Proceedings of the 2024 Winter Simulation Conference
H. Lam, E. Azar, D. Batur, S. Gao, W. Xie, S. R. Hunter, and M. D. Rossetti, eds.

DEPENDENCE BETWEEN ARRIVAL AND SERVICE PROCESSES IN HEALTHCARE SIMULATION MODELLING

Laura Boyle[1,2] and Nigel Bean[2]

[1]Mathematical Sciences Research Centre, Queen's University Belfast, Belfast, Northern Ireland, UK
[2]School of Mathematical and Computer Sciences, University of Adelaide, Adelaide, SA, AUSTRALIA

ABSTRACT

Analytical and simulated queuing models are useful for understanding and improving healthcare systems, including emergency departments (EDs), which serve as vital access points to healthcare systems internationally. EDs face challenges such as long waiting times, ambulance queuing, and bed-blocking, and simulation can be used to test strategies for improving these problems. Despite the widespread use of simulation, there is a lack of literature addressing the validity of queuing theory assumptions underpinning it, particularly the assumption of independence between the arrival and service processes. This paper employs semi-experiments to test these assumptions using real ED data. The results indicate that a correlation structure is present between the arrival and service processes in this data. The implications for simulation studies and directions for future work are discussed.

1 INTRODUCTION

Queuing theory (both analytical and simulation models) has been used extensively to model the dynamics of healthcare systems and experiment with improvement strategies. Emergency departments (EDs) are the most common component of healthcare systems studied in the simulation literature (Günal and Pidd 2010). EDs form a critical point of access to health care systems internationally. They frequently attract negative media attention in relation to long waiting times, ambulance ramping (queuing outside ED with paramedic staff caring for patients), and bed-blocking (queuing which arises from a lack of available beds in downstream health care facilities) (Vanberkel 2023).

In queuing theory, the arrival and service processes are assumed to be independent of each other. Despite the popularity of ED queuing models, there is a lack of literature that gives attention to the appropriateness of this assumption. Developing models that violate queuing theory assumptions could lead to incorrect conclusions being drawn about the real system. This paper contributes to the literature on ED simulation modelling through using semi-experiments (Ridoux et al. 2006) to test queuing assumptions on a motivating dataset, interpreting the experimental results, and discussing future research directions.

The rest of the paper is set out as follows. Section 2 discusses the relevant literature to this work and Section 3 sets out the semi-experiment method proposed for testing the feasibility of queuing assumptions in ED data. Section 4 presents information about the motivating data, results of the semi-experiments applied to the motivating data, and a semi-experiment with simulated data. Section 5 summarizes the implications of these results and suggests directions for future research.

2 BACKGROUND

2.1 Data-driven Queuing Models of Emergency Departments

There is a large volume of literature on emergency department (ED) queuing models (Salmon et al. 2018). The majority of papers use simulation over analytic methods due to the complex stochastic behavior present in EDs. There are two main approaches to developing ED simulation models.

979-8-3315-3420-2/24/$31.00 ©2024

The first approach focuses on creating a model that is specific to the ED being studied. Models commonly involve modelling entities (patients) as they progress through a sequence of activities, such as assessment by a nurse, medical testing, and treatment in a cubicle (Furian et al. 2018). Models which are specific to a particular ED are sometimes necessary to fulfil the simulation study objectives, for example if the aim is to optimize the use of the ED medical imaging facilities. The main disadvantages of this approach are that the model (i) will become outdated as processes change or the activity timing changes (Sinreich and Marmor 2004), (ii) cannot be used in another ED without carrying out clinical interviews and/or time-and-motion studies (Boyle et al. 2022), and (iii) it is difficult to adequately model the 'external' components (e.g., pharmacy, pathology, radiology) without also building a complete model of their behaviour, as they are subject to additional demands other than just those from the ED.

The second approach is data-driven, where the simulation model is developed using the data commonly collected in hospital electronic medical records. Examples of this approach include: Hoot et al. (2008) who developed a data-driven discrete event simulation that can be initialized in real-time and used to forecast several hours into the future; Boyle et al. (2022) who presented a data-driven framework for generalizable ED simulation modelling; and Heib et al. (2023) who presented a generalizable symbiotic simulation approach for EDs. The advantages of this approach are that the simulation can be updated with new data as it becomes available, and it is generalisable for use in most hospital EDs (Boyle et al. 2022; Heib et al. 2023). However, ED electronic medical record data presents several difficulties, further discussed in Section 2.2.

2.2 The Validity of Assumptions in Data-driven Healthcare Queuing Models

Simulation models in healthcare generally follow an $M/G/.$ queuing structure. There are fundamental differences between an $M/G/1$ queue and an $M/G/n$ queue, in that the output rate of an $M/G/n$ queue will appear to get faster with occupancy, while that of an $M/G/1$ queue will not. This does not carry over to service time, where they all still require their service time. But in a processor-sharing $M/G/1$ queue, everyone's service time will be significantly inflated because we can't tell when each arrival is getting (or not) service.

$M/G/.$ queueing models assume the arrival process to be Poisson-distributed and the arrival and service processes to be independent of each other. The first assumption is normally checked but Varney et al. (2019) highlighted that the second is often neglected. The authors present a semi-experiment approach (outlined in Section 3) and showed the presence of a correlation structure between the arrival process and length of stay distribution in 37 intensive care units, that could be attributed to occupancy management behaviors (for example demand-driven discharge). The study demonstrated that neglecting to account for dependence between the arrival and LOS processes in intensive care unit (ICU) simulation modelling leads to over-estimation of turnaway rates. A 'standard deviation ratio' approach was proposed as a means of accounting for this correlation structure, and was shown to improve simulation model result accuracy.

Occupancy management behaviors are well understood in the context of ICUs, where clinicians can discharge patients early to make space for more acutely unwell patients (Mallor and Azcárate 2014). In emergency departments (EDs) this type of behavior is less understood. A number of studies have hypothesized that there is a relationship between ED length of stay (LOS) and ED occupancy levels. McCarthy et al. (2009) used two approaches to study the effect of ED occupancy on each stage of ED LOS. The first 'static' approach used logistic regression to model the relationship between LOS and ED occupancy. The second 'time-varying' approach used a discrete-time survival model to investigate the relationship of between time-varying ED occupancy (measured every 15 minutes) and LOS. It was found that both the static and time-varying measures of occupancy were significant LOS predictors. Kao et al. (2015) developed a 'crowdedness index' (CI) to represent overcrowding in a Taiwanese ED. The CI was calculated as the ratio of 'current loading' to 'full capacity'. A DES of the ED was developed, where service times were modelled using gamma distributions. The authors indicated that the shape of these gamma distributions could be adjusted during the simulation depending on ED occupancy, but it is unclear whether, and if so how, this was achieved. Both of these papers showed a relationship between occupancy and LOS in EDs, but did not investigate whether the arrival and service processes are independent of each other.

The mechanisms through which load impacts service time have been summarized in the 'load effect on service times' (LEST) framework by Delasay et al. (2019). In the context of EDs it is natural to consider dependence between the service time distribution and system state, which could manifest in one or more of the following ways:

1. Server mechanisms include (i) 'demand-driven discharge' where patients are discharged at an increased rate to create space at high occupancy (Chan et al. 2019), (ii) prioritising the service of patients who are at risk of breaching time-based KPIs (Eatock et al. 2011), (iii) decrease in the number of diagnostic tests ordered as the number of waiting patients in ED increases (Batt and Terwiesch 2017) and (iv) reduction in productivity due to multitasking (simultaneous responsibility for the treatment of multiple patients) (Kc 2014).
2. Network mechanisms in the form of (i) sharing of resources between ED and other hospital units, for example radiology and pathology services, as well as specialist doctors (Hillier et al. 2009) and (ii) the ED is one node in a network of hospital units through which patient flow can become blocked by downstream congestion (Forster et al. 2003; Hillier et al. 2009).
3. Customer mechanisms for example (i) longer queue length in ED is associated with increased numbers of patients who did not wait for treatment (Batt and Terwiesch 2015).

The major barrier to understanding the effect of load on service time in EDs is the manner in which data is collected in hospitals. Administrative databases usually contain time stamps for arrival and/or triage, time first seen by a doctor and/or nurse, the time at which a decision is made to admit the patient, and the time of departure from ED. The patient journey through ED can then be separated into three LOS stages: (a) waiting, (b) treatment and (c) extended care (boarding for admitted patients or observation for discharged patients) (Boyle et al. 2022). Anecdotally, a further complication is a lack of accurate data recorded by hospitals on the number and composition of staff who have worked in each ED shift, making it difficult to understand and quantify the effect of service mechanisms without conducting empirical studies.

Distributions (b) and (c) contain both time spent waiting and time spent in service, with no way to decompose the amalgamated distributions. It is therefore difficult to identify and quantify the mechanisms through which load affects service and waiting times in ED using a data-driven approach. It is also difficult to determine because the state changes and so there is no clean way to measure the effect - effectively there is no clean definition of the server and hence of load.

Harper (2020) demonstrated that the service time distribution in a UK ED is related to the number of patients waiting for service, and proposed an analytical queueing model with switching thresholds to allow for a two-speed service time as a function of the number of waiting patients, defined here as the number of patients waiting in the emergency department for service at the time the patient starts service. It would be practically difficult to implement this model in a data-driven ED simulation, because data is not collected on the number, type, or tasks of servers in an ED.

In summary, some research has been conducted to explore the relationship between ED occupancy and service time distributions. However, there is no research which seeks to test the assumption of independence between ED arrival and service time processes. This will be tested using the semi-experiment method presented in Section 3.

3 METHODS

3.1 The Semi-experiment Method

The semi-experiment method was first introduced by Hohn et al. (2002) to investigate long-range dependence in time series of IP bytes and packets Hohn et al. (2002). The method of semi-experiments extended the concept of blockwise-shuffling, developed by Erramilli et al. (1996) to analyse long-range dependence in packet traffic. Ridoux et al. (2006) explained 'Typically, a semi-experiment involves replacing a single specific aspect of the real data with a simple, neutral model substitute. One then compares the statistics before and after, drawing conclusions on the role played by the structure removed by the manipulation'.

Varney et al. (2019) proposed the use of semi-experiments in a healthcare setting to investigate dependence between the arrival process and the length of stay (LOS) distribution in data from hospital Intensive Care Units (ICUs). The steps of this method are as follows:

1. Keep the arrival process stream of data exactly as it occurred.
2. Replace the original LOS data stream with a random permutation of the data.
3. Calculate the occupancy distribution of the hospital unit both before and after the manipulation was made.
4. Compare the statistics of the original and permuted occupancy distribution.

No significant change in the statistics of the original and permuted occupancy distributions indicates that there is no correlation structure between the arrival process and the LOS distribution. If the semi-experiment does show a change in the occupancy distribution statistics then 'these can only be attributed to the random permutation of the LOS stream of the semi-experiment, and thus the destruction of the delicate correlation structure between the arrival process and LOS stream' (Varney et al. 2019).

In EDs, the time between arrival and first seeing a doctor is waiting time. In this paper we consider the remaining time between first seeing a doctor and leaving the emergency department to be the service time. Therefore for the purpose of the semi-experiments we consider the time at which the patient first sees a doctor as the arrival process and the length of stay until they depart the ED as the service time process.

3.2 Patient Cohorts for Semi-experiments

In queuing theory, the series of inter-arrival times and service times are assumed to be independent and identically distributed (iid). This is often not a feasible assumption in healthcare data, but is adequately true when patients are categorized into appropriate 'cohorts' (Varney et al. 2019). To account for this as a possible source of correlation, the data can be segregated into patient cohorts before performing the semi-experiment. For example, triage category (representing different urgency levels) was used to divide the data into five cohorts. The semi-experiment is then performed by permuting LOS only within each patient cohort, recalculating the occupancy distributions, and comparing the statistics before and after the manipulation was made.

One important source of correlation could arise where the service time varies as a function of the queue length, as in Harper (2020). In our motivating dataset there is no information on the server rate or the service effort, so we cannot directly ascertain this information. We use ED occupancy and waiting room occupancy as a method of accounting for this source of correlation in the semi-experiments, as Varney et al. (2019) noted that 'assuming that the beds are the servers, and healing time and treatments are the service effort, then service effort should not be correlated to arrival rate other than through occupancy'.

- ED occupancy - calculate the total occupancy at each time point in the simulation. Create a variable that indicates the ED occupancy level at the time-point each patient arrived in the ED.
- Waiting room occupancy - calculate the number of patients in the waiting room at each time point in the simulation. Create a variable that indicates the waiting room occupancy level at the time-point each patient arrived in the ED.

Then use each occupancy level as a cohort in the semi-experiment by permuting LOS only within each patient group (e.g., LOS can only be swapped between patients who arrived when ED occupancy was 43, or LOS can only be swapped between patients who arrived when the waiting room occupancy was 12).

Note that both of these variables can be 'binned' e.g. ED occupancy could grouped into bins of size 5 to achieve 14 patient cohorts or size 10 to achieve 7 patient cohorts.

4 RESULTS

4.1 Motivating Data

The presence of a correlation structure between arrival and service processes was investigated using motivating data from the ED of a large teaching hospital in South Australia. The data contains information on the 119,306 patients who presented to the ED over the period of 20 months (Boyle et al. 2022). Clinical information (including triage category, primary disease classification and arrival method) were recorded in addition to time-stamps marking the progression of each patient through ED: triage, first occasion seen by doctor, admission and outcome. The work presented in this section was coded using R (R Core Team 2024).

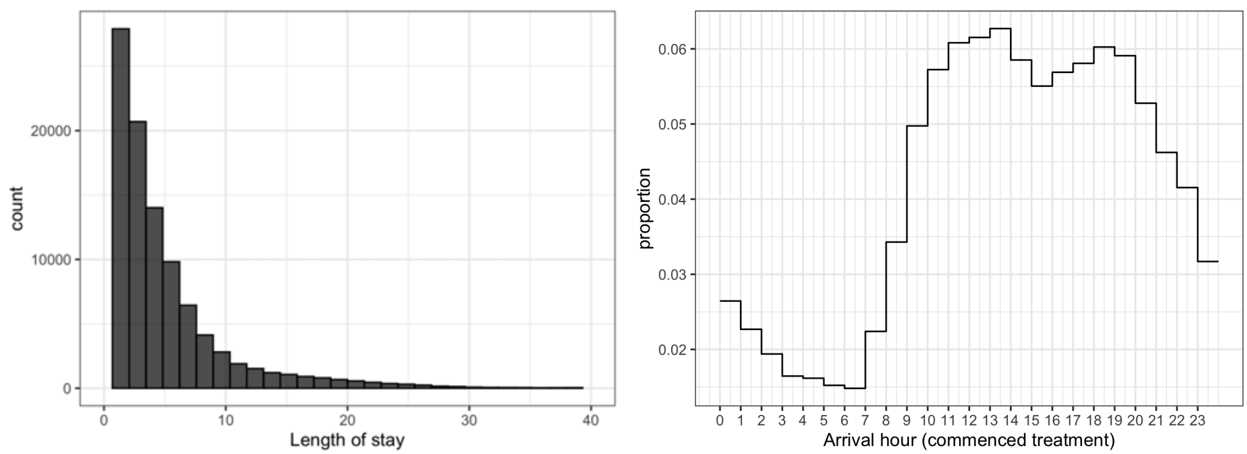

Figure 1: Emergency department length of stay in hours distribution between 'first seen by doctor' and departure (left); arrival process to 'first seen by doctor' timestamp by hour of day (right).

Figure 1 (left) shows the length of stay distribution between the 'first seen by doctor' and 'departure' timestamps. The distribution is strongly skewed. Figure 1 (right) shows the arrival process to the 'first seen by doctor' timestamp. The majority of arrivals occur during the day between 10am and 8pm with a drop in activity overnight. Both distributions exhibit typical ED behavior.

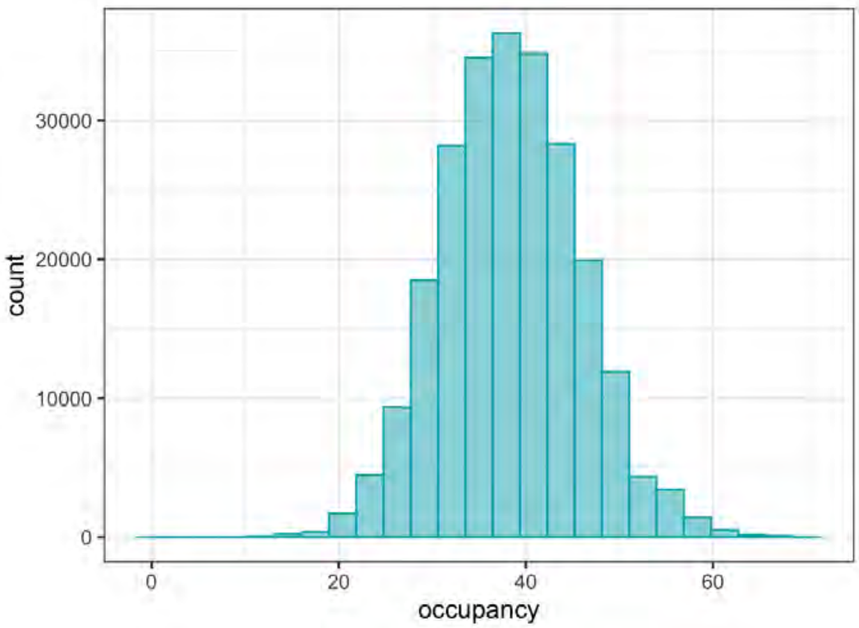

Figure 2: Distribution of emergency department occupancy.

Figure 2 displays the distribution of ED occupancy. The occupancy counts patients who have been seen by a doctor and have not yet departed from the ED. It excludes patients in the waiting room. The occupancy distribution is symmetric with a mean of 36.31 and a standard deviation of 7.64. Regression analysis was performed to investigate the relationship between both the ED occupancy and service time, and between waiting room occupancy and service time. ED occupancy and waiting room occupancy were both significant predictors of service time with a p-value of less than 0.01. However, the models explained less than 0.05% of the variation in service times.

4.2 Semi-experiments

The results of semi-experiments 1-5 are presented in Table 1 along with the statistics from the original data. Each of the semi-experiments was performed 20 times and a 95% confidence interval is included. Figure 3 shows the results from one run of semi-experiment 1 compared to the original data. Note that for semi-experiments 3 and 4 the data was not binned to achieve larger patient groups. This was tested during the analysis and did not change the conclusion of the semi-experiments.

For each of the five semi-experiments, a Kolmogorov-Smirnov (KS) test was applied to determine if the original occupancy distribution and the permuted occupancy distribution come from the same underlying distribution. The KS test calculates the maximum difference between the empirical cumulative distribution functions of the two samples under the null hypothesis H_0 that the two samples come from the same distribution. In all five cases, the KS test returned a p-value of less than 0.001%. There is therefore significant evidence to reject the null hypothesis and conclude that the semi-experiment results in a different distribution of occupancy than the original data stream.

Table 1: Statistics of the original occupancy distribution compared to five semi-experiments. 95% confidence intervals are indicated in brackets.

Semi-experiment	Mean	Standard deviation
Original data stream (no permutation)	36.31	7.64
Semi-experiment 1: permutation with no cohorts	36.31 (36.31,36.31)	9.48 (9.44,9.52)
Semi-experiment 2: permutation within triage category	36.31 (36.31,36.31)	9.46 (9.42,9.49)
Semi-experiment 3: permutation within ED occupancy level	36.31 (36.31,36.31)	8.87 (8.81,8.92)
Semi-experiment 4: permutation within waiting room occ	36.31 (36.31,36.31)	8.96 (8.94,8.98)
Semi-experiment 5: permutation within departure destination	36.31 (36.31,36.31)	9.20 (9.13,9.27)

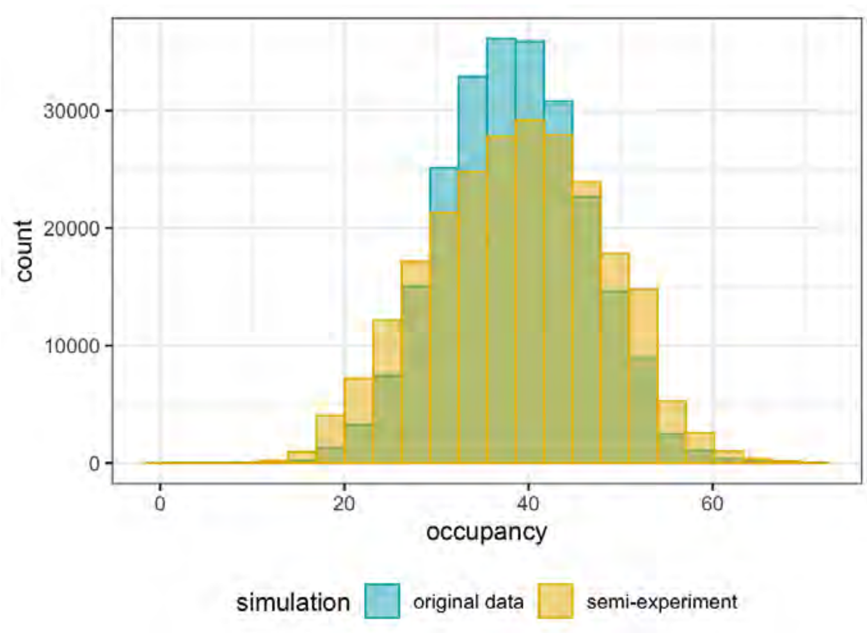

Figure 3: Occupancy distribution from 1 run of semi-experiment 1 compared to the original occupancy distribution from the data. The darker colour indicates where the two histograms overlap.

Table 1 and Figure 3 demonstrate that:

- In all semi-experiments the mean of the resulting occupancy distribution matched the mean of the original data.
- Figure 3 and the Kolmogorov-Smirnov test results demonstrate that the semi-experiment results in a different shape of occupancy distribution to the original data.
- The standard deviation of the occupancy after performing semi-experiments is approximately 16-24% greater than the standard distribution of the occupancy in the data.
- The addition of cohorts in semi-experiments 2-5 did not substantially reduce the standard deviation of the occupancy.
- Including ED occupancy as a cohort did not substantially reduce the standard deviation of the occupancy distribution.

Table 2 displays a closer look at semi-experiment five (permutation within departure destination) broken down by category into the discharged and admitted groups. Examination of the data by group shows that the standard deviation increases by 21% in the admitted group but decreases by 2% in the discharged group. Therefore, the standard deviation effect is much stronger in the admitted group.

Table 2: Statistics of the original occupancy distribution compared to semi-experiment five. 95% confidence intervals are indicated in brackets.

	Original data		Semi-experiment	
	Mean	Standard deviation	Mean	Standard deviation
All data	36.31	7.64	36.31 (36.31,36.31)	9.20 (9.13,9.27)
Discharged group	10.96	5.03	10.96 (10.96,10.96)	4.92 (4.90,4.94)
Admitted group	25.35	5.31	25.35 (25.35 ,25.35)	6.43 (6.37,6.48)

4.3 Semi-experiments on a Simulated Queue

The semi-experiments presented in Section 4.2 were performed with real ED data. We also performed semi-experiments on a simulated queue for comparison. For simplicity, and considering the difficulty with defining and measuring servers in EDs, an $M/M/\infty$ queue was used. Parameters were chosen to be a relatively close match to the data; the service time distribution was modelled as an exponential distribution with rate parameter 0.0038 and the inter-arrival time distribution was modelled as an exponentially distribution with rate parameter 0.1370. The number of observations generated was 10,000 and the semi-experiment was performed 20 times with confidence intervals.

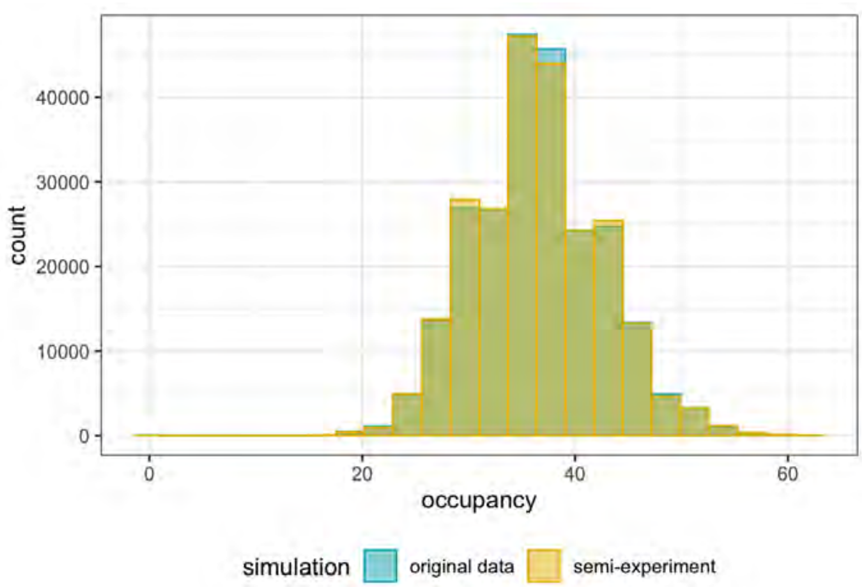

Figure 4: Occupancy distribution from 1 run of the semi-experiment compared to the occupancy distribution of the simulated $M/M/\infty$ queue.

Table 3 and Figure 4 demonstrate that both the mean and the standard deviation of the original simulated occupancy distribution are recovered in the semi-experiments. The results presented here show one realisation of a simulated $M/M/\infty$ queue, but simulating from different random number seeds yields the same conclusion.

Table 3: Statistics of the simulated $M/M/\infty$ occupancy distribution compared to the semi-experiment. 95% confidence intervals are indicated in brackets.

Semi-experiment	Mean	Standard deviation
Original data stream	36.41	6.24
Permutation with no cohorts	36.41 (36.41,36.41)	6.21 (6.19,6.24)

5 CONCLUSIONS AND FUTURE WORK

The semi-experiments presented in Section 4.2 show that keeping the emergency department (ED) arrival data stream fixed and randomly permuting the service times results in different occupancy distribution to the original data - the mean is recovered but the standard deviation is inflated. When the service time data is only permuted within cohorts (e.g., triage category, departure destination, occupancy level), the standard deviation is still inflated. This demonstrates that the arrival and service time processes are not independent. It also shows that the average of the distribution is captured but not the variability. Semi-experiments 3 (permutation within ED occupancy level) and 4 (permutation within waiting room occupancy) show that this correlation structure is not simply a function of service time varying with queue length.

A closer look at semi-experiment 5 (permutation within departure destination) in Table 2 shows that the inflation in standard deviation exists only in the admitted group of patients. This suggests that the correlation structure between the arrival and service time processes is present in the admitted group of patients and not the discharged group. As a comparison, Section 4.3 shows the results of semi-experiments on data from an $M/M/\infty$ queue. The semi-experiments here show that both the mean and the standard deviation of the occupancy distribution were recovered. This is further evidence that the inflation in standard deviation is present.

Failing to account for this correlation structure in designing queuing models can result in inaccurate outcome measures and subsequent errors in drawing conclusions from scenario analysis (Varney et al. 2019). A limitation of this study is that it was isolated to data from one ED. Future work will seek to apply the semi-experiment method to multiple case study EDs for comparison. Although the study is limited to one ED, similar results have been detected in intensive care units (Varney et al. 2019) and this correlation structure is likely to exist in many healthcare systems. Finding a method to accurately detect, interpret, and model it is crucial. The next steps in this project will be to (i) perform further semi-experiments and hypothesize the source of the correlation, (ii) conduct a simulation study with this correlation structure to draw conclusions about the results, and (iii) develop methods to model this dependency structure in healthcare simulation models.

REFERENCES

Batt, R. J. and C. Terwiesch. 2015. “Waiting Patiently: An Empirical Study of Queue Abandonment in an Emergency Department”. *Management Science* 61(1):39–59 https://doi.org/10.1287/mnsc.2014.2058.

Batt, R. J. and C. Terwiesch. 2017. “Early Task Initiation and Other Load-adaptive Mechanisms in the Emergency Department”. *Management Science* 63(11):3531–3551 https://doi.org/10.1287/mnsc.2016.2516.

Boyle, L. M., A. H. Marshall, and M. Mackay. 2022. “A Framework for Developing Generalisable Discrete Event Simulation Models of Hospital Emergency Departments”. *European Journal of Operational Research* 302(1):337–347 https://doi.org/10.1016/j.ejor.2021.12.033.

Chan, C. W., L. V. Green, S. Lekwijit, L. Lu and G. Escobar. 2019. “Assessing the Impact of Service Level when Customer Needs are Uncertain: An Empirical Investigation of Hospital Step-down Units”. *Management Science* 65(2):751–775 https://doi.org/10.1287/mnsc.2017.2974.

Delasay, M., A. Ingolfsson, B. Kolfal, and K. Schultz. 2019. “Load Effect on Service Times”. *European Journal of Operational Research* 279(3):673–686 https://doi.org/doi.org/10.1016/j.ejor.2018.12.028.

Eatock, J., M. Clarke, C. Picton, and T. Young. 2011. “Meeting the Four-hour Deadline in an A&E Department”. *Journal of Health Organization and Management* 25(6):606–624 https://doi.org/10.1108/14777261111178510.

Erramilli, A., O. Narayan, and W. Willinger. 1996. “Experimental Queueing Analysis with Long-range Dependent Packet Traffic”. *IEEE/ACM Transactions on Networking* 4(2):209–223 https://doi.org/10.1109/90.491008.

Forster, A. J., I. Stiell, G. Wells, A. J. Lee and C. Van Walraven. 2003. “The effect of hospital occupancy on emergency department length of stay and patient disposition”. *Academic Emergency Medicine* 10(2):127–133 https://doi.org/doi.org/10.1197/aemj.10.2.127.

Furian, N., D. Neubacher, M. O’Sullivan, C. Walker and C. Pizzera. 2018. “GEDMod–Towards a Generic Toolkit for Emergency Department Modeling”. *Simulation Modelling Practice and Theory* 87:239–273 https://doi.org/doi.org/10.1016/j.simpat.2018.07.010.

Günal, M. M. and M. Pidd. 2010. “Discrete Event Simulation for Performance Modelling in Health Care: A Review of the Literature”. *Journal of Simulation* 4:42–51 https://doi.org/10.1057/jos.2009.25.

Harper, P. R. 2020. "Server Behaviours in Healthcare Queueing Systems". *Journal of the Operational Research Society* 71(7):1124–1136 https://doi.org/10.1080/01605682.2019.1567653.

Heib, A. R., C. S. Currie, B. S. Onggo, H. K. Smith and J. Kerr. 2023. "A Generalized Symbiotic Simulation Model of an Emergency Department for Real-Time Operational Decision-Making". In *2023 Winter Simulation Conference (WSC)*, 1042–1053. IEEE https://doi.org/https://doi.org/10.1109/WSC60868.2023.10407981.

Hillier, D. F., G. J. Parry, M. W. Shannon, and A. M. Stack. 2009. "The Effect of Hospital Bed Occupancy on Throughput in the Pediatric Emergency Department". *Annals of Emergency Medicine* 53(6):767–776 https://doi.org/doi.org/10.1016/j.annemergmed.2008.11.024.

Hohn, N., D. Veitch, and P. Abry. 2002. "Does Fractal Scaling at the IP Level Depend on TCP Flow Arrival Processes?". In *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet Measurment*, 63–68 https://doi.org/10.1145/637201.637208.

Hoot, N. R., L. J. LeBlanc, I. Jones, S. R. Levin, C. Zhou, C. S. Gadd *et al*. 2008. "Forecasting Emergency Department Crowding: A Discrete Event Simulation". *Annals of Emergency Medicine* 52(2):116–125 https://doi.org/10.1016/j.annemergmed.2007.12.011.

Kao, C.-Y., J.-C. Yang, and C.-H. Lin. 2015. "The Impact of Ambulance and Patient Diversion on Crowdedness of Multiple Emergency Departments in a Region". *PLoS One* 10(12):e0144227 https://doi.org/10.1371/journal.pone.0144227.

Kc, D. S. 2014. "Does Multitasking Improve Performance? Evidence from the Emergency Department". *Manufacturing & Service Operations Management* 16(2):168–183 https://doi.org/10.1287/msom.2013.0464.

Mallor, F. and C. Azcárate. 2014. "Combining Optimization with Simulation to Obtain Credible Models for Intensive Care Units". *Annals of Operations Research* 221:255–271 https://doi.org/10.1007/s10479-011-1035-8.

McCarthy, M. L., S. L. Zeger, R. Ding, S. R. Levin, J. S. Desmond, J. Lee *et al*. 2009. "Crowding Delays Treatment and Lengthens Emergency Department Length of Stay, Even Among High-acuity Patients". *Annals of Emergency Medicine* 54(4):492–503 https://doi.org/10.1016/j.annemergmed.2009.03.006.

R Core Team 2024. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.

Ridoux, J., A. Nucci, and D. Veitch. 2006. "Seeing the Difference in IP Traffic: Wireless Versus Wireline". In *Proceedings IEEE INFOCOM 2006. 25TH IEEE International Conference on Computer Communications*, 1–12. April 23rd-29th, Barcelona, Spain. https://doi.org/10.1109/INFOCOM.2006.292.

Salmon, A., S. Rachuba, S. Briscoe, and M. Pitt. 2018. "A Structured Literature Review of Simulation Modelling Applied to Emergency Departments: Current Patterns and Emerging Trends". *Operations Research for Health Care* 19:1–13 https://doi.org/10.1016/j.orhc.2018.01.001.

Sinreich, D. and Y. N. Marmor. 2004. "A Simple and Intuitive Simulation Tool for Analyzing Emergency Department Operations". In *2004 Winter Simulation Conference (WSC)*, Volume 2, 1994–2002. IEEE https://doi.org/10.1109/WSC.2004.1371561.

Vanberkel, P. 2023. "Emergency Department Crowding Has Gone Beyond Hallways onto Ambulance Ramps. Now there's Nowhere Left to Wait". https://theconversation.com/emergency-department-crowding-has-gone-beyond-hallways-onto-ambulance-ramps-now-theres-nowhere-left-to-wait-200420. Accessed: 2024-04-12.

Varney, J., N. Bean, and M. Mackay. 2019. "The Self-regulating Nature of Occupancy in ICUs: Stochastic Homoeostasis". *Health Care Management Science* 22:615–634 https://doi.org/10.1007/s10729-018-9448-4.

AUTHOR BIOGRAPHIES

LAURA BOYLE is an Lecturer in Data Analytics in the Mathematical Sciences Research Centre at Queen's University Belfast and an adjunct lecturer in the School of Computer and Mathematical Sciences at the University of Adelaide, Australia. Her research interests include simulation, data analytics, and operational research with applications in healthcare. Her email address is laura.boyle@qub.ac.uk and her website is https://pure.qub.ac.uk/en/persons/laura-boyle/.

NIGEL BEAN is an adjunct Professor in the School of Computer and Mathematical Sciences at the University of Adelaide, Australia. His research interests include Markov chains, matrix analytic methods, optimisation, queuing theory, and stochastic modelling. His email address is nigel.bean@adelaide.edu.au and his website is https://researchers.adelaide.edu.au/profile/nigel.bean/.