Proceedings of the 2017 Winter Simulation Conference
W. K. V. Chan, A. D'Ambrogio, G. Zacharewicz, N. Mustafee, G. Wainer, and E. Page, eds.

RATIONALIZING HEALTHCARE BUDGETING WHEN PROVIDING SERVICES WITH MANDATED MAXIMUM DELAYS: A SIMULATION MODELING APPROACH

Philip Troy
Lisa Westaway
Anthonie Grondin
Thaddeus Rezanowicz

Centre intégré universitaire de santé et de services sociaux du Centre-Ouest-de-l'Île-de-Montréal
3755 Ch de la Côte-Sainte-Catherine
Montreal, QC H3T 1E2, CANADA

ABSTRACT

To determine the budget needed by a healthcare network to provide government mandated mental health services, a simulation model of those services was built, verified and validated; it was then used to identify where mandated delivery times were not being met and where staff should be reallocated. In addition to the obvious benefits of this approach, a less obvious benefit was that the discovery process needed to build the model identified additional opportunities for providing better care with less resources. A further benefit of this work was the potential, recognized by the chief financial officer, that this approach could be used throughout the network to rationalize staffing levels, and thus make it possible to provide more, better or timelier outcomes with the same resources throughout the network.

1 INTRODUCTION

In 2005, Canada's provinces and territories, which are responsible for managing and funding public healthcare, announced common goals for performing several medical procedures within pre-defined periods of time. In particular, those governments announced (Ontario 2005) goals for performing "radiation therapy to treat cancer within four weeks of patients being ready to treat, hip fracture fixation within 48 hours, hip replacements within 26 weeks, knee replacements within 26 weeks, surgery to remove cataracts within 16 weeks for patients who are at high risk, breast cancer screening for women aged 50 to 69 every two years, and cervical cancer screening for women aged 18 to 69 every three years after two normal tests." Not surprisingly, since that announcement many of those governments have struggled to provide the resources needed to achieve those goals (CIHI 2017, Sutherland and Repin 2014). Contributing to that struggle is the difficulty of finding an appropriate approach for determining healthcare budgets.

In an ideal world it would be easy to argue that zero based budgeting (Wetherbe and Montanari 1981) should be used to budget healthcare expenditures. In zero based budgeting, budgets for products or services are based on the cost of the resources used to provide those products or services, multipled by the number of those products or services provided per fiscal period.

Unfortunately, there are many challenges to applying zero based budgeting, particularly when government provided information systems do not track important process drivers such as the time spent preparing for client interactions, charting those interactions, and training and administrative activities. Because of this, in Quebec it is not always possible to determine the amount of resources used to address each client's needs, or the time actually available for providing services.

Another challenge to applying zero based budgeting to healthcare delivery arises when services, whose demand varies over time, need to be provided within pre-defined periods of time. Sources of demand variability include the instantaneous variablity around the mean rate at which new requests occur, as well as the variability in the mean rate itself which may change with the season or the day of the week. They also include the variability in the amount of resources needed to provide each type of service, which is a function of patients' ages, levels of risk, co-morbidities, and family support, all of which are inherently highly variable. And while it is conceivably possible to achieve near to 100% utilization of resources when there is no variability, it is more difficult to do so when there is variability and when resource levels are set to address peak or near peak demand levels. This is because variability under those conditions will almost always result in periods of time during which some resources are not used. And while it might seem useful to just adjust budgets using a utilization percentage, the value this percentage should be set to is highly contingent on the length of time during which the service needs to be provided, the size of the pool of resources providing the services, and the extent of each type of variability. Thus there do not appear to any readily available rules for setting this percentage.

While the challenges described above are faced by many healthcare groups, in the spring of 2016 they were specifically faced by the management of the mental health and addiction directorate at the Centre intégré universitaire de santé et de services sociaux (CIUSSS) du Centre-Ouest-de-l'Île-de-Montreál, the McGill University affiliated integrated health and social services network of west central Montreal. In particular, they needed to determine staffing requirements in context of the Mental Health Action Plan of 2015 - 2020, which required centralization of access to mental health services within the CIUSSS via an intake center, without providing additional resources for staffing that center. This plan also required that all intake requests be oriented to first line specialized mental health services within seven days, or 48 hours for more urgent requests, and that when needed, first line psychiatric nurses, social workers, or psychologists meet with clients within 30 days, or within one week in the case of urgent needs. Compounding these challenges was that demand for mental health services provided by the CIUSSS differed by location within the CIUSSS, in addition to varying for the reasons discussed above. Finally, these challenges also needed to be addressed in context of an ongoing search for budget cuts by the CIUSSS's chief financial officer.

Given the limitations of zero based budgeting for such an environment, a member of the CIUSSS's analytics team proposed, and mental health management agreed to, a project in which a discrete event simulation would be built to model the interactions between resources used by the process. As part of the project, after the model was verified and validated, it would be used to determine the number of each type of staff needed in each location, in context of the variability in demand and the government's mandates for the period of times during which each type of service was to be provided. In order to validate the model it would be necessary to statistically determine the parameters of existing mental health processes, or to otherwise obtain estimates of those parameters where data for statistical analysis did not exist.

After the model was built, verified and validated, it was used to rigorously determine where there was an excess of resources and where addditional resources were needed. In addition, the act of developing an understanding of the current processes to build the simulation led to questions about the current processes. This in turn led to a search for ways of improving the processes.

This is not the first time that simulation has been applied to healthcare. In fact, the use of simulation for healthcare goes back more than five decades to Fetter and Thompson (1965), who used it for evaluating hospital tactical decisions. More recently, Jacobson, Hall, and Swisher (2013) authored a detailed "overview of discrete-event simulation modeling applications to health care clinics and integrated health care systems (e.g. hospitals, outpatient clinics, emergency departments, and pharmacies) over the past forty years." Particularly relevant to this project is a tutorial on the use of discrete event simulation for health policy design and decision making by Ramwadhdoebea, Buskensb, Sakkers, and Stahl (2009), as well as articles on the use of discrete event simulation for mental health and community planning services (Meadows et al. 1997, Kim et al. 2013). Also relevant to this project, Bedoya-Valencia and Kirac (2016) discussed the successful reallocation of resources in an emergency department using discrete event simulation to reduce patient length of stay and time to be seen by a physician or physician assistant. In addition, in an attempt to apply lessons learned from the use of simulation for manufacturing, Fowler and Monch (2015) compared the use of simulation for manufacturing and healthcare and pointed out that for healthcare "there is a larger need to model the stochastic behavior of resources and working objects" and that "the more complicated service structure of healthcare is a serious modelling issue."

In context of this literature, it appears that this is the first time that simulation has been applied on such a granular level for a large healthcare network to determine the number of each type of staff needed each day of the week at each location to ensure that maximum delays are not exceeded. And when this approach was presented to the CIUSSS's chief financial officer, her reaction was that all of the CIUSSS's operating units should consider similar analyses for improving processes and determining staffing needs.

This paper describes different aspects of the project, with an explicit focus on describing them, and an implicit focus on suggesting the potential value of similar projects for other healthcare services. In particular, section 2 discusses the request for data, section 3 discusses the simulation model, section 4 discusses the steps taken to prepare data for use in the model, and section 5 discusses verification and validation of the model. Then, section 6 discusses experimentation performed with the model, and section 7 discusses conclusions and the opportunities found for further improvements.

2 THE REQUEST FOR DATA

To build a discrete event simulation model for this project, it was necessary to determine the relative likelihoods of each possible value that would affect the amount of time between events, as well as the characteristics of those events. For this simulation, the type of events are listed in Table 1.

Table 1: Types of events modeled by the simulation.

Type Of Event
Receipt of new requests each day
Start and end of staff sick days
Start and end of staff meetings
Start and end of staff vacations
Enqueuing of requests
Start and end of interventions
Start and end of planned delays

Table 2: Information needed for the model.

Status	For All	Information Needed For The Model
* ***	Requests	date of request, episode id* ***
* ***	Interventions	staff type, intervention id, amount of staff time required, episode id
* **	Interventions	intervention id, amount of staff time required for transportation
* **	Interventions	intervention id, amount of staff time required for reporting
	Interventions	intervention id, intervention type
	Interventions	intervention id, location
	Staff at each location	staff type, staff id, weekdays worked
***	Staff at each location	staff id, date and time taken off for illness
***	Staff at each location	staff id, date and time taken off for vacation
***	Meetings	start and length of each meeting for each staff type at each location
	Statuary holidays	date

During meetings subsequent to requesting data (see Table 2), it was discovered that some of the data was not available (*), that it would be more feasible to use average times specified by mental health team members for some types of data(**), or that another approach should be used to create some of the data (***). With respect to the last grouping, new episodes were identified as occurring for a particular client whenever there was an intervention for that client after an interval of six months or more without interventions. It was also decided that staff members' maximum allotment of ten days of sick time would be randomly distributed throughout the year with a heavier likelihood during the winter months, that staff meetings would be spaced regularly with small random pertubations so that part time staff would occasionally attend those meetings, and that vacation times would be staggered to preclude, except for the summertime, two staff members of the same type at the same location taking vacation at the same time.

It is worth noting that for future analysis it would be better if all fields except for the vacation and illness data were to be collected.

3 THE MODEL

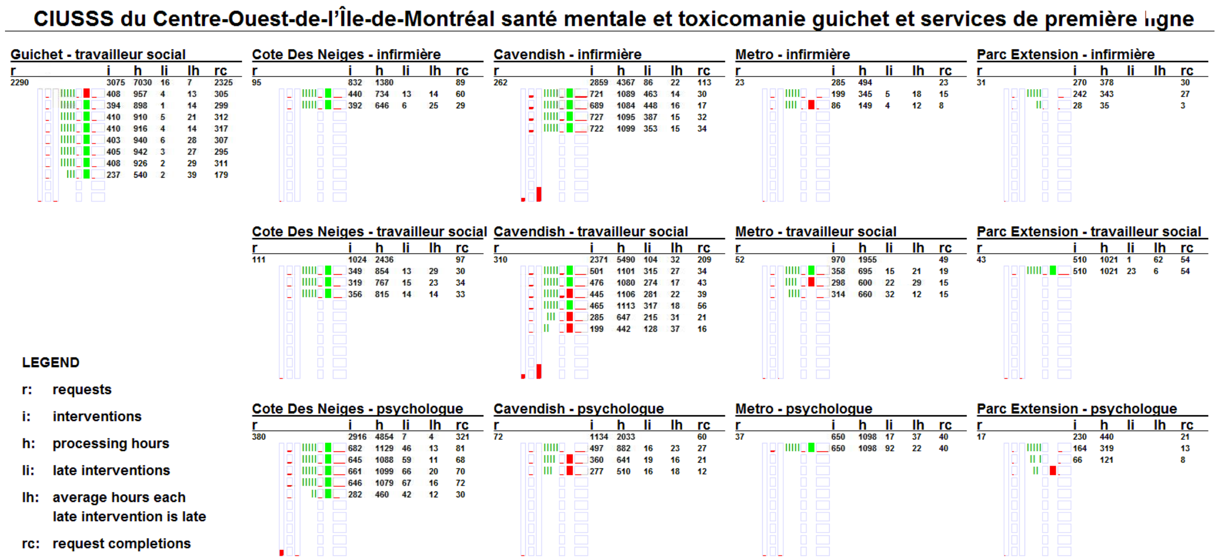

Figure 1: A visual representation of the model.

The model is best understood by viewing its display (see Figure 1) which is organized so there is one location (on the left) for the intake center (referred to as the guichet), and an additional location for each of the four physical locations at which first line services are provided. Within each of the latter locations are areas for the types of mental health staff that work at those locations (see Figure 2 which displays the area for psychiatric nurses at the Cavendish location). Just below the row of headers (r, i, h, li, lh, and rc) displayed for each area, the number of requests, the number of interventions, the hours worked, the number of late interventions, the average late hours of each late intervention, and the requests completed year to date for the whole area are displayed. Below that row, there are two queues displayed as tall and narrow rectangles that span the remaining rows of the area; each red dot displayed in those queues (and in the other queues discussed below) represents one request. The leftmost common queue is for requests for the area that have not yet been assigned to specific staff. The rightmost common queue is for unassigned requests that have been prioritized because they have already waited the maximum allowable time.

In the same rows in which the common queues are displayed, parameters and information for each staff member in the area are displayed. The leftmost staff specific rectangle displays the queue of requests previously assigned to the staff member that are currently waiting for a secondary intervention. The narrow green vertical bars represent the days of the week the staff member works. The next staff specific rectangle displays the queue of all requests that have been assigned to the staff member and given higher priority because they have already waited the maximum allowable time. The following staff specific rectangle represents the state of the staff member and is red, green or clear when the staff member is not available, is working, or has nothing to do, respectively. The wide staff specific rectangle to the right displays one red dot per request being delayed for a week until it is enqueued for an additional intervention by the same staff member. The statistics displayed next are the number of interventions completed, the number of hours worked, the number of late interventions, the average late hours for each late intervention, and the number of requests completed for the staff members.

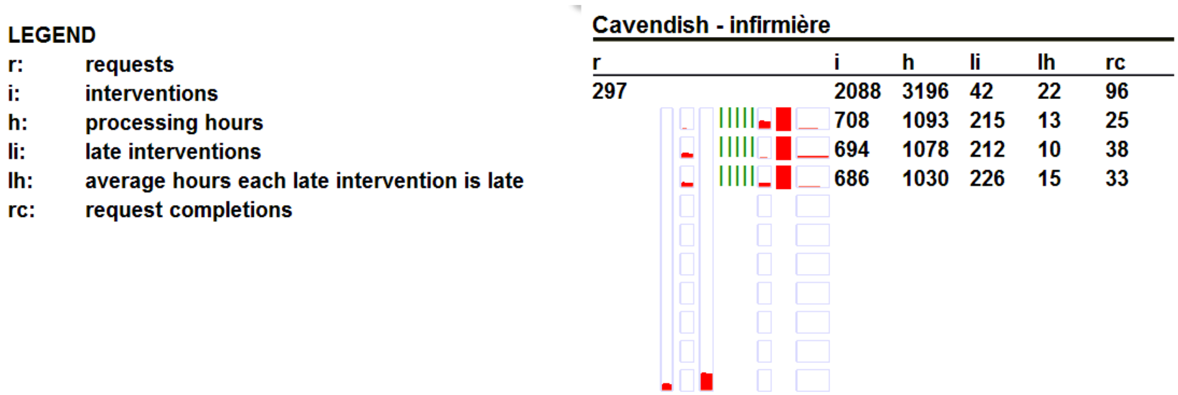

Figure 2: A visual representation of the model for one type of staff in one location.

As is typical of many discrete event simulation models, this model is driven by activities that are triggered by events. Each workday morning there is an event that generates new requests for each location and each client profile, based on the analysis of new requests described in the next section. At that time, all random quantities for each new request, including the urgency, number of interventions from each type of staff for each request and amount of time for each intervention, are randomly generated. The requests are then moved to the unallocated queue in the guichet, the leftmost common queue, where they wait to be assigned to and then processed by a guichet staff member. If a request waits longer in the unallocated request queue than the time mandated for it by the government, an event is generated that moves the request to the second unallocated request queue in the guichet, the rightmost common queue, where it waits with a higher priority for a guichet staff member. In response to requests being added to the unallocated request queue, logic is executed to determine if there currently is a guichet staff member that can process a request immediately. If there is, the request is moved to that staff member's high priority request queue which immediately feeds it to the staff member's activity (since the staff member is not busy at the time of the move) for processing.

When a staff member activity completion event occurs, logic is used to determine whether all activities required by the request have been completed, and if so the request is counted as a completed request. If the activities have not been completed, logic is used to determine whether the request's remaining activities can be processed by the same staff member in the guichet, by determining whether it only requires four or less interventions by a social worker and the extent of backlogs in the guichet. If the request requires additional activities that can be performed by the staff member in the guichet, the request is sent to that staff member's delay area where it is held for a week, after which it is sent to that staff member's allocated request queue. If the request requires additional interventions that cannot be handled at that time by the staff member in the guichet, it is sent to the low priority (leftmost) common request queue of the correct staff member type at the appropriate location, where it is handled in pretty much the same way as it was handled in the guichet. If the request needs processing by more than one type of staff member, the request is cloned to the low priority common request queue of each appropriate staff type. In contrast to the limit for the guichet, there is no limit on the number of interventions first line staff members can perform for specific clients, though larger numbers of interventions for individual clients are likely to lead to longer delays for the start of interventions for other clients.

Vacations, sick leave, and meetings are also handled by the start of day event handler which invokes the simulation package's breakdown capability. That capability then turns off the appropriate staff members for the time of the vacation, sick leave or meeting, and schedules an event for turning the staff members back on at the appropriate time.

While building the model, special attention was made to ensuring that the model's animation provided management with information needed to reallocate staff, i.e., staff utilization, the number of requests not handled on time, and the average lateness of late requests. Thus the animation of each area (see the first row of numbers in Figure 2) also includes statistics (discussed above) for the area as a whole as well as for each individual staff member. These indicators can be set to be updated either while the simulation is running, at the end of each run, or at the end of each trial. In addition to encouraging managers to view the simulation's animation while it was running and at the end of a trial, managers were encouraged to scroll through and look at the results of individual runs so they could obtain a better feeling for the variability of results between runs in a trial.

4 DATA CLEANING AND CHARACTERIZATION

A considerable amount of effort was required to clean and characterize the data so that it could be used by the model. This effort started with data cleaning operations listed in Table 3.

Table 3: Data cleaning operations.

Operation
Harmonizing client ids (so that leading letters and zeroes where removed)
Changing all text to lower case (so that appropriate data not in the same case could be matched)
Filling in empty cells (since provided data was formatted for reporting instead of analysis)
Standardizing date formats
Standardizing values in other fields
Replacing birth date data with birth year data (for anonymyzation purposes)
Replacing 6 digit postal codes with three digit postal codes (for anonymyzation purposes)
Converting intervention times from hh:mm format to minute format
Consolidating multiple records of data for a single intervention into a single record
Removing irrelevant interventions
Recoding long telephone interventions as face to face interventions
Assigning episode numbers to interventions
Assigning profile numbers to individual interventions

After cleaning the data, the number of face to face interventions and the amount of time between the start of those interventions was computed for each episode and staff type. Then, a cluster analysis was performed on the episodes based on these factors. (See Figure 3 for the results of clustering episodes by the number of interventions and the average time of all interventions between the start of face to face interventions for nursing and social worker interventions at one location.) For the clustering, all values were normalized to between 0 and 1. For the display, the height and width of bubbles are divided by a constant so as to minimize the amount large bubbles hide smaller bubbles, the x axis is used to indicate the number of psychiatric nurse (infirmiere) face to face interventions, the y axis is used to indicate the number of face to face social worker (travailleur social) interventions, the height of bubbles is used to indicate the average intervention time between the start of psychiatric nurse face to face interventions, and the width of bubbles is used to indicate the average intervention time between the start of social worker face to face interventions. As the clustering did not seem to effectively characterize the episodes, episodes were instead assigned profiles based on the mix of staff types that performed interventions during an episode (see Table 4). All subsequent characterization of the data was performed separately for episodes grouped into those profiles at each location.

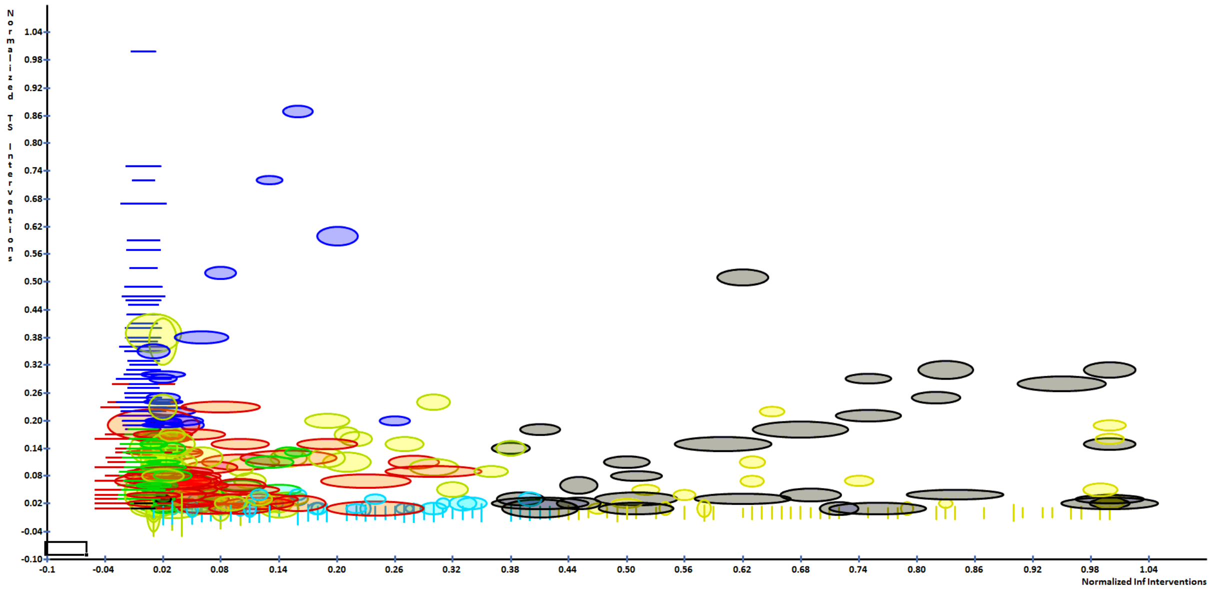

Figure 3: Bubble diagram displaying results of clustering of episodes.

Table 4: Episode profiles based on the type of staff that intervened for a particular episode of care.

Profile Number	Profiles
1	Psychiatric nurse
2	Social worker
3	Social worker and psychiatric nurse
4	Psychologist
5	Psychologist and psychiatric nurse
6	Psychologist and social worker
7	Psychologist and social worker and psychiatric nurse

Given the profiles used to group client interventions, the next step was to characterize the number of new requests received each day into the guichet for each profile and location, for the two fiscal years during which data was collected. Note that there were 251 working days in each fiscal year and that the fiscal years started at the beginning of April. (See Figure 4.) As initial attempts to detect seasonality by month of the year or by the fall, winter, summer and spring seasons did not yield statistically significant results, a twenty one day moving average of the data was computed and a seasonal pattern that occurred both years was identified. (See Figure 5.) (Note that there only appears to be one line between days 201 and 236 because there was only data for one of those years during that period of time, and thus the average for that period of time exactly overlaps the line for that data.)

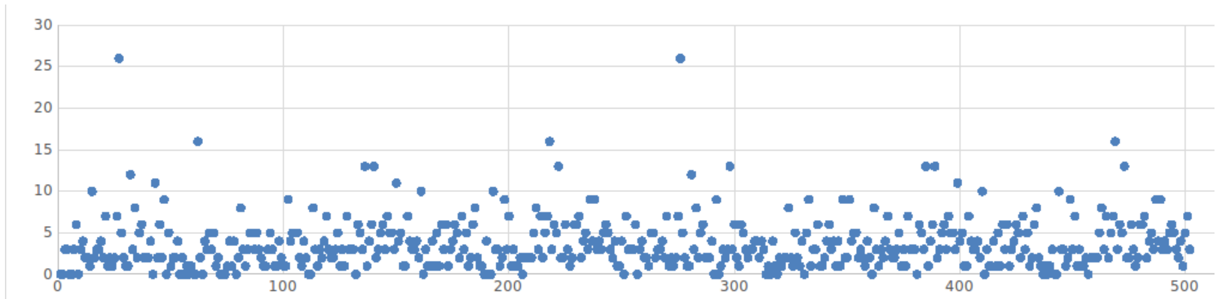

Figure 4: Scatter diagram of the number of reqests received at one of the locations on each of the 502 working days analyzed.

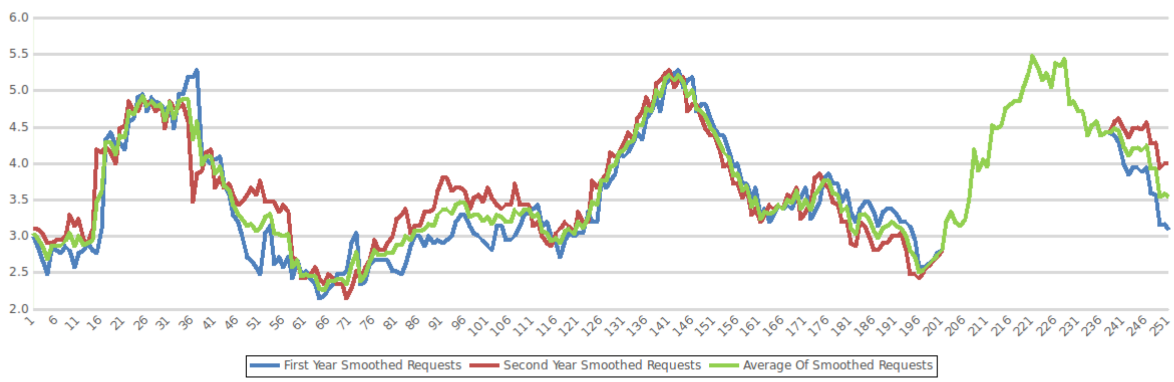

Figure 5: Line graph of twenty one working day moving average of the total number of received requests.

Having identified that pattern, the twenty one day moving average of new requests for each profile and location was used to normalize the data. A linear regression analysis was performed on the results of the normalization using day of the week dummy variables; that regression detected a small day of the week effect. Since there was no expectation that they would be normally distributed, an inverse Bezier distriution (Wagner and Wilson 1993) of the residuals was computed to make it possile to randomly generate the number of new requests each day.

A similar but simpler approach was used to characterize the number of face to face interventions, and the total intervention time between the start of face to face interventions, for episodes of each location and profile. In particular, the mean of both quantities and an inverse Bezier distribution of both sets of residuals was computed to make it possible to randomly generate those quantities.

5 VERIFICATION AND VALIDATION OF THE MODEL

Sargent (2011) explicitly states that "computerized model verification ensures that the computer programming and implementation of the conceptual model are correct." The steps taken to ensure this include: displaying the state of the simulation as well as the cumulative results to make it easy to check the state of the simulation at any time during a simulation run; stepping through the model to make sure that individual requests were fully processed and that staff members attended meetings, went on vacations and got sick; presenting the model for review to other simulation modelers; carefully reviewing the model logic several times; testing the model with too few staff members to ensure that backlogs were created; and testing the model with too many staff members to ensure that significant backlogs did not occur.

Schlesinger et al. (1979) defines validation as "substantiation that a computerized model within its domain of applicability possesses a satisfactory range of accuracy consistent with the intended application of the model." For this analysis, the validation of the model was considerably more difficult than its verification because the process being modelled had new components for which results data did not exist. In particular, there was no historical data for transportation, charting, and reporting times, requests received per day data was missing for some months and locations, and existing data to recreate the state of the system that existed at the start of the data collection period was not readily available.

Steps taken to validate the model included informing stakeholders in advance of the need for their involvement in the validation of the model, reviewing with stakeholders the analysis of seasonality and day to day variability in the number of new requests, the profiling mechanism, the variability in the number of each type of intervention, and the variability in the amount of total intervention time between the start of face to face interventions. Results of simulation trials were also reviewed with stakeholders who found that in all but one case, the number of psychiatric nurses needed at the Cavendish location, results made sense and corresponded with their expectations. These results included the number of new requests per year, the number of interventions performed each year by each staff type at each location, and the time spent performing those interventions. Where the results did not make sense, further effort is being expended to identify the issues, and adjustments to the results discussed in this paper will be made based on an analysis of additional simulation runs.

6 EXPERIMENTATION

The results of simulating the existing configuration can be seen by looking at the numbers for each staff type at each location in Table 5. Note that all results, except for full time equivalent staff counts (FTEs), are rounded to the nearest integer; thus it is possible to get an average of 0 late interventions with an average greater than 0 of late hours per intervention. These results were tabulated from the second year of each simulation run, after first running the simulation for one year so that the simulation at the start of the data collection period would more closely resemble the state of the real system.

Table 5: Experimentation Results - Current Staff Allocation.

Location	Staff Type	FTEs	Requests	Interventions	Hours	Late Interventions	Late Hours	Requests Finished
CDN	Inf	2.0	90	626	1080	0	0	89
Cavendish	Inf	4.0	279	3312	5059	84	31	122
Metro	Inf	1.8	27	271	491	0	0	27
Parc Ex	Inf	1.4	36	248	359	0	0	36
		9.2						
Guichet	TS	7.6	2662	3520	8027	27	7	2659
CDN	TS	3.0	137	1358	3132	0	0	135
Cavendish	TS	5.0	353	2789	6364	130	46	230
Metro	TS	2.6	56	1189	2344	0	0	56
Parc Ex	TS	1.0	51	570	1119	2	14	47
		19.2						
CDN	Psy	4.4	428	3346	5592	89	31	364
Cavendish	Psy	2.4	78	1420	2550	0	0	71
Metro	Psy	1.0	40	683	1170	5	21	36
Parc Ex	Psy	2.0	21	310	625	0	0	22
		9.8						

Using this information, we can see that there are not enough social workers (TS) in the guichet (there are 27 late requests averaging 7 hours late), there are too many nurses (Inf) at Cote Des Neiges (there are no late requests and the number of hours each nurse is busy is very low compared to that in the guichet), there are not enough nurses at Cavendish, and there are too many nurses at both Metro and Parc Extension. Likewise, there are not enough social workers at Cavendish, too many social workers at Metro, a small need for more social workers at Parc Extension, not enough psychlogists (Psy) at Cote Des Neiges, too many psychologists at Cavendish, and too many psychologists at Parc Extension.

These results led to simulation experiments in which underutilized staff were reallocated to locations that didn't have enough of the same type of staff. The results are displayed in Table 6 where we see that with the same amount of staff we have improved the nursing situation at Cavendish, somewhat improved the situation for social workers at Cavendish, and addressed the psychologist situation at Cote Des Neiges. Should budget for additional staff become available, additional experiments can be performed to see how many additional staff are required to meet the needs of all areas at each location.

Table 6: Experimentation Results - Proposed Staff Reallocation.

Location	Staff Type	FTEs	Requests	Interventions	Hours	Late Interventions	Late Hours	Requests Finished
CDN	Inf	1.2	85	579	1001	0	0	83
Cavendish	Inf	7.0	279	5626	8642	2	4	205
Metro	Inf	0.6	28	250	450	0	11	27
Parc Ex	Inf	0.4	35	232	337	0	0	34
		9.2						
Guichet	TS	7.6	2670	3514	8001	36	22	2672
CDN	TS	2.6	134	1296	2979	1	4	130
Cavendish	TS	5.8	351	3182	7357	107	46	261
Metro	TS	2.0	58	1151	2279	0	4	49
Parc Ex	TS	1.2	54	636	1268	0	1	51
		19.2						
CDN	Psy	5.6	428	3840	6416	0	0	407
Cavendish	Psy	2.2	82	1404	2517	3	13	73
Metro	Psy	1.2	41	771	1321	2	13	38
Parc Ex	Psy	0.8	21	285	570	0	0	21
		9.8						

7 CONCLUSIONS AND OPPORTUNITIES FOR FUTURE WORK

This paper suggests an approach for rationalizing budgets when providing healthcare services within mandated delivery periods. This approach is to: collect appropriate data on demanded services; build, verify and validate a simulation model of the process; use the model to identify where mandated delivery times are not being met; reallocate staff to ensure that delivery times are just being met; and identify possibilities for reducing or pooling needed resources without negatively impacting services.

In addition to rationalizing budgeting, another benefit of this approach is the realization, achieved due to the discovery process needed for building, testing and using the model, that there are aspects of the process that need further examination to identify opportunities for providing more, better or timelier care with the same or less resources. These opportunities include: the possible use of tablets, with forms carefully designed to reduce the time needed for charting and reporting; the recording of time needed to support client activities; better categorization of client needs at the start of their journey through the system; categorization of the type of care provided each client; measurement of outcomes; the use of newly collected data to design mental health care pathways and to determine which staff members are most appropriate for specific client needs; dynamically deciding the interval of time between face to face interventions so as to level demand (and thus reduce variability and waiting times); dynamically making decisions as to whether clients remain in the guichet (for small number of interventions) based on the relative workloads of staff in the guichet and in the relevant first line areas and thus reduce the effects of variability of demand; merging the two smaller locations, at least from a service perspective, to reduce the effects of variability of demand.

A final benefit of this work is the potential, recognized by the chief financial officer, that our approach could be used throughout the CIUSSS to rationalize staffing levels while also improving service levels.

REFERENCES

Bedoya-Valencia, L., and E. Kirac. 2016. "Evaluating Alternative Resource Allocation In An Emergency Department Using Discrete Event Simulation". *Simulation* 92(12).

CIHI 2017. "Wait Times For Priority Procedures In Canada, 2017". Ottawa, Ontario: Canadian Institute For Health Information.

Fetter, R. B., and J. D. Thompson. 1965. "The Simulation Of Hospital Systems". *Operations Research*:689–711.

Fowler, J. W., and L. Monch. 2015. "A Comparison Of Discrete-Event Simulation Approaches For Complex Manufacturing Systems And Healthcare Systems". In *Simulation in Production and Logistics 2015*, edited by M. Rabe and U. Clausen. Stuttgart, Germany: Fraunhofer IRB Verlag.

Jacobson, S., S. Hall, and J. Swisher. 2013. "Discrete-Event Simulation Of Health Care Systems". In *Patient Flow: Reducing Delay In Healthcare Delivery*, edited by R. W. Hall, 211–252. Boston, Massachusetts.

Kim, B., Y. Elstein, B. Shiner, R. Konrad, A. S. Pomerantz, and B. V. Watts. 2013. "Use Of Discrete Event Simulation To Improve A Mental Health Clinic". *General Hospital Psychiatry, Psychiatry, Medicine and Primary Care* 35(6):668–670.

Meadows, G., H. G. H, B. Falconer, H. Kelly, L. Joubert, and M. Clarke. 1997. "The pattern-Of-Care Model: A Tool For Planning Community Mental Health Services". *Psychiatr Serv.* 48(2):218–223.

Ontario 2005. "First Ever Common Benchmarks Will Allow Canadians To Measure Progress In Reducing Wait Times". *https://news.ontario.ca/archive/en/2005/12/12/First-ever-common-benchmarks-will-allow-Canadians-to-measure-progress-in-reducin.html*.

Ramwadhdoebea, S., E. Buskensb, R. J. Sakkers, and J. E. Stahl. 2009. "A Tutorial On Discrete-Event Simulation For Health Policy Design And Decision Making: Optimizing Pediatric Ultrasound Screening For Hip Dysplasia As An Illustration". *Health Policy*:143–150.

Sargent, R. G. 2011. "Verification And Validation Of Simulation Models". In *Proceedings of the 2011 Winter Simulation Conference*, edited by S. Jain, R. Creasey, J. Himmelspach, K. White, and M. Fu, 183–198. Piscataway, New Jersey.

Schlesinger, S. R., R. E. Crosbie, R. E. Gagné, G. S. Innis, C. S. Lalwani, J. Loch, R. J. Sylvester, R. D. Wright, N. Kheir, and D. Bartos. 1979. "Terminology For Model Credibility". *Simulation* 32(3):103–104.

Sutherland, J., and N. Repin. 2014. "Current Hospital Funding In Canada Policy Brief". Vancouver, British Columbia: UBC Centre for Health Services and Policy Research.

Wagner, M. A. F., and J. R. Wilson. 1993. "Using Univariate Bezier Distributions To Model Simulation Input Processes". In *Proceedings of the 1993 Winter Simulation Conference*, edited by G. W. Evans, M. Mollaghasemi, E. C. Russell, and W. E. Biles, 365–373. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Wetherbe, J. C., and J. R. Montanari. 1981. "Zero Based Budgeting In The Planning Process". *Strategic Management Journal* 2(1):1–14.

AUTHOR BIOGRAPHIES

PHILIP TROY is a senior analytics advisor for Centre intégré universitaire de santé et de services sociaux du Centre-Ouest-de-l'Île-de-Montréal. Prior to that he focused his efforts on analyzing, simulating and making recommendations for improving processes at the Sir Mortimer B. Davis Jewish General Hospital. Efforts for both organizations include analyzing the need for surgical Intensive Care Unit beds; co-developing a simulation based optimization model for a pre-surgical screening clinic; supporting logistics team planning for operations in a new pavilion; supervising Time Driven Activity Based Costing analysis of cataract procedures and laundry operations; analyzing the cause of cancer surgery delays and an approach to preclude them; analyzing the need for radiology technicians in the O.R. Suite; and supporting outcome measurement efforts. Dr. Troy earned his doctorate in Operations Research/Management Science from Yale University. His e-mail address is Phil.Troy.CCOMTL@ssss.gouv.qc.ca.

LISA WESTAWAY is coordinator of the youth-adult-mental health continuum at Centre intégré universitaire de santé et de services sociaux du Centre-Ouest-de-l'Île-de-Montréal after previously being head of program for adult mental health for the same organization, and manager at the MAB-Mackay Rehabilitation Center. She has been working in health care for over 20 years, leading several successful restructuring projects aimed at improving access and quality health services while reducing costs. In addition to earning her M.Sc. in Vision Science at the University Of Montreal, Ms. Westaway completed the micro-proram in Process and Quality Improvement within the Health Care Sector at HEC Montreal. Her email address is Lisa.Westaway.CCOMTL@ssss.gouv.qc.ca.

ANTHONIE GRONDIN is a program manager for first line integrated services for Centre intégré universitaire de santé et de services sociaux du Centre-Ouest-de-l'Île-de-Montréal, after previously having been a coordinator for the intake office for adult mental health services for the same organization. A social worker by profession, he has been working in health and social services for more than 12 years. His email address is Anthonie.Grondin.CVD@ssss.gouv.qc.ca.

THADDEUS REZANOWICZ is director of the Mental Health and Addiction Program for Centre intégré universitaire de santé et de services sociaux du Centre-Ouest-de-l'Île-de-Montréal. Having worked in a number of Quebec healthcare institutions, Mr. Rezanowicz has over 25 years of management experience. In 1990, he created his first geographical interdisciplinary team in mental health, partnering with general practitioners in the area and the Chief of Psychiatry at Anna Laberge Hospital. From 2005 until 2010, he engaged institutional and community partners in a mental health action plan, which was implemented in close collaboration with the Centre hospitalier de l'Université de Montréal. Mr. Rezanowicz has a Master's Degree in Health Administration. His email address is Thaddeus.Rezanowicz.CVD@ssss.gouv.qc.ca