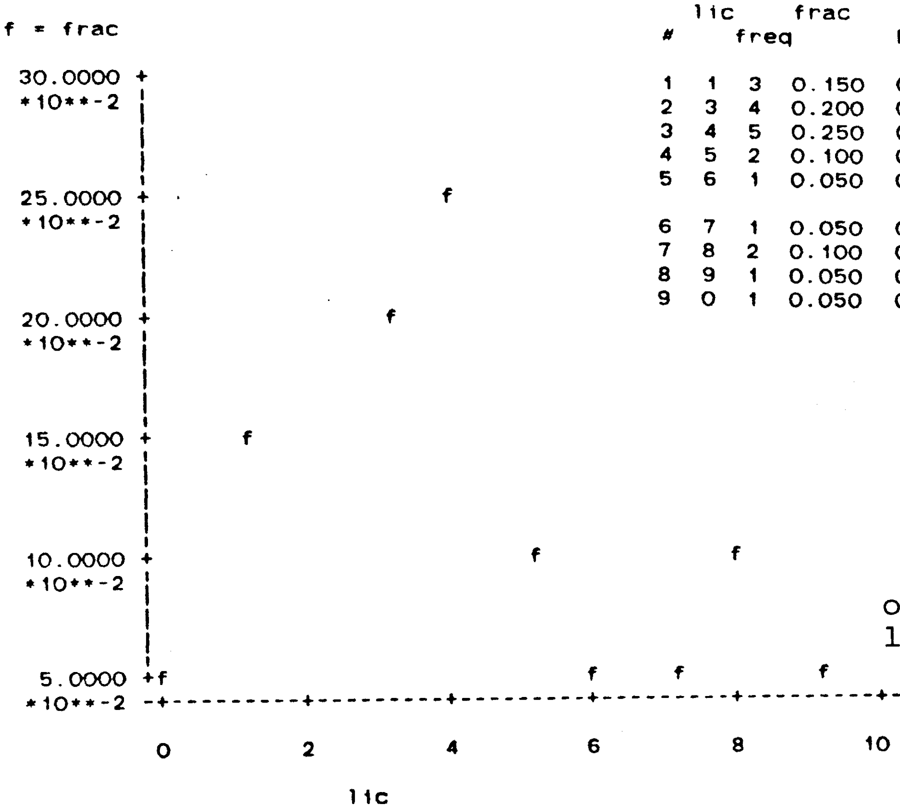

observed pmf for
licensed work orders

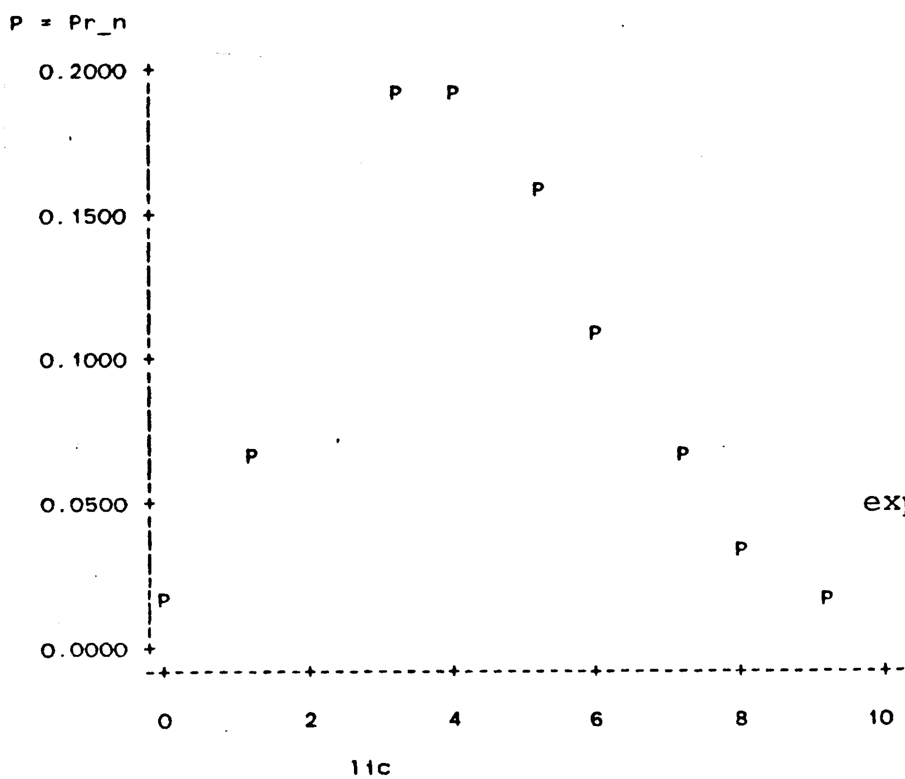

expected Poisson pmf

figure 1d

#	sk1	freq	frac	Pr_n
1	1	1	0.050	0.008
2	2	3	0.150	0.027
3	4	2	0.100	0.101
4	5	1	0.050	0.137
5	6	2	0.100	0.154
6	7	3	0.150	0.148
7	8	3	0.150	0.125
8	11	3	0.150	0.039
9	12	1	0.050	0.022
10	13	1	0.050	0.011

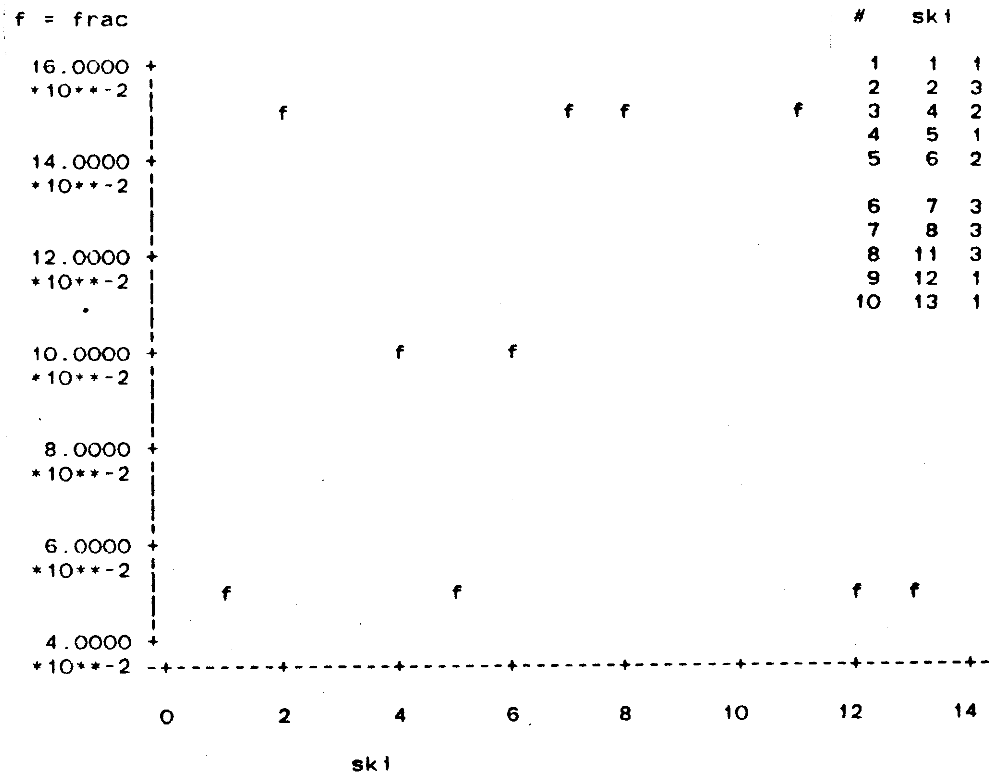

observed pmf for skilled work orders

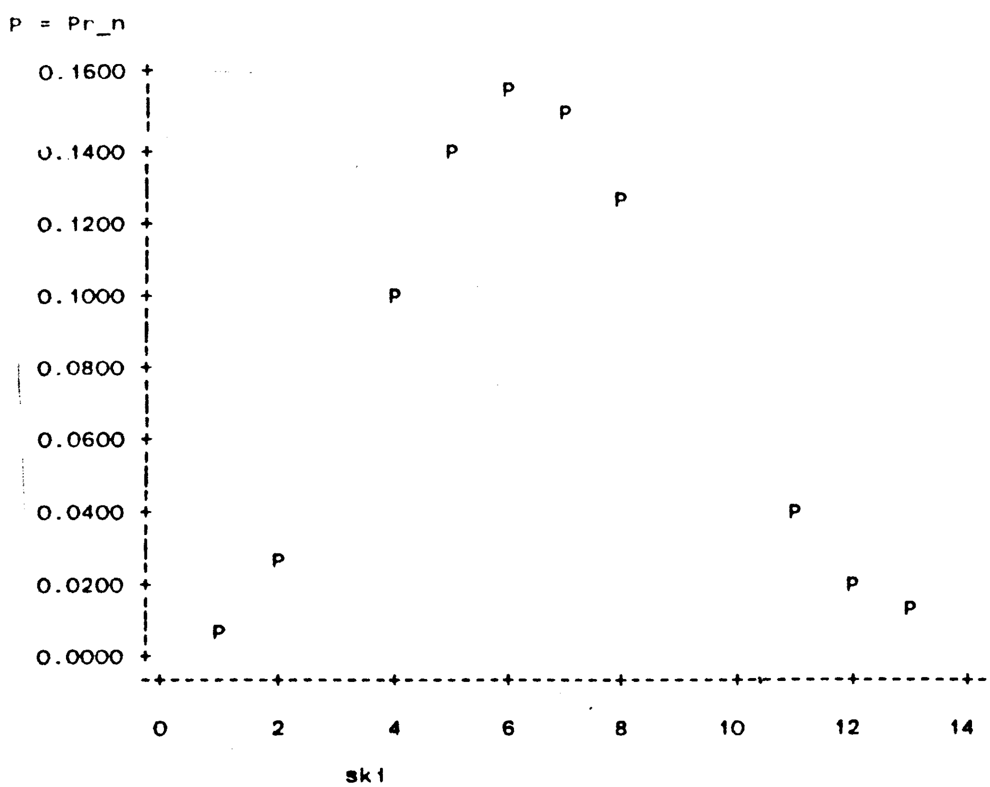

expected Poisson pmf

table 3
Regression Sensitivity Test
Comparing Random Uniform & Poisson Distributions

Regression Sensitivity Test Comparing Random Uniform & Poisson Distributions**

Range: work orders per day	Sample size days	Poisson on Uniform			Uniform on Poisson		
		r^2	B0	B1	r^2	B0	B1
1 to 12	20	.12	.06	.33	.12	.05*	.36
1 to 18	20	.07	.04	.26	.07	.04	.25
1 to 12	40	.16	.04*	.47	.16	.06	.34
1 to 18	40	.02	.05	.13	.02	.05	.14
1 to 12	60	.35	.03*	.69	.35	.04	.52
1 to 18	60	.00	.06	-.06	.00	.06	-.04
1 to 12	100	.02	.06*	.25	.02	.08	.06
1 to 18	100	.00	.05*	.12	.00	.05	.03

* - insignificant B0 coefficient, indicating the fitted line may pass through the origin.

** - mean arrival rate used for calculating Poisson distribution is that implied by the random uniform distribution.

TABLE 4 Appendix III Kolmogorov-Smirnov Goodness of Fit Test for Arrivals, by Class & Craft, Development 108

Ho: observed=Poisson
H1: observed≠Poisson

Class/Craft Category	D statistic	n (sample size)	Critical Values a=.20	Critical Values a=.05	Decision
All work orders	.213	20	.231	.294	accept
Emergency	.184	20	.231	.294	accept
Routine	.199	20	.231	.294	accept
Licensed	.086	20	.231	.294	accept
Skilled	.164	20	.231	.294	accept
Emlic	.038	20	.231	.294	accept
Emski	.100	20	.231	.294	accept
Roulic	.156	20	.231	.294	accept
Rouski	.145	20	.231	.294	accept

table 4
Kolmogorov-Smirnov Goodness of Fit Test for Arrivals
by class & craft
development 108

table 5 Kolmogorov-Smirnov Test Calculations
by class & craft
development 108

cumPr_n	cumfrac2	allKS
0.005	0.050	-0.045
0.035	0.200	-0.165
0.075	0.250	-0.175
0.137	0.350	-0.213
0.330	0.400	-0.070
0.567	0.500	0.067
0.678	0.600	0.078
0.772	0.650	0.122
0.902	0.700	0.202
0.940	0.800	0.140
0.965	0.950	0.015
0.981	1.000	-0.019

em	freq	cumfrac	cumPr	eKS
1	3	0.150	0.061	-0.089
2	3	0.300	0.189	-0.111
3	3	0.450	0.372	-0.078
4	6	0.750	0.566	-0.184
6	1	0.800	0.847	0.047
7	1	0.850	0.918	0.068
8	1	0.900	0.956	0.056
10	1	0.950	0.981	0.031
12	1	1.000	0.985	-0.015

rou	freq	cumfrac	cumPr	rKS
2	1	0.050	0.037	-0.013
3	5	0.300	0.101	-0.199
4	1	0.350	0.206	-0.144
5	2	0.450	0.346	-0.104
6	2	0.550	0.502	-0.048
7	1	0.600	0.649	0.049
8	1	0.650	0.772	0.122
9	2	0.750	0.863	0.113
10	1	0.800	0.923	0.123
11	2	0.900	0.960	0.060
12	1	0.950	0.980	0.030
13	1	1.000	0.990	-0.010

lic	freq	cumfrac	cumPr	lKS
1	3	0.150	0.131	-0.019
3	4	0.350	0.351	0.001
4	5	0.600	0.546	-0.054
5	2	0.700	0.708	0.008
6	1	0.750	0.820	0.070
7	1	0.800	0.886	0.086
8	2	0.900	0.920	0.020
9	1	0.950	0.936	-0.014
0	1	1.000	0.952	-0.048

ski	freq	cumfrac	cumPr	sKS
1	1	0.050	0.009	-0.041
2	3	0.200	0.036	-0.164
4	2	0.300	0.197	-0.103
5	1	0.350	0.334	-0.016
6	2	0.450	0.488	0.038
7	3	0.600	0.636	0.036
8	3	0.750	0.761	0.011
11	3	0.900	0.957	0.057
12	1	0.950	0.979	0.029
13	1	1.000	0.990	-0.010

emlic	freq	cumfrac	cumPr	elKS
1	7	0.350	0.329	-0.021
2	4	0.550	0.584	0.034
3	3	0.700	0.716	0.016
7	1	0.750	0.788	0.038
0	5	1.000	1.000	-0.000

emski	freq	cumfrac	cumPr	esKS
1	5	0.250	0.187	-0.063
2	4	0.450	0.435	-0.015
3	4	0.650	0.654	0.004
4	1	0.700	0.800	0.100
5	3	0.850	0.877	0.027
9	1	0.900	0.929	0.029
0	2	1.000	1.000	-0.000

roulic	freq	cumfrac	cumPr	rlKS
1	3	0.150	0.199	0.049
2	4	0.350	0.453	0.103
3	4	0.550	0.669	0.119
4	2	0.650	0.806	0.156
5	4	0.850	0.876	0.026
0	3	1.000	0.955	-0.045

rouski	freq	cumfrac	cumPr	rsKS
1	3	0.150	0.056	-0.094
2	2	0.250	0.178	-0.072
3	5	0.500	0.355	-0.145
4	2	0.600	0.548	-0.052
5	1	0.650	0.715	0.065
6	2	0.750	0.837	0.087
7	2	0.850	0.912	0.062
8	2	0.950	0.953	0.003
10	1	1.000	0.982	-0.018

figure 2

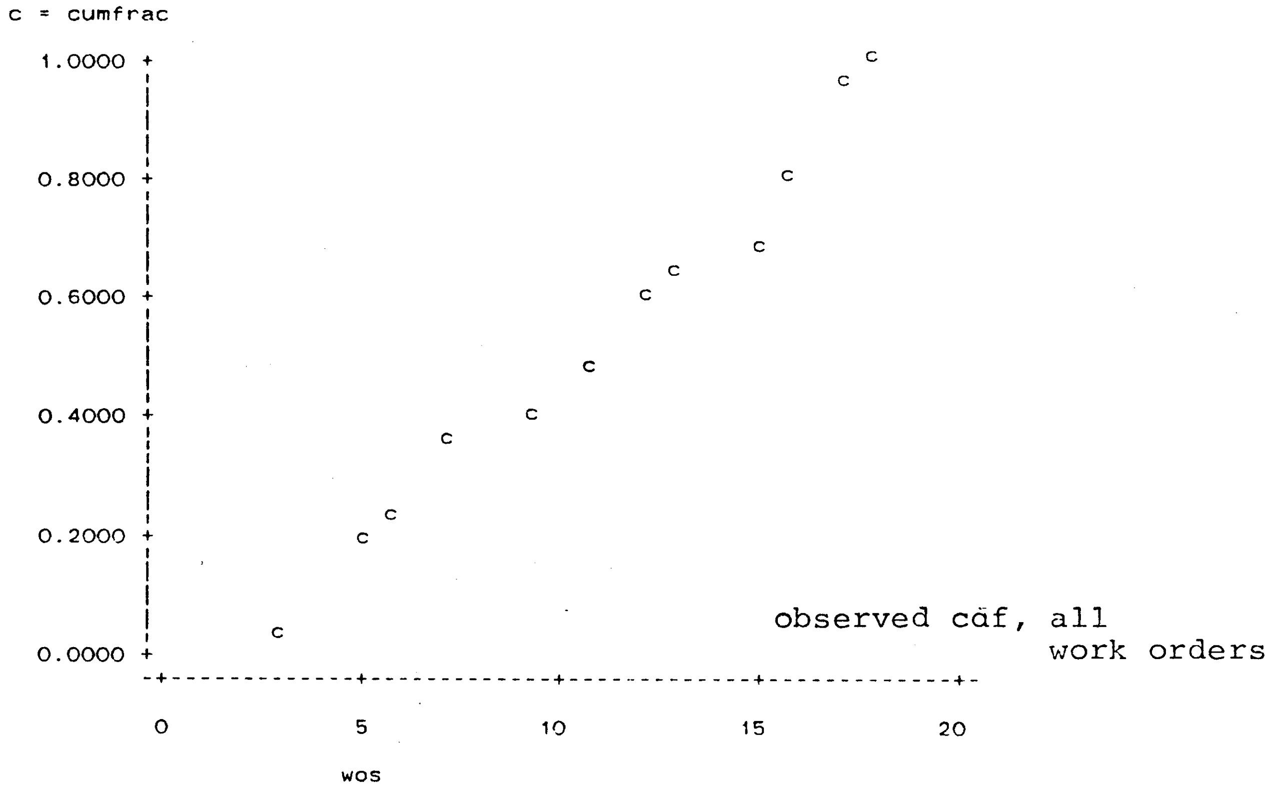
c = cumfrac
1.0000
0.8000
0.6000
0.4000
0.2000
0.0000
observed cdf, all
work orders
0
5
10
15
20
wos

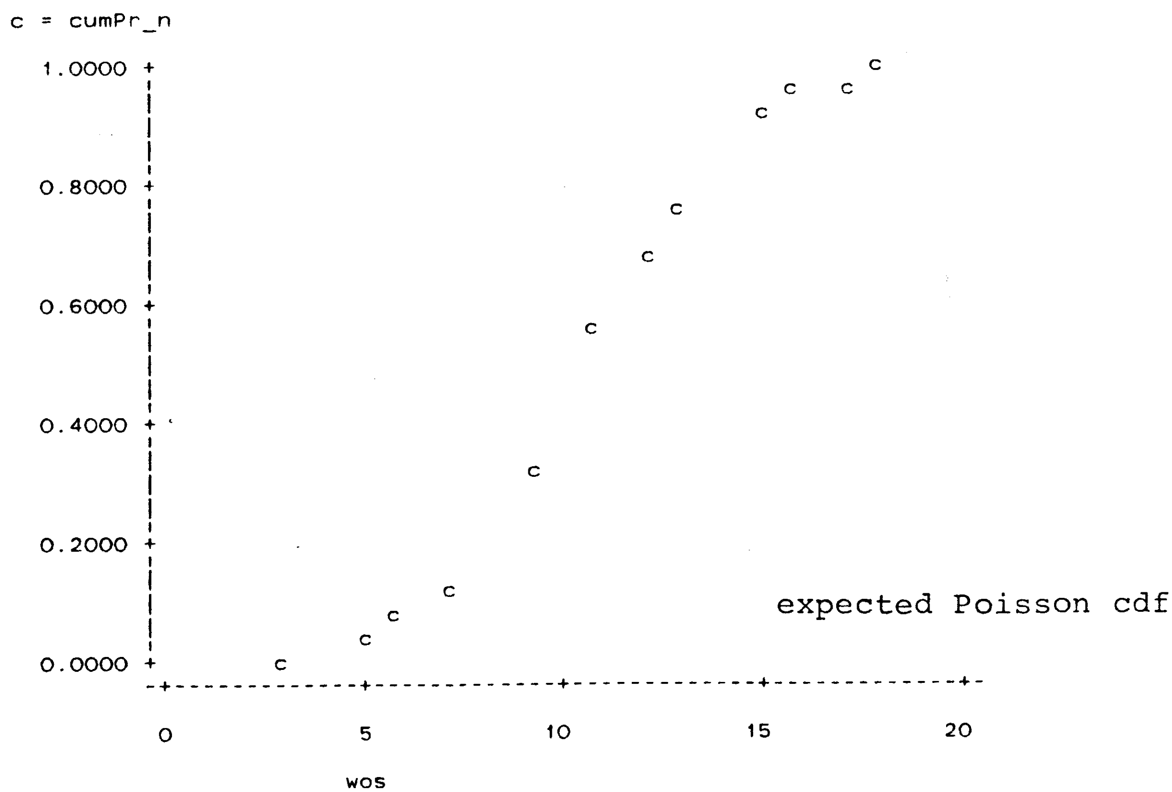
c = cumPr_n
1.0000
0.8000
0.6000
0.4000
0.2000
0.0000
expected Poisson cdf
0
5
10
15
20
wos

Table 6 Kolmogorov-Smirnov Sensitivity Test Comparing Random Uniform and Poisson Distributions *

Range: work orders per day	Sample size n	D statistic D	Critical Values a = .20	a = .05	Decision
1 to 12	20	.147	.231	.294	accept Ho
1 to 18	20	.236	.231	.294	reject at .20 accept at .05
1 to 12	40	.166	.169	.215	accept Ho
1 to 18	40	.176	.169	.215	reject at .20 accept at .05
1 to 12	60	.104	.138	.176	accept Ho
1 to 18	60	.227	.138	.176	reject Ho
1 to 12	100	.152	.107	.136	reject Ho
1 to 18	100	.199	.107	.136	reject Ho

Ho: uniform = Poisson (test insensitive to differences at this n)

H1: test sensitive at least to differences between uniform & Poisson at this n

* - mean arrival rate used for calculating Poisson distribution is that implied by the random uniform distribution.

table 7a

Kolmogorov-Smirnov Sensitivity Test Calculations

#	wos	un20	un40	un60	un100	frac20	frac40	frac60	frac100
1	1.000	1.000	3.000	1.000	7.000	0.050	0.075	0.017	0.070
2	2.000	3.000	3.000	6.000	6.000	0.150	0.075	0.100	0.060
3	3.000	0.000	8.000	5.000	12.000	0.000	0.200	0.083	0.120
4	4.000	2.000	2.000	6.000	8.000	0.100	0.050	0.100	0.080
5	5.000	4.000	4.000	10.000	7.000	0.200	0.100	0.167	0.070
6	6.000	1.000	3.000	4.000	10.000	0.050	0.075	0.067	0.100
7	7.000	3.000	5.000	8.000	9.000	0.150	0.125	0.133	0.090
8	8.000	1.000	3.000	3.000	5.000	0.050	0.075	0.050	0.050
9	9.000	1.000	5.000	7.000	9.000	0.050	0.125	0.117	0.090
10	10.000	1.000	2.000	7.000	14.000	0.050	0.050	0.117	0.140
11	11.000	2.000	2.000	1.000	8.000	0.100	0.050	0.017	0.080
12	12.000	1.000	0.000	2.000	5.000	0.050	0.000	0.033	0.050

uncdf20	uncdf40	uncdf60	uncdf100	poi_pmf20	poi_pmf40	poi_pmf6	po
0.050	0.075	0.017	0.070	0.013	0.020	0.012	
0.200	0.150	0.117	0.130	0.040	0.057	0.038	
0.200	0.350	0.200	0.250	0.083	0.107	0.079	
0.300	0.400	0.300	0.330	0.127	0.150	0.124	
0.500	0.500	0.467	0.400	0.156	0.169	0.154	
0.550	0.575	0.533	0.500	0.160	0.159	0.160	
0.700	0.700	0.667	0.590	0.141	0.128	0.142	
0.750	0.775	0.717	0.640	0.108	0.090	0.111	
0.800	0.900	0.833	0.730	0.074	0.056	0.077	
0.850	0.950	0.950	0.870	0.046	0.032	0.048	
0.950	1.000	0.967	0.950	0.025	0.016	0.027	
1.000	1.000	1.000	1.000	0.013	0.008	0.014	

i_pmf100 0	poi_cdf20	poi_cdf40	poi_cdf60	poi_cdf100	D20	D40	D60	D100
0.010	0.013	0.020	0.012	0.010	-0.037	-0.055	-0.004	-0.060
0.032	0.053	0.077	0.050	0.042	-0.147	-0.073	-0.066	-0.088
0.070	0.136	0.184	0.130	0.112	-0.064	-0.166	-0.070	-0.138
0.113	0.263	0.334	0.253	0.224	-0.037	-0.066	-0.047	-0.106
0.146	0.420	0.503	0.408	0.370	-0.080	0.003	-0.059	-0.030
0.158	0.580	0.662	0.567	0.528	0.030	0.087	0.034	0.028
0.146	0.721	0.790	0.710	0.674	0.021	0.090	0.043	0.084
0.118	0.829	0.879	0.821	0.792	0.079	0.104	0.104	0.152
0.085	0.903	0.936	0.897	0.878	0.103	0.036	0.064	0.148
0.055	0.949	0.967	0.945	0.933	0.099	0.017	-0.005	0.063
0.033	0.974	0.983	0.972	0.965	0.024	-0.017	0.006	0.015
0.018	0.987	0.991	0.986	0.983	-0.013	-0.009	-0.014	-0.017

table 7b

Kolmogorov-Smirnov Sensitivity Test Calculations

#	wos2	un20	un40	un60	un100	frac20	frac40	frac60	frac100
1	1.000	1.000	0.000	3.000	1.000	0.050	0.000	0.050	0.010
2	2.000	1.000	3.000	1.000	10.000	0.050	0.075	0.017	0.100
3	3.000	1.000	3.000	5.000	10.000	0.050	0.075	0.083	0.100
4	4.000	3.000	1.000	6.000	4.000	0.150	0.025	0.100	0.040
5	5.000	0.000	2.000	3.000	7.000	0.000	0.050	0.050	0.070
6	6.000	0.000	3.000	0.000	6.000	0.000	0.075	0.000	0.060
7	7.000	1.000	2.000	1.000	6.000	0.050	0.050	0.017	0.060
8	8.000	1.000	0.000	6.000	5.000	0.050	0.000	0.100	0.050
9	9.000	2.000	0.000	2.000	7.000	0.100	0.000	0.033	0.070
10	10.000	3.000	2.000	2.000	2.000	0.150	0.050	0.033	0.020
11	11.000	2.000	6.000	7.000	8.000	0.100	0.150	0.117	0.080
12	12.000	0.000	6.000	2.000	5.000	0.000	0.150	0.033	0.050
13	13.000	1.000	1.000	2.000	7.000	0.050	0.025	0.033	0.070
14	14.000	2.000	3.000	4.000	7.000	0.100	0.075	0.067	0.070
15	15.000	1.000	1.000	6.000	2.000	0.050	0.025	0.100	0.020
16	16.000	1.000	5.000	4.000	4.000	0.050	0.125	0.067	0.040
17	17.000	0.000	1.000	6.000	4.000	0.000	0.025	0.100	0.040
18	18.000	0.000	1.000	0.000	5.000	0.000	0.025	0.000	0.050

uncdf20	uncdf40	uncdf60	uncdf100	pmf20	pmf40	pmf60
0.050	0.000	0.050	0.010	0.001	0.000	0.001
0.100	0.075	0.067	0.110	0.006	0.002	0.003
0.150	0.150	0.150	0.210	0.018	0.007	0.008
0.300	0.175	0.250	0.250	0.039	0.018	0.021
0.300	0.225	0.300	0.320	0.068	0.036	0.041
0.300	0.300	0.300	0.380	0.099	0.061	0.067
0.350	0.350	0.317	0.440	0.123	0.087	0.094
0.400	0.350	0.417	0.490	0.135	0.110	0.116
0.500	0.350	0.450	0.560	0.131	0.124	0.127
0.650	0.400	0.483	0.580	0.115	0.125	0.125
0.750	0.550	0.600	0.660	0.091	0.115	0.112
0.750	0.700	0.633	0.710	0.067	0.097	0.092
0.800	0.725	0.667	0.780	0.045	0.075	0.070
0.900	0.800	0.733	0.850	0.028	0.054	0.049
0.950	0.825	0.833	0.870	0.016	0.036	0.032
1.000	0.950	0.900	0.910	0.009	0.023	0.020
1.000	0.975	1.000	0.950	0.005	0.014	0.011
1.000	1.000	1.000	1.000	0.002	0.008	0.006

pmf100	poi_cdf20	poi_cdf40	poi_cdf60	poi_cdf100	D20	D40	D60	D100
0.001	0.001	0.000	0.001	0.001	-0.049	0.000	-0.049	-0.009
0.005	0.007	0.003	0.003	0.007	-0.093	-0.072	-0.064	-0.103
0.016	0.025	0.010	0.011	0.022	-0.125	-0.140	-0.139	-0.188
0.035	0.064	0.027	0.032	0.058	-0.236	-0.148	-0.218	-0.192
0.063	0.132	0.063	0.073	0.121	-0.168	-0.162	-0.227	-0.199
0.094	0.230	0.124	0.140	0.214	-0.070	-0.176	-0.160	-0.166
0.119	0.354	0.211	0.234	0.333	0.004	-0.139	-0.083	-0.107
0.133	0.489	0.322	0.350	0.466	0.089	-0.028	-0.067	-0.024
0.132	0.620	0.445	0.477	0.598	0.120	0.095	0.027	0.038
0.117	0.735	0.570	0.602	0.715	0.085	0.170	0.118	0.135
0.095	0.826	0.685	0.714	0.811	0.076	0.135	0.114	0.151
0.071	0.893	0.782	0.806	0.881	0.143	0.082	0.172	0.171
0.049	0.938	0.857	0.875	0.930	0.138	0.132	0.208	0.150
0.031	0.966	0.911	0.924	0.961	0.066	0.111	0.191	0.111
0.018	0.982	0.948	0.956	0.979	0.032	0.123	0.123	0.109
0.010	0.991	0.971	0.976	0.990	-0.009	0.021	0.076	0.080
0.005	0.996	0.984	0.987	0.995	-0.004	0.009	-0.013	0.045
0.003	0.998	0.992	0.994	0.998	-0.002	-0.008	-0.006	-0.002

table 8
Regressions of Mean Arrivals on Variances by class & craft

Regressions of Mean Arrivals on Variances, by Class & Craft

Class/Craft Category	Variance on Mean			Mean on Variance			Var. on Mean		Mean on Var.	
	r^2	B0	B1	r^2	B0	B1	r^2	B1	r^2	B1
All work orders	.86	-9.65	3.50	.86	3.75	.25	.80	2.63	.61	.36
Emergency	.93	-1.70*	2.74	.93	.81	.34	.91	2.43	.89	.39
Routine	.93	-4.08	3.31	.93	1.80	.28	.90	2.84	.86	.33
Licensed	.74	0	1.46	.74	1.07*	.50	.74	1.46	.66	.65
Skilled	.84	-21.7	6.44	.84	3.80	.13	.58	3.08	error	.26
Emlic	.96	-.25*	1.68	.96	.21*	.57	.95	1.59	.94	.62
Emski	.93	-.50*	2.31	.93	.34	.40	.92	2.14	.90	.45
Roulic	.89	-.06*	1.48	.89	.31*	.60	.89	1.46	.88	.65
Rouski	.74	-8.20	4.94	.74	2.23	.15	.65	3.54	.29	.21

* -insignificant B0 coefficient, indicating fitted line may pass through the origin.

APPENDIX IV

1) <u>Use of Regression to Test Exponential Service Times</u>

Using linear regression, the exponential distribution was compared with observed probabilities representing the fraction of actual service times that have a particular length. Regressions were run on the noncumulative probabilities just as they were for arrivals, using the same series of simple regression equations. Table 2 presents the results of these tests, and table 3 gives the actual distributions. Slopes given by the "observed on expected" regressions are in all cases fairly close to one, but an analysis of variation in the other direction (expected on observed) suggests that service times are non-Poisson. With the conservative decision rule that Poisson processes should have r-squares over 50 and slopes between 60 and 140 in each of the four regressions, we conclude that service times are not Poisson. In no cases, however, do the observed service time distributions differ wildly from the Poisson model. These distributions are also more homogenous across classes and crafts than are the arrival rates.

Once more, however, the problem of inhomogenous variance makes the results of these regression tests unreliable. A weighted least squares approach should be used to account for non-constant variances.

Following is a collection of data, calculations and scatter plots used in the service time analyses.

table 1

Work Orders
by
service time, class & craft

#	hrs	wos	em	rou	lic	ski	emlic	emski	roulic	rouski
1	0.500	352	138	210	127	222	68	67	59	158
2	1.000	719	275	441	295	421	100	174	195	250
3	1.500	463	134	327	191	271	30	104	161	168
4	2.000	477	174	301	202	275	71	103	129	174
5	2.500	168	57	110	76	92	27	30	48	63
6	3.000	206	73	133	88	117	42	30	46	88
7	3.500	41	14	27	24	17	9	5	15	12
8	4.000	152	60	91	59	92	24	35	35	58
9	4.500	26	11	15	17	9	6	5	11	4
10	5.000	41	14	27	18	22	10	4	8	19
11	5.500	9	2	7	4	5	1	1	3	4
12	6.000	51	26	24	23	28	15	11	7	18
13	6.500	9	3	6	4	5	1	2	3	3
14	7.000	30	6	24	8	21	3	3	5	19
15	7.500	7	1	6	2	5	1	0	1	5
16	8.000	123	39	81	53	70	24	15	28	56
17	8.500	5	0	5	0	5	0	0	0	5
18	9.000	4	2	2	0	4	0	2	0	2
19	9.500	2	1	1	2	0	1	0	1	0
20	10.000	8	3	5	2	6	0	3	2	3
21	10.500	2	1	1	0	2	0	1	0	1
22	11.000	6	2	4	2	4	1	1	1	3
23	11.500	1	1	0	1	0	1	0	0	0
24	12.000	18	6	12	4	14	2	4	2	10
25	13.000	5	0	5	3	2	0	0	3	2
26	14.000	17	6	11	1	16	0	6	1	10
27	15.000	7	2	5	1	6	1	1	0	5
28	16.000	28	10	18	7	21	5	5	2	16
29	20.000	1	0	1	0	1	0	0	0	1
30	21.000	4	0	4	0	4	0	0	0	4
31	22.000	2	1	1	1	1	1	0	0	1
32	23.000	2	0	2	1	1	0	0	1	1
33	24.000	8	1	7	1	7	0	1	1	6
34	25.000	2	1	1	0	2	0	1	0	1
35	26.000	2	1	1	1	1	1	0	0	1
36	27.000	2	0	2	0	2	0	0	0	2
37	28.000	5	2	3	0	5	0	2	0	3
38	29.000	2	1	1	1	1	1	0	0	1
39	30.000	4	0	4	0	4	0	0	0	4
40	32.000	3	1	2	0	3	0	1	0	2
41	34.000	1	0	1	0	1	0	0	0	1
42	38.000	2	0	2	0	2	0	0	0	2
43	39.000	2	1	1	1	1	1	0	0	1
44	40.000	2	0	2	0	2	0	0	0	2
45	43.000	1	0	1	0	1	0	0	0	1
46	45.000	1	0	1	0	1	0	0	0	1
47	48.000	3	2	1	0	3	0	2	0	1
48	52.000	3	0	3	0	3	0	0	0	3
49	54.000	1	0	1	0	1	0	0	0	1
50	64.000	2	0	2	0	2	0	0	0	2

table 2

Linear Regressions of Observed on Expected Exponential Service Times by class & craft development 108

Linear Regressions of Observed Service Times on Expected

Class/Craft Category	Expected on Observed			Observed on Expected			Exp. on Obs.		Obs. on Exp.	
	r^2	B0	B1	r^2	B0	B1	r^2	B1	r^2	B1
All work orders	.63	.02	.51	.63	0	1.24	.49	.64	.62	1.17
Emergency	.57	.02	.48	.57	0	1.19	.43	.62	.57	1.15
Routine	.65	.02	.52	.65	0	1.26	.54	.65	.65	1.19
Licensed	.44	.02	.41	.44	0	1.07	.25	.56	.44	1.09
Skilled	.59	.02	.43	.59	-.01*	1.36	.39	.57	.58	1.23
Emlic	.32	.03	.31	.32	0	1.03	error	.48	.32	1.07
Emski	.64	.02	.49	.64	0	1.30	.51	.61	.63	1.21
Roulic	.45	.02	.41	.45	0	1.08	.28	.55	.45	1.10
Rouski	.56	.02	.40	.56	-.01*	1.39	.32	.55	.55	1.24

* -insignificant B0 coefficient, indicating fitted line may pass through the origin.

table 3

Observed & Expected Service Time Probability Density Functions
by class & craft
development 108

hrs	emfrac	epdf	rfrac	rpdf	lfrac	lpdf	sfrac	spdf
0.500	0.080	0.140	0.085	0.140	0.061	0.148	0.098	0.138
1.000	0.250	0.119	0.223	0.118	0.140	0.124	0.299	0.117
1.500	0.110	0.101	0.149	0.100	0.105	0.104	0.161	0.100
2.000	0.130	0.085	0.176	0.085	0.263	0.087	0.098	0.085
2.500	0.030	0.072	0.032	0.072	0.044	0.073	0.023	0.072
3.000	0.140	0.061	0.096	0.061	0.175	0.061	0.069	0.061
3.500	0.000	0.052	0.005	0.052	0.009	0.051	0.000	0.052
4.000	0.070	0.044	0.048	0.044	0.061	0.043	0.046	0.044
4.500	0.010	0.037	0.016	0.037	0.018	0.036	0.011	0.038
5.000	0.020	0.032	0.037	0.032	0.026	0.030	0.029	0.032
5.500	0.010	0.027	0.000	0.027	0.000	0.025	0.006	0.027
6.000	0.040	0.023	0.032	0.023	0.053	0.021	0.023	0.023
6.500	0.000	0.019	0.011	0.019	0.009	0.018	0.006	0.020
7.000	0.000	0.016	0.011	0.016	0.000	0.015	0.006	0.017
8.000	0.070	0.023	0.048	0.024	0.018	0.021	0.080	0.024
9.000	0.010	0.017	0.000	0.017	0.000	0.015	0.006	0.018
10.000	0.000	0.012	0.005	0.012	0.000	0.010	0.006	0.013
11.000	0.000	0.009	0.005	0.009	0.000	0.007	0.006	0.009
12.000	0.010	0.006	0.000	0.006	0.000	0.005	0.006	0.007
16.000	0.020	0.002	0.000	0.002	0.009	0.001	0.006	0.002
22.000	0.000	0.000	0.005	0.000	0.000	0.000	0.006	0.000
24.000	0.000	0.000	0.011	0.000	0.009	0.000	0.006	0.000
30.000	0.000	0.000	0.005	0.000	0.000	0.000	0.006	0.000

hrs	elfrac	elpdf	esfrac	espdf	rlfrac	rlpdf	rsfrac	rspdf
0.500	0.027	0.132	0.113	0.147	0.079	0.156	0.092	0.131
1.000	0.189	0.113	0.290	0.123	0.118	0.129	0.294	0.113
1.500	0.081	0.097	0.129	0.103	0.118	0.107	0.174	0.097
2.000	0.162	0.083	0.113	0.087	0.303	0.089	0.092	0.083
2.500	0.027	0.071	0.032	0.073	0.053	0.073	0.018	0.071
3.000	0.243	0.061	0.081	0.061	0.145	0.061	0.064	0.061
3.500	0.000	0.052	0.000	0.051	0.013	0.050	0.000	0.052
4.000	0.081	0.045	0.048	0.043	0.053	0.042	0.046	0.045
4.500	0.000	0.039	0.016	0.036	0.026	0.035	0.009	0.039
5.000	0.027	0.033	0.016	0.030	0.026	0.029	0.037	0.033
5.500	0.000	0.028	0.016	0.026	0.000	0.024	0.000	0.028
6.000	0.081	0.024	0.016	0.021	0.039	0.020	0.028	0.024
6.500	0.000	0.021	0.000	0.018	0.013	0.016	0.009	0.021
7.000	0.000	0.018	0.000	0.015	0.000	0.013	0.009	0.018
8.000	0.054	0.026	0.081	0.021	0.000	0.018	0.083	0.026
9.000	0.000	0.019	0.016	0.015	0.000	0.013	0.000	0.019
10.000	0.000	0.014	0.000	0.011	0.000	0.009	0.009	0.014
11.000	0.000	0.010	0.000	0.008	0.000	0.006	0.009	0.011
12.000	0.000	0.008	0.016	0.005	0.000	0.004	0.000	0.008
16.000	0.027	0.002	0.016	0.001	0.000	0.001	0.000	0.002
22.000	0.000	0.000	0.000	0.000	0.000	0.000	0.009	0.000
24.000	0.000	0.000	0.000	0.000	0.013	0.000	0.009	0.000
30.000	0.000	0.000	0.000	0.000	0.000	0.000	0.009	0.000

table 4

Kolmogorov-Smirnov Test Calculations

by class & craft

development 108

hrs	cepdf	cefrac	eks	crpdf	crfrac	rks	clpdf	clfrac	lks	cspdf	csfrac	sks
0.500	0.140	0.080	0.060	0.140	0.085	0.054	0.148	0.061	0.086	0.138	0.098	0.040
1.000	0.259	0.330	-0.071	0.258	0.309	-0.051	0.271	0.202	0.070	0.255	0.397	-0.142
1.500	0.360	0.440	-0.080	0.358	0.457	-0.099	0.375	0.307	0.068	0.355	0.557	-0.203
2.000	0.445	0.570	-0.125	0.444	0.633	-0.189	0.462	0.570	-0.108	0.439	0.655	-0.216
2.500	0.518	0.600	-0.082	0.516	0.665	-0.149	0.535	0.614	-0.079	0.511	0.678	-0.167
3.000	0.579	0.740	-0.161	0.577	0.761	-0.184	0.596	0.789	-0.193	0.573	0.747	-0.174
3.500	0.631	0.740	-0.109	0.629	0.766	-0.137	0.648	0.798	-0.150	0.625	0.747	-0.122
4.000	0.675	0.810	-0.135	0.673	0.814	-0.141	0.691	0.860	-0.169	0.669	0.793	-0.124
4.500	0.712	0.820	-0.108	0.711	0.830	-0.119	0.727	0.877	-0.150	0.707	0.805	-0.098
5.000	0.744	0.840	-0.096	0.742	0.867	-0.125	0.757	0.904	-0.146	0.739	0.833	-0.094
5.500	0.771	0.850	-0.079	0.769	0.867	-0.098	0.783	0.904	-0.121	0.766	0.839	-0.073
6.000	0.793	0.890	-0.097	0.792	0.899	-0.107	0.804	0.956	-0.152	0.789	0.862	-0.073
6.500	0.813	0.890	-0.077	0.812	0.910	-0.098	0.822	0.965	-0.143	0.809	0.868	-0.059
7.000	0.829	0.890	-0.061	0.828	0.920	-0.092	0.837	0.965	-0.128	0.826	0.874	-0.048
8.000	0.852	0.960	-0.108	0.852	0.968	-0.116	0.858	0.982	-0.125	0.850	0.954	-0.104
9.000	0.869	0.970	-0.101	0.869	0.968	-0.099	0.873	0.982	-0.110	0.868	0.960	-0.092
10.000	0.881	0.970	-0.089	0.881	0.973	-0.092	0.883	0.982	-0.099	0.880	0.966	-0.085
11.000	0.890	0.970	-0.080	0.890	0.979	-0.089	0.890	0.982	-0.092	0.889	0.971	-0.082
12.000	0.896	0.980	-0.084	0.896	0.979	-0.083	0.896	0.982	-0.087	0.896	0.977	-0.081
16.000	0.898	1.000	-0.102	0.898	0.979	-0.081	0.897	0.991	-0.094	0.898	0.983	-0.085
22.000	0.898	1.000	-0.102	0.898	0.984	-0.086	0.897	0.991	-0.094	0.898	0.989	-0.090
24.000	0.898	1.000	-0.102	0.898	0.995	-0.096	0.897	1.000	-0.103	0.898	0.994	-0.096
30.000	0.898	1.000	-0.102	0.898	1.000	-0.102	0.897	1.000	-0.103	0.898	1.000	-0.102

hrs	celpdf	celfrac	elks	cespdf	cesfrac	esks	crlpdf	crlfrac	rlks	crspdf	crsfrac	rsks
0.500	0.132	0.027	0.105	0.147	0.113	0.034	0.156	0.079	0.077	0.131	0.092	0.040
1.000	0.245	0.216	0.028	0.270	0.403	-0.134	0.285	0.197	0.088	0.244	0.385	-0.141
1.500	0.341	0.297	0.044	0.373	0.532	-0.159	0.392	0.316	0.077	0.341	0.560	-0.219
2.000	0.425	0.459	-0.035	0.460	0.645	-0.185	0.481	0.618	-0.137	0.424	0.651	-0.228
2.500	0.496	0.486	0.009	0.533	0.677	-0.145	0.555	0.671	-0.116	0.495	0.670	-0.175
3.000	0.557	0.730	-0.173	0.594	0.758	-0.164	0.615	0.816	-0.200	0.556	0.734	-0.178
3.500	0.609	0.730	-0.120	0.645	0.758	-0.113	0.666	0.829	-0.163	0.608	0.734	-0.126
4.000	0.654	0.811	-0.157	0.689	0.806	-0.118	0.708	0.882	-0.174	0.653	0.780	-0.126
4.500	0.693	0.811	-0.118	0.725	0.823	-0.098	0.742	0.908	-0.166	0.692	0.789	-0.097
5.000	0.726	0.838	-0.112	0.755	0.839	-0.083	0.771	0.934	-0.163	0.725	0.826	-0.101
5.500	0.754	0.838	-0.084	0.781	0.855	-0.074	0.794	0.934	-0.140	0.753	0.826	-0.072
6.000	0.779	0.919	-0.140	0.803	0.871	-0.068	0.814	0.974	-0.160	0.778	0.853	-0.075
6.500	0.799	0.919	-0.119	0.821	0.871	-0.050	0.830	0.987	-0.157	0.799	0.862	-0.064
7.000	0.817	0.919	-0.102	0.836	0.871	-0.035	0.844	0.987	-0.143	0.817	0.872	-0.055
8.000	0.844	0.973	-0.129	0.857	0.952	-0.094	0.862	0.987	-0.125	0.843	0.954	-0.111
9.000	0.863	0.973	-0.110	0.872	0.968	-0.096	0.875	0.987	-0.112	0.863	0.954	-0.091
10.000	0.877	0.973	-0.096	0.883	0.968	-0.085	0.884	0.987	-0.103	0.877	0.963	-0.086
11.000	0.888	0.973	-0.085	0.890	0.968	-0.077	0.890	0.987	-0.097	0.888	0.972	-0.085
12.000	0.895	0.973	-0.077	0.896	0.984	-0.088	0.894	0.987	-0.093	0.895	0.972	-0.077
16.000	0.898	1.000	-0.102	0.897	1.000	-0.103	0.895	0.987	-0.092	0.898	0.972	-0.075
22.000	0.898	1.000	-0.102	0.897	1.000	-0.103	0.895	0.987	-0.092	0.898	0.982	-0.084
24.000	0.898	1.000	-0.102	0.897	1.000	-0.103	0.895	1.000	-0.105	0.898	0.991	-0.093
30.000	0.898	1.000	-0.102	0.897	1.000	-0.103	0.895	1.000	-0.105	0.898	1.000	-0.102

Regressions of Mean Service Times on Standard Deviations, by Class & Craft

Class/Craft Category	St. D. on Mean			Mean on St. D.			St. D. on Mean		Mean on St. D.	
	r^2	B0	B1	r^2	B0	B1	r^2	B1	r^2	B1
All work orders	.63	-1.12*	1.90	.63	1.55	.33	.61	1.60	.12	.61
Emergency	.85	-1.21*	1.60	.85	1.12	.52	.82	1.34	.69	.71
Routine	.79	-.22*	1.59	.79	.93*	.50	.79	1.54	.72	.62
Licensed	.38	.13*	1.02	.38	1.61	.38	.38	1.06	error	.87
Skilled	.73	-1.25	1.98	.73	1.42	.37	.71	1.65	.46	.58
Emlic	.92	-.38*	1.10	.92	.62*	.84	.92	1.04	.91	.92
Emski	.67	.25*	1.12	.67	.82*	.60	.67	1.18	.60	.77
Roulic	.79	.30*	.76	.79	.22*	1.04	.78	.87	.79	1.10
Rouski	.89	.12*	1.49	.89	.44*	.59	.89	1.50	.88	.65

* -insignificant B0 coefficient, indicating fitted line may pass through the origin.

table 5

Regressions of Mean Service Times on Standard Deviations by class & craft

APPENDIX V

table 1a

#	tat	wos	emer	rout	licen	skill	lntat	lnwos	lnemer	lnrout	lnlicen	lnskill
1	1	1135	394	735	511	621	0.000	7.034	5.976	6.600	6.236	6.431
2	2	391	148	242	176	213	0.693	5.969	4.997	5.489	5.170	5.361
3	3	198	67	128	77	120	1.099	5.288	4.205	4.852	4.344	4.787
4	4	145	46	96	59	85	1.386	4.977	3.829	4.564	4.078	4.443
5	5	108	29	79	41	67	1.609	4.682	3.367	4.369	3.714	4.205
6	6	106	55	51	33	73	1.792	4.663	4.007	3.932	3.497	4.290
7	7	78	16	61	28	50	1.946	4.357	2.773	4.111	3.332	3.912
8	8	65	22	43	26	39	2.079	4.174	3.091	3.761	3.258	3.664
9	9	42	11	31	17	25	2.197	3.738	2.398	3.434	2.833	3.219
10	10	49	12	36	22	27	2.303	3.892	2.485	3.584	3.091	3.296

figure 1

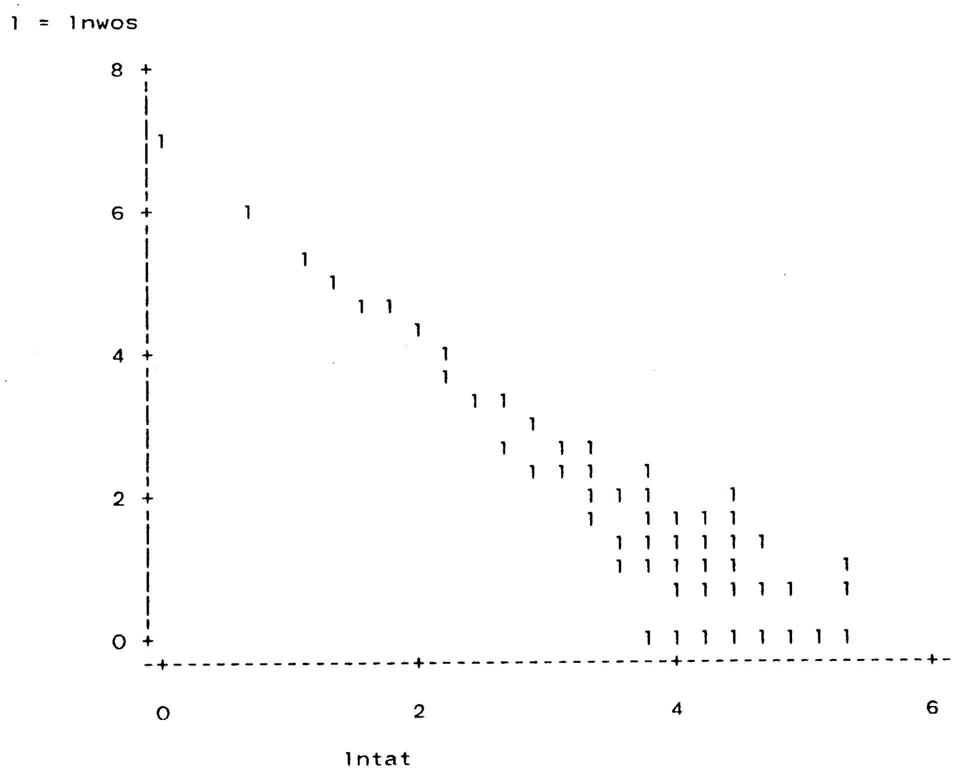

table 1b

	coef	stand_err	dgfd	F_ratio	sig
constant	6.500	0.200	129.000	1055.291	0.000
lntat	-1.298	0.049	129.000	712.540	0.000

	sum_of_squares	dgf	mean_square	F_ratio	sig
total	269.886	130.000	2.076	*	*
regression	228.515	1.000	228.515	712.540	0.000
error	41.371	129.000	0.321	*	*

f = fitted

l = lntat

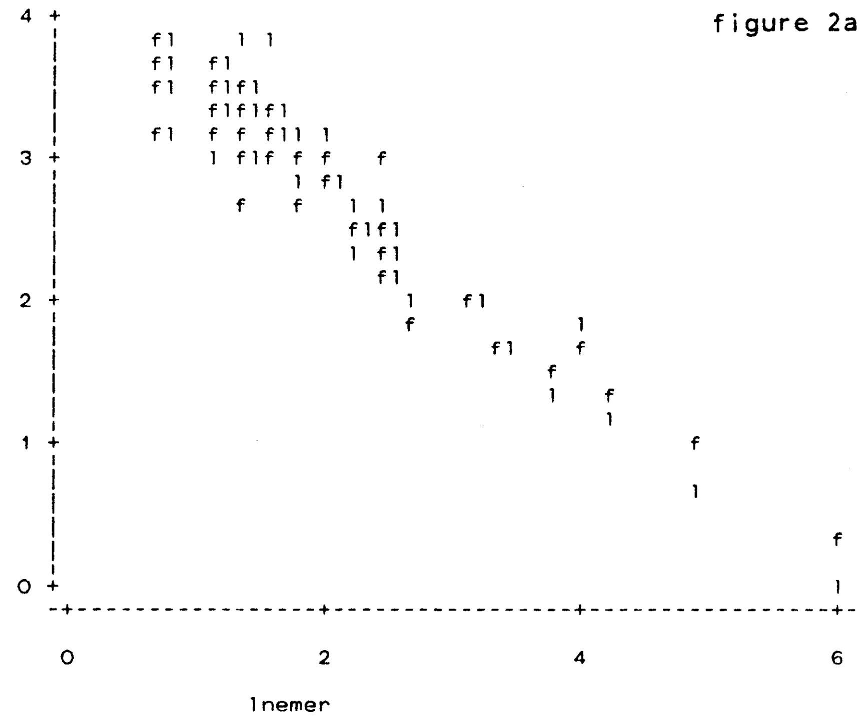

figure 2a

f = fitted

l = lntat

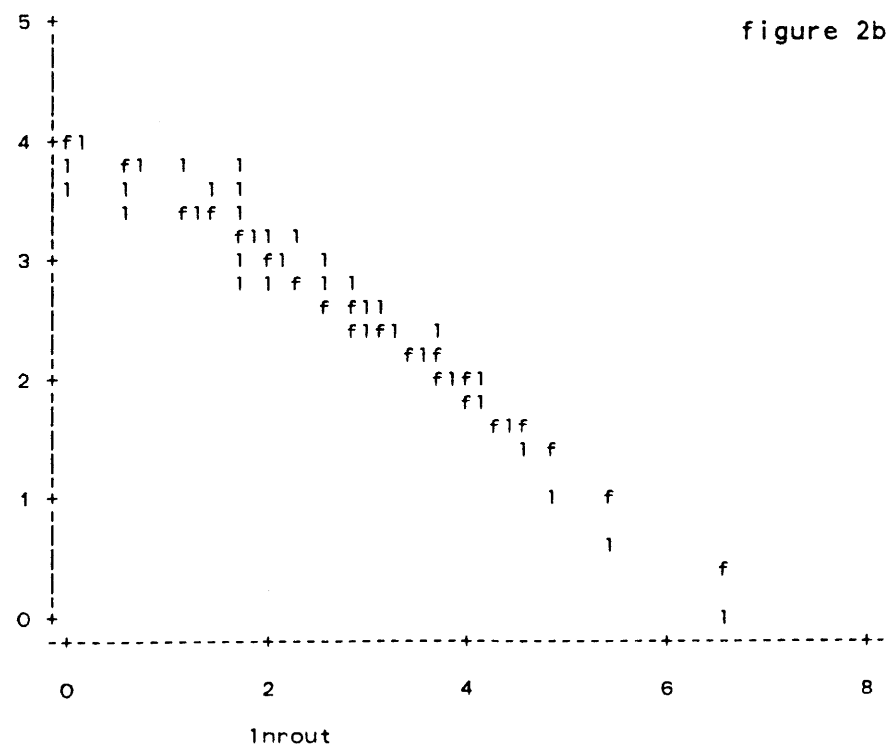

figure 2b

f = fitted

l = lntat

figure 2c

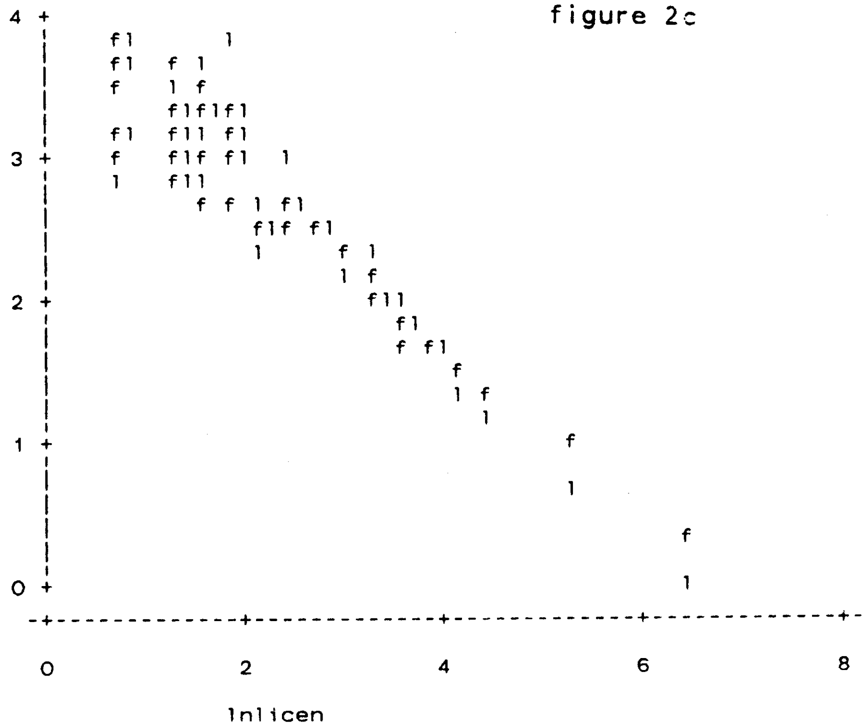

f = fitted

l = lntat

figure 2d

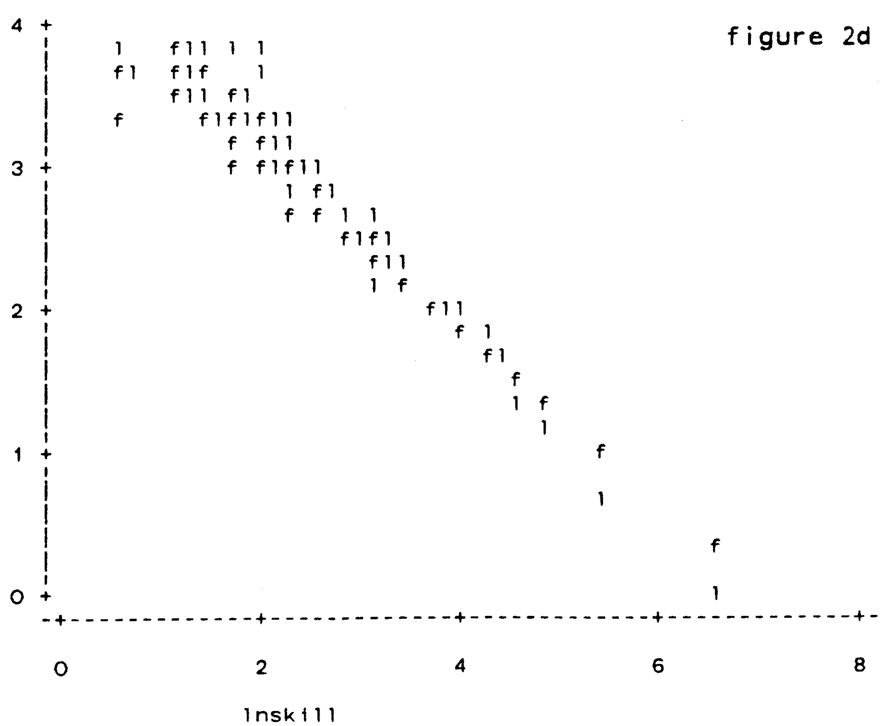

BIBLIOGRAPHY

Allen, A.O. (1978) Probability, Statistics and Queueing Theory
New York: Academic Press

Boston Housing Authority, "Monthly Management Report" Vol 1, 12/83, 1/84
________________ "Living Unit Inspections", Gilmore, D. I., 8/83
________________ "Maintenance Service", Gilmore, D. I., 3/84
________________ "Dallas Housing Authority Work Order Procedures" Gilmore, D. I., 9/83
________________ "The New Work Order System", Friend, G., 3/84
________________ "Memorandum of Agreement with the Boston Building Trades Council", 1/84

Ferreira, J. (1973) Analysis of Public Systems, Ch. 15, MIT Press

Forthofer, R. N. & Lehnen, R. G. (1981) Public Program Analysis: A New Categorical Data Approach, Belmont, CA: Wadsworth

Hoos, I. R. (1972) Systems Analysis in Public Policy: A Critique
Berkeley: UC Press

Kaplan, E. H. & Larson, R. C. (1981) "Decision-Oriented Approaches to Program Evaluation", Operations Research Center, MIT

Larson, C. L. & Odoni, A. R. (1981) Urban Operations Research
Englewood Cliffs, NJ: Prentice Hall

Larson, R. C. (1978) "Improving the Effectiveness of New York City's 911"
Operations Research Center, MIT

MIT, Laboratory of Architecture & Planning, Consistent System:
"Elementary Statistical Analysis" 1980
"CS Program Summaries" 1981
"Elementary Janus Manual" 1978
"Janus Reference Manual" 1982

Massey, F. J. (1951) "The Kolmogorov-Smirnov Test for Goodness of Fit"
J. of Amer. Stat. Assoc. 46

Newell, G. F. (1971) Applications of Queueing Theory, London: Chapman & Hall

Saaty, T. (1961) Elements of Queueing Theory, New York: McGraw Hill

Stokey, E. & Zeckhauser, R. (1978) A Primer for Policy Analysis
New York: Norton

Tufte, E. R. (1974) Data Analysis for Politics and Policy,
Englewood Cliffs, NJ: Prentice Hall

Weisberg, S. (1980) Applied Linear Regression, New York: Wiley