Proceedings of the 2018 Winter Simulation Conference
M. Rabe, A.A. Juan, N. Mustafee, A. Skoogh, S. Jain, and B. Johansson, eds.

MODELING THE IMPACT OF MAKE-AHEAD CHEMOTHERAPY DRUG POLICIES THROUGH DISCRETE-EVENT SIMULATION

Donald B. Richardson
Amy E.M. Cohn

University of Michigan
1205 Beal Ave
Ann Arbor, MI 48105, USA

ABSTRACT

During an outpatient chemotherapy infusion visit, patients typically have blood work done, see their oncologist in the clinic, wait for the pharmacy to prepare their drugs, and receive their infusion. The time variability at each of these steps can introduce delays, which not only negatively impact the patient but propagate through the system to negatively impact other patients and staff as well. One major opportunity to reduce patient waiting time is by *pre-mixing* (i.e., making drugs before the patient arrives for their infusion appointment) at the pharmacy. This, however, requires careful consideration of the trade-off between time savings versus the potential cost of wasting a drug if the patients are deemed ineligible for treatment on the day of their appointment. We present a discrete-event simulation model to predict the effectiveness of various make-ahead drug policies utilizing data from our collaborators at the University of Michigan Rogel Cancer Center (UMRCC).

1 INTRODUCTION

Healthcare operations, care delivery, and treatment decision making are just a few examples where data-driven modeling has greatly impacted performance and grown in popularity as data access has improved (Thara et al. 2016). Discrete-event simulation stands as one of the most tractable methods practiced to improve overall healthcare system performance (Katsaliaki and Mustafee 2011; Barjis 2011; Rohleder et al. 2011; Seminelli et al. 2016). This increasing trend of performance improvement models exists for cancer treatment centers as well (Woodall et al. 2013; Liang et al. 2015). With an aging population that continues to grow, the estimated number of new cases of cancer has grown from 1.5 million in 2010 to 1.7 million in 2018 (ACS 2018). While most patients undergo a combination of chemotherapy infusion (i.e., intravenous administered medication), radiation therapy, and surgery, over half of cancer patient in the US will require some form of chemotherapy. Additionally, cancer treatments continue to advance and increase in complexity. This results in an increase in the frequency of patient infusion visits for a given treatment regimen as well as an increase in cancer survivors who require follow-up visits (Levit and Patlak 2009). Consequently, this increased demand of patients at outpatient chemotherapy infusion centers leads to increased patient waiting time, patient dissatisfaction, and overworked staff. Additionally, patient and nurse safety becomes a concern as demand increases. The risk of the hazardous chemotherapy agents spilling on patients or nurses increases when the nurses are overworked (Schrijvers 2003; He et al. 2017; Harrison et al. 2016). Therefore, any process improvement on the system has the potential for multifaceted benefits.

Cancer patients who require chemotherapy infusion undergo exhausting and lengthy infusion sessions over the course of their treatment. These session lengths can increase even more during peak demand hours. Through observations of work flow and interviews with pharmacists at UMRCC, we determined peak drug demand hours on certain days of the week. During these peak hours, the pharmacy can get backed up to the point of taking up to two hours to get a drug out to a patient. Therefore, one major opportunity to reduce

patient delays is by optimizing drug preparation at the pharmacy. Drugs can be prepared (i.e., compounded) ahead of time to prevent patients from waiting as their chemotherapy drug(s) is/are compounded as done in Masselink et al. (2012). However, patients scheduled for outpatient chemotherapy infusion may be deferred for treatment (i.e., last minute cancellation due to not meeting medical parameters or personal reasons such as a family member not being available to support the patient) after arriving for their appointment (Fuentes and Frödin 2015). Consequently, the infusion center may incur a waste cost if a drug is made ahead and the patient defers treatment. Infusion centers must implement policies, determining which drugs to make before patients arrive given a fixed window of time, to balance this potential waste cost with the time savings for their patients and staff. In support of this effort, we developed a discrete-event simulation, using UMRCC data, that has widespread applicability to evaluate the effectiveness of pre-mixing chemotherapy drugs as well as the drug planning process in general. This simulation allows us to take multiple sets of drugs determined for pre-mix, along with the patient schedule for the day, and simulate the performance (i.e., time in system and staff utilization) of these sets (i.e., compare the pre-mix policies).

Through our computational experiments, we test proposed methods to create pre-mix policies. For example, if a drug is below a certain cost threshold and one or more patients are scheduled to receive it below a certain probability of deferral, then we can pre-mix the drug. We incorporate patient deferral probabilities from the prediction model in (Richardson et al. 2017). The simulation allows us to test various make-ahead policies for mixing drugs throughout the day. Improvements made at the pharmacy may reduce patient delays and nurse overtime in the infusion area.

The rest of the paper is structured as follows: Section 2 defines the problem we are solving, Section 3 describes the simulation modeling construction and all assumptions, Section 4 presents our computational experiments, and Section 5 provides discussion and conclusion.

2 PROBLEM DESCRIPTION

2.1 Description of Chemotherapy Infusion Process

Chemotherapy infusion patients' visits consist of getting blood work done, seeing their oncologist, waiting for the pharmacy to prepare their chemotherapy drug, and receiving their infusion. However, recurring patients do not always have a clinic visit as seen in Figure 1. We note that our focus will be on patients who do not have a clinic visit on the day of their appointment, however future work will consider two patient arrival streams, those with and those without a clinic appointment on the same day as their infusion. Patients

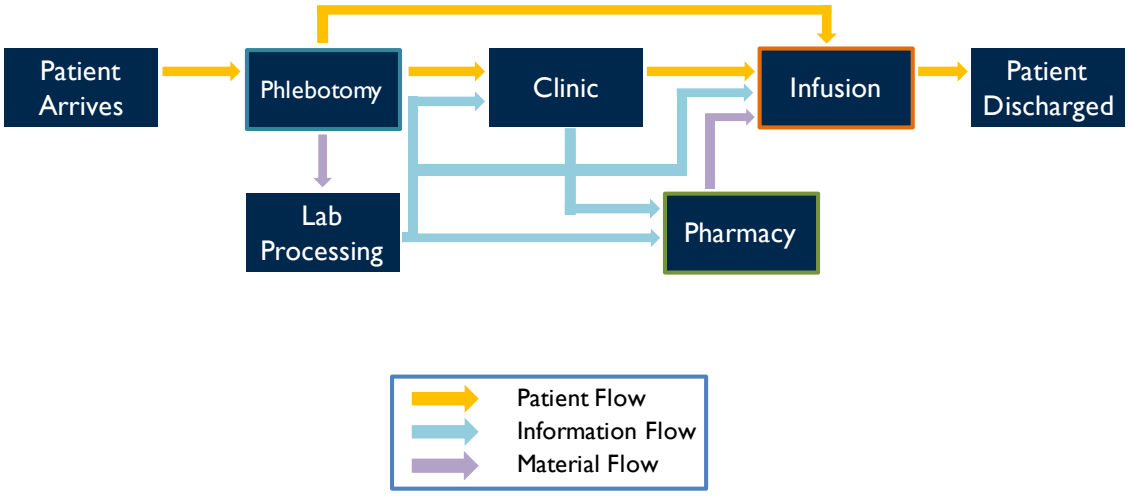

Figure 1: This figure presents the patient flow for outpatients and their information. This was created based off of observations and interviews conducted with the medical staff at UMRCC.

with a same-day clinic appointment provide our model additional information to better predict treatment deferrals as well as reduce the sources for lateness to the infusion center (e.g., bad weather, travel time, etc.)

At UMRCC, the patients' orders are only sent to the pharmacy once the patients have finished their clinic visit or are checked-in to the infusion area. The pharmacy has a goal to keep the order turnaround-time (TAT) under 1 hour for each patient. However, through observation and historical data, we see that the TAT can be as much as 2 hours. Our focus is to improve the drug turnaround time in the pharmacy and in turn reduce the overall time in the system, illustrated in Figure 1, for patients. We propose various policies ranging in risk tolerance. More conservative policies mimic the current state of pre-mix which are solely cost based. Conversely, less conservative policies consider pre-mixing higher-cost drugs if the probability of wasting them is low enough. Although all policies save time, the purpose of the simulation is to determine how much time and whose time (patient or pharmacist/tech).

2.2 Pharmacy Process Flow Description

The drug mixing process at the pharmacy consists of a series of order verifications, the actual compounding of the drug, and safety checks to ensure safe delivery to patients. The process we model begins when a drug order has been received by the pharmacy (i.e., the patient has arrived at infusion and is ready to be treated). Figure 2 illustrates the various steps taken and various checks needed to complete a chemotherapy drug order. If a drug is pre-mixed, all steps are performed in advance except for the final safety check. The drug mixing process is carried out by various pharmacist and technical staff. Pharmacists conduct both the order verification and safety check of all drugs. They are assigned to one of these two tasks for the first half of the day then switch. Pharmacy technical staff serve one of two roles: compounding drugs under the hood or printing all labels and collecting supplies.

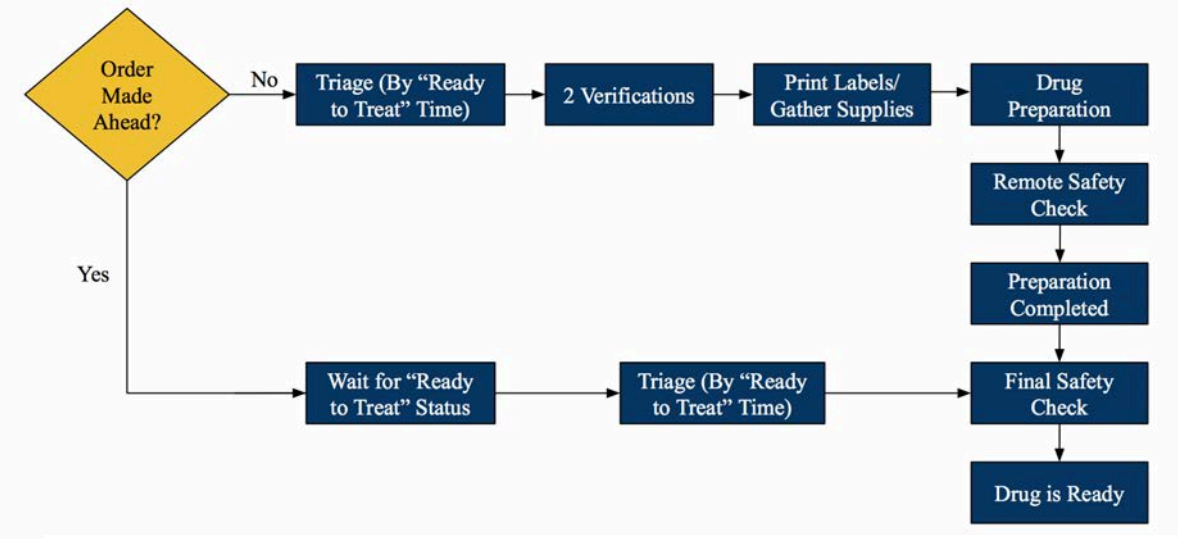

Figure 2: Pharmacy process flow for chemotherapy infusion drug orders.

For simulation purposes, we simplified this process flow into five main steps.

1. Check if the drug was made ahead (pre-mixed), if yes skip to Step 5 otherwise continue to Step 2.
2. Complete two drug verifications; must be done by two different pharmacists
3. Print labels and collect supplies for the drug order
4. Compound the drug and perform the first safety check. If order is mixed incorrectly, the same tech must remake the order (i.e., repeat Step 4)
5. Perform the safety check and deliver the completed order.

Arrivals are determined based on the patients' appointment time plus some random deviation where a negative deviation means the patient was early and positive means they were late. Each time a new order arrives, the orders are sorted first by arrival time then by appointment time (i.e., if two orders arrive at the same time before the pharmacy opens, the order with the earlier appointment time will be processed first). Orders are then released to follow the process in Figure 2. Once a drug is verified twice (Step 2) and all supplies have been gathered (Step 3), it then can proceed to the mixing hood. After the drug is mixed, the remote safety check is performed by the pharmacist. Before compounding the drugs, the mixing tech must take a photo of all materials and another photo after the drug has been fully compounded. The safety check pharmacist must review these images before the mixing process can be marked as complete. However, during this check, the mixing tech cannot conduct any other work. Since the historical data is only captured when a drug entered and exited the mixing hood, we included this initial safety check in the mixing time for our drugs. We note this safety check is performed in addition to the final safety check/drug delivery in step 5 of our simulation.

3 SIMULATION MODEL

Our simulation was built using the *SimPy* module in Python version 3.6. The model is initiated with a set of orders scheduled for the day, with some being flagged as pre-mixed, each of which has a scheduled arrival time for the corresponding patient. The actual arrival time of orders, the length of each service, and the potential for the compounding service to fail are all stochastic. Due to the interest of our collaborators, variable staffing was outside the scope of this project. Our metrics are the time in system for each order (i.e., perceived patient wait time) and the utilization of all resources (i.e., pharmacist and techs at the pharmacy).

3.1 Model Input Parameters

Our simulation input parameters are determined using a combination of observations, historical data, and expert pharmacist opinion. Using the *SciPy* module in Python, we fit all distributions used for arrival time and service time estimation as seen in Table 1.

Table 1: All input distribution and parameters used in our simulation model.

Process	**Distribution**	**Description**
Patient Arrivals	JohnsonSU (-0.428,1.41,-2.767,45.511)	Negative values=early arrival Positive=late arrival
First Verification	Triangular (1,2,15)	Expert Opinion in min
Second Verification	Triangular (1,2,5)	Expert Opinion in min
Print Labels/Kit	Triangular (1,3,5)	Expert Opinion in min
Drug Mix Time	Beta (1.461, 1376723443.471, 1.019, 7036129537.303)	Historical Data
Safety Check	Pearson3(2.509, 3.583, 3.240)	Historical Data

While the arrivals are appointment driven, most patients' actual arrival time will deviate from their scheduled appointment time due to the stochastic nature of any previous appointments as well as general tardiness or earliness. An additional delay can occur since a pharmacy order is only initiated once a patient has arrived and checked-in for their infusion appointment. We approach modeling the arrival process by first determining the distribution of the deviation from scheduled appointment time to actual arrival time (check-in time). Figure 3 presents a histogram of the arrival deviation data along with the JohnsonSU distribution used for arrival estimations in the simulation similar to (Rohleder et al. 2011).

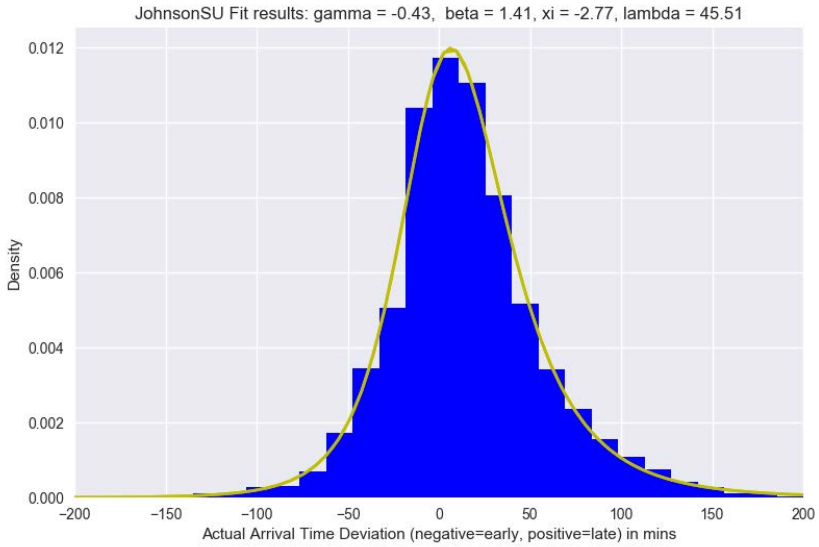

Figure 3: JohnsonSU Arrival Deviation Distribution Fit. Historical Data taken from UMRCC MiChart (Epic Product) Medical Records from 2016.

The model also incorporates patients deferrals/no-shows. From the pharmacy perspective, this will only affect the metrics if a patient's drug is pre-mixed and they have a same day deferral/no-show. To determine this probability, we trained a Bayesian Additive Regression Tree (BART) model with demographic (e.g., race, gender, age, sex, and marital status), scheduling history (i.e., number of previous cancellations, previous appointments, and past treatment protocols), and medical treatment patient data to make the treatment deferral/no-show predictions for each patient. This model was chosen based on a comparative procedure, discussed in Richardson et al. (2017), which tested multiple binary decision models to determine the model with the best out-of-sample prediction based on AUC, Briar Score, and F-1 Score. With a decision threshold of probability .75, we correctly predicted 93% of completed appointments and 21% of deferrals/no-shows with an overall accuracy of 84%. We note that while a false negative is not ideal, there is no waste cost associated with it (i.e., the drug would be mixed as planned.) However, a false positive could result in a drug being wasted if pre-mixed.

We emphasize that the UMRCC system was used as an example case study. The presented computational experiments utilized historical data from 2016 and expert opinion from the UMRCC pharmacy to estimate model inputs and validity. We note our modeling approach can be replicated for other facilities with minor modifications to the pharmacy process flow as well as the appropriate data sources.

After observing the simulated no pre-mix case discussed in Section 4, it was determined a reasonable estimate if the pharmacy did not pre-mix based on the observed data. We highlight that we focus on comparing various make-ahead policies as a proof of concept.

3.2 Modeling Conditions

The infusion center pharmacy opens at 6:00 am to complete all pre-mix drug orders. We assume that all pre-mixed orders will not expire before they are administered if the patient appointment is at noon or before (most drugs have an 8-12 hour lifetime). At 7:30am they finish pre-mixing and begin making orders as patients arrive (earliest appointment time is 7:30am). We assume if a drug is pre-mix that it will be completed in this window of time. This is when our simulation starts and runs until all patients are served each day. The simulation only considers a single arrival stream of patients based on their appointment time in infusion. There are some cases when a same day appointment can occur but the majority of appoints will be scheduled ahead of time. The simulation then models a single drug order for each patient with a drug

compounding probability of failure of 5%. Through observations and interviews with pharmacists, we found this to be a reasonable assumption. There is also a chance for pre-mixed orders to be wasted if a patient defers or doesn't show. We assume constant patient volumes for each iteration of the simulation buy vary them by day of week to best represent the cancer centers current schedule.

4 COMPUTATIONAL EXPERIMENTS

We present three tractable and easily implementable pre-mix policies. Then we compare their performance with the baseline scenario (i.e., no pre-mix policy). Table 2a and 2b highlight our performance metrics (total time in system and staff utilization) compared across scenarios. Since staffing is out of the scope of this paper, we keep a constant schedule across all scenarios based on the current staffing schedule at UMRCC. We simulate a week (i.e., Monday through Friday) with a constant amount of pre-mixed drugs each day since the pharmacy has the same limited time to pre-mix drugs each day. This week was simulated 20 times to insure we stayed within an error of 5 minutes.

4.1 Scenario 1

In the first scenario, we only consider pre-mixing drugs for the first 20 (this number is adjustable and dependent on the time allotted for pre-mix) patients who have a probability of deferral/no-show of 0.1 or lower. This threshold is used as an example; it is ultimately determined by the decision maker, depending on their risk tolerance. The probability of deferral/no-show for patients scheduled to receive the drug orders were determined using a BART model with the input data mentioned in Section 3.1.

4.2 Scenario 2

For our second scenario, we disregard the probability of deferral/no-show and focus solely on patient appointment times when deciding to pre-mix. Since the pharmacy schedule is heaviest before noon, we focus on pre-mixing only for patients with scheduled appointments from 8am-12pm . This also ensures that all pre-mixed drugs won't expire before the patients' appointments. Next, we determine the proportion of appointments in each hour block from the first appointment until noon. This proportion is used to determine the number of pre-mixed drugs in a respective hour block (e.g., suppose from 8-9am there are 20 patients out of a total of 100 that morning, then given we can only pre-mix 20 drugs, the first 5 drugs in that time window will be pre-mixed).

4.3 Scenario 3

In our last scenario, we combine ideas from both Scenarios 1 and 2. Using the proportions for each hour block determined in Scenario 2, we then only assign the allocated number of pre-mixed drugs to patients that fall below the probability of deferral/no-show threshold from Scenario 1. This policy should incorporate the benefits of pre-mixing throughout the first half of the day while being more risk adverse in regards to wasting drugs. This also may spread out the appointments that have pre-mix drugs (assuming the first five patient in our previous example do not all fall under the probability of deferral/no-show threshold).

4.4 Results/Discussion

Referring to Table 2a, it is clear that pre-mixing chemotherapy drugs has a significant impact on the average time in the system for a drug order as Scenarios 1, 2, and 3 all outperform the No Pre-mix case. On Day 3 we see that Scenario 1 out performs Scenario 2. This is an example where being more risk-seeking by not considering the probability of deferral/no-show had less reward (i.e., drugs may have been pre-mixed for patients who actually deferred treatment resulting in no benefit for pre-mixing) for the patients. However, all other days we see Scenario 2 dominates 1. This results from Scenario 1 only mixing the first 20 drugs that fit the criteria. This reduces the potential propagated time savings compared to spreading the pre-mixed drugs throughout the schedule. While not always significant, we see that the average time in system for

Scenario 3 and the average number of wasted drugs is lower than Scenarios 1 and 2. We hypothesized this to be the case since Scenario 3 is a more efficient "rule of thumb" policy by utilizing the more conservative approach of Scenario 1 but also by lightning the pharmacy load throughout the morning instead of just in the first couple of hours similar to Scenario 2.

Table 2a: Contains the comparison of time in system for the various scenarios in minutes.

Drug Order Time in System and Wasted Drug Results

Metrics		Scenarios			
Day of Week		No Pre-mix	1	2	3
1	Average	52.79	30.70	26.64	26.17
	C.I.	(49.51, 56.07)	(29.53, 31.87)	(26.23, 27.05)	(25.65, 26.68)
2	Average	85.63	46.60	41.73	38.19
	C.I.	(80.79, 90.46)	(44.43, 48.8)	(39.21, 44.25)	(36.19, 40.19)
3	Average	58.04	35.44	37.69	27.47
	C.I.	(54.74, 61.34)	(33.65, 37.22)	(34.87, 40.51)	(26.59, 28.35)
4	Average	38.10	24.78	22.82	22.43
	C.I.	(36.3, 39.89)	(24.18, 25.37)	(22.35, 23.33)	(22.06, 22.81)
5	Average	47.86	28.32	25.73	25.70
	C.I.	(44.95, 50.78)	(27.53, 29.09)	(25.27, 26.19)	(24.71, 26.7)
Avg. # of drugs wasted per day		0	2.81	3.13	2.32

Looking at Day 2 (highest drug order demand day of the week) in Table 2b, we also notice a significant effect on staff utilization. For example, both Verification Pharmacists and Printing Technicians have a very high utilization in the No Pre-mix case. By simply pre-mixing 20 drugs before the patient rush, we are able to decrease their utilization by almost 10% and 5% respectively. This supports the idea that pre-mixing chemotherapy drugs will benefit both the patients and the staff to better ensure safe delivery of such high hazardous drugs. Since we assume the same distribution for mixing all chemotherapy drugs and that arrivals are appointment driven, the utilization across all pre-mix policies is relatively the same. We note that our utilization calculation for the Safety Check Pharmacist are an extreme under-estimate. As discussed in Section 3, we were only able to capture the start time for compounding a drug to the completed time. As seen in Figure 2, this also incorporated the Remote Safety check which is performed by the Safety Pharmacist. It is here where almost all of the 5% of failed drug orders are determined and cycled back for re-compounding. Refer to the Appendix for utilization results on all other days.

Table 2b: Utilization comparison between policies on Day 2 of our simulation as percentages.

Day 2 Utilization

Resource	Scenarios											
	No Pre-mix			1			2			3		
	Avg.	CI-	CI+	Avg.	CI-	CI+	Avg.	CI-	CI+	Avg.	CI-	CI+
Verification Pharm	0.89	0.80	0.99	0.82	0.77	0.87	0.81	0.77	0.85	0.81	0.77	0.86
Print Tech	0.71	0.56	0.85	0.65	0.62	0.67	0.65	0.61	0.69	0.65	0.62	0.68
Mix Tech	0.30	0.28	0.32	0.28	0.27	0.29	0.26	0.25	0.27	0.27	0.26	0.28
Safety Pharm	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01

5 CONCLUSION

Using discrete event simulation, we evaluated various pre-mix policies to determine which most benefited both the individual patients and the entire outpatient infusion system. We used estimated probabilities of deferral/no-show from our BART model to develop these policies for chemotherapy drugs at an outpatient

chemotherapy infusion center pharmacy. These experiments were not meant to determine the best approach for UMRCC but to demonstrate the idea of pre-mix as a technique that can be utilized by other institutions.

The discussed model serves multiple purposes both for our current and future work. While we can test current "rule of thumb" polices as done in this paper, we also can evaluate optimization models that determine, within a fix window of time, what set of drugs are optimal to pre-mix to minimize the expected time in system as well as expected waste cost. Our immediate next steps will incorporate the drug cost into the simulation model as well as introduce additional patient arrival streams from those with a clinic visit. Future extensions will also allow for more dynamic policies to be tested that provide unattainable improvements from a static model.

This work provides an invaluable tool to both engineers and medical professional working to reduce patient waiting time in an outpatient chemotherapy infusion center by helping insure the safety of the patients and improves their overall satisfaction.

ACKNOWLEDGMENTS

We would like to thank the University of Michigan-Rackham Graduate School, the Center for Healthcare Engineering and Patient Safety (CHEPS), and the Seth Bonder Foundation for funding this research. We would also like to thank our clinical collaborators, Vincent Larroca and Carolina Typaldos. Also we thank Francisco Aldarondo, Matt See, Hwon Tak, Alisson Winnik and all other CHEPS students who assisted with this work.

APPENDIX - SIMULATION LOGIC

This simulation consists of ten queues:

- The order arrival queue (A_q) which contains one entry per patient pharmacy order and is sorted by actual arrival time of the order into the pharmacy queue
- The order verification queue (V_q) is initially empty but later contains all orders waiting for verification. These orders are sorted by scheduled appointment time
- The Available-Pharmacists-to-Verify Queue (VP_q) contains one entry for each pharmacist available for the verification task, sorted by time that they become available (i.e., when their shift starts or when they have completed a task)
- The print labels queue (L_q) follows a First in First Out (FIFO) policy. This queue contains an entry for each order after it finishes two verifications
- The Available-Tech-to-Print Queue (PT_q) contains one entry for the single print tech working during that shift
- The mixing drug queue (M_q) follows a FIFO policy. This queue contains an entry for each order after finishing printing the labels
- The Available-Tech-to-Mix-Drug Queue (MT_q) contains one entry for each mixing tech available for this task, sorted by the time they become available
- The safety check and sorting queue (S_q) also follows a FIFO policy. This queue contains an entry for each order after drug mixing is complete
- The Available-Pharmacist-to-Safety-Check- Queue (PS_q) contains one entry for contains one entry for each safety check pharmacist tech available for this task, sorted by the time they become available
- The Event queue is used to sequentially order and trace all events in the simulation

These queues are initialized as follows:

- For each drug order, we generate the actual arrival time from a perturbation of the scheduled arrival time (i.e., some patients arrive before and some after the scheduled appointment time). Each order is then placed into $A_{q,}$
- All order with an actual arrival time before the start-of-operations are placed into V_q by order of their appointment time
- VP_q, PT_q, $MT_{q,}$ and PS_q are then generated based on the staff schedule for all pharmacy techs and pharmacist.
- If a drug is pre-mixed, skip to step 6 below

Loop until through the entire arrival queue:

1. Any drug order with an actual arrival time at or before the simulation clock time will then move into the verification queue.
    - This queue is then sorted by scheduled appointment time
2. Verification pharmacist then take the order from the top of the queue to complete service. After the first verification is complete, the order is then placed back in the verification queue with a higher priority than the orders who have not received a first verification yet.
3. Next the order must go through a second verification done by a different pharmacist than their first verification. After the second verification, the drug order is then sent to the print labels/drug kit queue.
4. Print techs than take orders in their queue at a first come first serve basis. After service is complete, the order along with all supplies are placed in the drug mixing queue.
5. Mixing tech again grab the orders by FIFO and start working on the drug. There is a chance the tech makes a mistake which is caught by one of the safety checks. If this is the case, the tech must re mix the entire drug. After mixing service is successfully completed. The drug is placed into the sort/safety check queue
6. Another group of pharmacist will pull from the sort/safety check queue by FIFO. Once this service is complete the drug is ready for the patient.

Utilization Tables

Table a: Utilization comparison between policies on Day 1 of our simulation as percentages.

Utilization Day 1

Resource	Scenarios											
	No Pre-mix			1			2			3		
	Avg.	CI-	CI+	Avg.	CI-	CI+	Avg.	CI-	CI+	Avg.	CI-	CI+
Verification Pharm	0.65	0.61	0.69	0.54	0.51	0.58	0.54	0.52	0.57	0.54	0.52	0.55
Print Tech	0.52	0.49	0.54	0.43	0.41	0.45	0.43	0.41	0.45	0.43	0.40	0.46
Mix Tech	0.21	0.20	0.21	0.17	0.16	0.18	0.17	0.16	0.17	0.16	0.16	0.17
Safety Pharm	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01

Table b: Utilization comparison between policies on Day 3 of our simulation as percentages.

Utilization Day 3

Resource	Scenarios											
	No Pre-mix			1			2			3		
	Avg.	CI-	CI+	Avg.	CI-	CI+	Avg.	CI-	CI+	Avg.	CI-	CI+
Verification Pharm	0.87	0.83	0.90	0.76	0.72	0.80	0.75	0.70	0.80	0.75	0.72	0.79
Print Tech	0.69	0.66	0.72	0.60	0.57	0.63	0.60	0.57	0.63	0.60	0.58	0.63
Mix Tech	0.28	0.27	0.29	0.25	0.23	0.26	0.23	0.23	0.24	0.24	0.23	0.25
Safety Pharm	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01

Table c: Utilization comparison between policies on Day 4 of our simulation as percentages.

Utilization Day 4

Resource	Scenarios											
	No Pre-mix			1			2			3		
	Avg.	CI-	CI+	Avg.	CI-	CI+	Avg.	CI-	CI+	Avg.	CI-	CI+
Verification Pharm	0.56	0.51	0.61	0.46	0.42	0.51	0.46	0.42	0.51	0.46	0.41	0.51
Print Tech	0.41	0.34	0.47	0.34	0.28	0.39	0.33	0.28	0.39	0.33	0.28	0.38
Mix Tech	0.19	0.18	0.20	0.16	0.15	0.16	0.15	0.14	0.16	0.15	0.14	0.15
Safety Pharm	0.01	0.00	0.01	0.01	0.00	0.01	0.01	0.00	0.01	0.01	0.00	0.01

Table d: Utilization comparison between policies on Day 5 of our simulation as percentages.

Utilization Day 5

Resource	Scenarios											
	No Pre-mix			1			2			3		
	Avg.	CI-	CI+	Avg.	CI-	CI+	Avg.	CI-	CI+	Avg.	CI-	CI+
Verification Pharm	0.69	0.64	0.73	0.58	0.55	0.62	0.58	0.55	0.62	0.58	0.54	0.62
Print Tech	0.55	0.52	0.58	0.47	0.44	0.50	0.46	0.44	0.49	0.47	0.45	0.49
Mix Tech	0.22	0.22	0.23	0.20	0.19	0.21	0.19	0.18	0.20	0.19	0.19	0.20
Safety Pharm	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01

REFERENCES

American Cancer Society (ACS). 2018. American Cancer Society - Cancer Facts & Figures 2018. https://www.cancer.org/research/cancer-facts-statistics/all-cancer-facts-figures/cancer-facts-figures-2018.html, accessed 20.01.2018.

Barjis, J. 2011. “Healthcare Simulation and its Potential Areas and Future Trends”. *SCS M&S Magazine* 2(5):1-6.

Fuentes, S. and J. E. Frödin. 2015. “Why is Intravenous Chemotherapy Cancelled and How Often. Could it be Prevented. A Prospective Analysis of all Planned and Given Intravenous Anti-tumor Treatments at the Department of Oncology, Karolinska University Hospital, Stockholm during One Month”. *Acta Oncologica* 54(7):1056–1062.

Harrison, J. M., B. Lavasseur, P. J. Stella, P. T. Adams, L. Swafford, J. Lewis, K. Mendelsohn-Victor, and C. R. Friese. 2016. “Factors Associated with Toxicity-related Service Use Among Community Oncology Patients”. *Journal of Clinical Oncology* 34(7):133-133.

He, B. Y., K. Mendelsohn-Victor, M. C. McCullagh MC, and C. R. Friese. 2017. "Personal Protective Equipment Use and Hazardous Drug Spills among Ambulatory Oncology Nurses: Results from a Mailed Survey". *Oncology Nursing Forum* 44(1):60-65.

Katsaliaki, K. and N. Mustafee. 2011. "Applications of Simulation within the Healthcare Context". *Journal of the Operational Research Society* 62(8):1431-1451.

Levit, L. and M. Patlak. 2009. "Ensuring Quality Cancer Care through the Oncology Workforce: Sustaining Care in the 21st Century: Workshop Summary". Institute of Medicine (US) National Cancer Policy Forum. Source Washington DC, USA: National Academies Press.

Liang, B., A. Turkcan, M. E. Ceyhan, and K. Stuart. 2015. "Improvement of Chemotherapy Patient Flow and Scheduling in an Outpatient Oncology Clinic." *International Journal of Production Research 53*(24):7177-7190.

Masselink, I. H., T. L. van der Mijden, N. Litvak, and P. T. Vanberkel. 2012. "Preparation of Chemotherapy Drugs: Planning Policy for Reduced Waiting Times." *Omega* 40(2):181-187.

Richardson, D. B., S. D. Guikema, and A. E. Cohn. 2017. "Predicting Patient Treatment Deferrals at an Outpatient Chemotherapy Infusion Center: A Statistical Approach". *JCO Clinical Cancer Informatics* 1:1-8.

Rohleder, T. R., P. Lewkonia, D. P. Bischak, P. Duffy, and R. Hendijani. 2011. "Using Simulation Modeling to Improve Patient Flow at an Outpatient Orthopedic Clinic". *Health Care Management Science* 14(2):135-145.

Schrijvers, D. L. 2003. "Extravasation: A Dreaded Complication of Chemotherapy". *Annals of Oncology* 14(supplement 3):26–30.

Seminelli, M. D., J. W. Wilson, and B. M. McConnell. 2016. "Implementing Discrete Event Simulation to Improve Optometry Clinic Operations." In *Proceedings of the 2016 Winter Simulation Conference*, edited by T. M. K. Roeder et al., 2157-2168. Piscataway, New Jersey: IEEE.

Thara, D. K., B. G. Premasudha, V. R. Ram, and R. Suma. 2016. "Impact of Big Data in Healthcare: A Survey". In *Proceedings of the 2nd International Conference on Contemporary Computing and Informatics (IC3I*), 729-735. Piscataway, New Jersey: IEEE.

Woodall, J. C., T. Gosselin, A. Boswell, M. Murr, and B. T. Denton. 2013. "Improving Patient Access to Chemotherapy Treatment at Duke Cancer Institute". *Interfaces*, *43*(5):449-461.

AUTHOR BIOGRAPHIES

DONALD B. RICHARDSON is a PhD candidate working under the advisement Dr. Amy Cohn in the Department of Industrial and Operations Engineering at the University of Michigan. He is also an affiliate member of the Center for Healthcare Engineering and Patient Safety. His research interests involve applying optimization, predictive modeling, and simulation in healthcare. He is interested in working in the healthcare industry continuing to help improve the delivery and efficiency of patient care. He is also particularly interested in helping increase the participation of minorities in STEM. Donald's email is donalric@umich.edu.

AMY E.M. COHN is an Alfred F. Thurnau Professor in the Department of Industrial and Operations Engineering at the University of Michigan, where she also holds an appointment in the Department of Health Management and Policy in the School of Public Health. Dr. Cohn is the Associate Director of the Center for Healthcare Engineering and Patient Safety (CHEPS) and serves on the leadership team of the Institute for Healthcare Policy and Innovation (IHPI). She holds an A.B. in applied mathematics, magna cum laude, from Harvard University and a PhD in operations research from the Massachusetts Institute of Technology. Her primary research interests are in applications of combinatorial optimization, particularly to healthcare and aviation, and to the challenges of optimization problems with multiple objective criteria. Her email address is amycohn@umich.edu.