To: Amy Cohn, Dan Moyer, Directors of Michigan Medicine Primary Care Clinic

From: Kaden Todjo, Varsha Rajesh, Vasiliki Tsirukis, Justin Gates, University of Michigan Engineering Students, Founders of KVVJ Solutions

Subject: Primary Care Clinic Access Problem at Michigan Medicine

Date: April 21, 2025

FOREWORD

As a research and consulting firm, KVVJ Solutions, we've developed this final report to convey the deterministic event simulation and analytical evaluation needed to properly address the patient access issues within our client, Michigan Medicine's clinic. Operations research and system analysis are interdisciplinary fields committed to improving complex systems through evidence-based decision-making, efficiency optimization, and human-centered design. This paper looks at our firm's underlying principles, particularly in human systems integration and computer-based modeling, to outline our implementation of discrete event simulation to address issues within patient access care in primary care.

SUMMARY

This report looks at computational modeling, specifically in its role in improving healthcare systems. While the computational work KVVJ does is quite essential to our proposed solution, almost as important is the understanding of how a primary care clinic operates, and how this understanding guides our decision-making when attempting to find a solution. As a primary care clinic, patients with a wide range of concerns enter, presenting the clinic with a high relative standard deviation of demand, thus it is critical to schedule patients within the goal metric, even in addition to the development of a triage system, because some patients may need to see a physician before others, based off the severity of their emergency. KVVJ strives to increase healthcare accessibility and efficiency by tackling issues such as access to healthcare by reducing how many slots are needed to address scheduling issues, strategically developing ways around low budget constraints, having enough replications, meeting target metrics like getting 95 percent of patients scheduled in the same period (three weeks), and creating a long enough time full to produce the serviceable results. Our overall goal is to develop a solution for patient access to health care. This solution depends on the average number of patients arriving each week and on the variance of our data. The final report will provide a comprehensive analysis of KVVJ's impact on our client's primary care clinic, supported by our work in computational modeling and our communication of our results.

WHAT IS KVVJ SOLUTIONS?

Description

KVVJ Solutions is a cutting-edge computational modeling firm focused on revolutionizing the healthcare system, run by founders Kaden Todjo, Varsha Rajesh, Valia Tsirukis, and Justin Gates. At KVVJ, we are committed to optimizing patient access and care while reducing the cost of running primary care clinics. KVVJ uses innovative, evidence-based modeling software to ensure that our patients are given the optimal results that lead to the best possible care, access, and lives saved. Our core values are focused on innovation, collaboration, and integrity. We will continue to pioneer advancements to set new standards in the healthcare industry, build solutions with clinics that are both practical and sustainable, and continue to focus on using evidence-based decision-making.

Mission Statement

At KVVJ, our mission is to revolutionize patient care by designing and implementing computer-based modeling and human systems integration research for our clients so that they can best pursue their patients' care. Employing state-of-the-art technologies, we strive to ensure timely, effective care for every patient while maintaining sustainability and cost-effectiveness. Together, we will develop system outlook innovation that generates a patient-first mindset while advancing the future of healthcare in the field of system operations.

Our Focus on Primary Care Clinics

Our firm focuses on streamlining primary care because it is the foundation of an effective healthcare system. By optimizing patient visit times, we can improve early detection, reduce specialty visits, and enable better treatment, affecting other sectors of healthcare, leading to a more efficient healthcare system. Strengthening primary care through discrete event simulation can lighten the burden on emergency care and specialty sectors, creating a positive ripple effect on the entire system.

OUR CLINIC

KVVJ has been selected to work with a primary care clinic. These clinics provide general medical care across a diverse population of possible patients, ranging from children to the elderly. Primary care clinics cater to a mix of semi-urgent and non-urgent needs. For example, some patients may need relatively immediate care for their condition, while others may just be coming in for a physical or follow-up visit. Keeping this in mind, our access goal for our clinic is to ensure that 95 percent of our patients are scheduled within three weeks. This goal strikes a balance between urgent and non-urgent needs, where three weeks is the maximum acceptable wait time for more urgent needs. This patient-centered approach aligns with KVVJ’s values and beliefs. These patient access goals are unique to primary care clinics and couldn’t be applied to other clinics, such as newborn health or the emergency department, where it would be much more likely that immediate care would be needed. Each one of these clinics has unique problems, constraints, and needs, so it is important to carefully analyze the constraints of our clinic to ensure that its specific needs are met. By addressing and implementing the proper solutions to these challenges, our clinic ensures that timely and affordable care is accessible to our community.

PARAMETERS

Modeling patient demand with a Gaussian distribution with a mean of 100 patients per week and a standard deviation of 20 is appropriate, given the widespread services of primary care delivery. The variety in types of appointments primary care clinics receive has differing required meeting schedules, since some are annual, while others, like treatment for chronic illness, could be weekly. There is also variance between weeks since there could be spikes during flu season or slower weeks during the holidays. Gaussian modeling is well-suited for primary care clinics due to their operational consistency and manageable patient volume. Additionally, while there is some variance, most weekly demands are centered around the mean.

Primary care clinics typically offer routine services like preventive care, chronic disease management, and general health maintenance, which lead to a steady stream of patient visits. The chosen mean of 100 patients per week represents a realistic patient volume for a primary care provider, as the average provider sees 93 patients per week (Washington Post, 2014).

The standard deviation of 20 allows for the natural variability in weekly demand, as seasonal illness trends, holidays, or community events can largely affect the number of people seeking treatment. Because people seek primary care for a large variety of reasons—from acute, sudden conditions to ongoing chronic care—the demand can fluctuate significantly week to week. This justifies the use of a higher standard deviation to capture realistic variability in patient flow while still maintaining the underlying consistency that makes Gaussian modeling appropriate for primary care clinics.

ANALYSIS AND COST

Initial Capacity Testing

While finding the optimal daily average supply, our approach was to calculate one standard deviation above and below the mean, giving a range of [80, 120]. This allowed for a minimal range to test which supply value was the smallest that still satisfied both our target metrics.

We began by simulating five replications, each over four weeks, with 120 slots per week. This exceeded the Michigan Medicine target metrics, with 99 percent of patients being seen within one week. We then lowered it to 110 patients, which had only 92 percent of people seen within the first week. Testing 112, 113, 114, and 115, we found that the lowest number of slots that satisfied the Michigan Medicine target metric was 114 patients. However, the error for the mean queue was high, and to lower that, we had to change the number of periods and replications to reduce error while keeping computational costs low.

Reducing Error: Replication Length and Count

We first focused on the number of periods (weeks) per replication, experimenting with increased replication lengths to observe their effect on variability. Extending from 4 to 8 and then to 12 weeks significantly reduced the standard error of the mean queue length, while maintaining minimal outliers. However, increasing beyond this led to diminishing returns and increased computational cost. Next, we optimized the number of replications. Starting from five (which yielded high error), we tested fifteen (low error but high cost), and then intermediate values. Ultimately, we found that ten replications provided acceptable error at a feasible computational cost.

Figure 1. Michigan Medicine Metric: Graphs of Queue, Individual Replications (left) and Mean with Error (right)

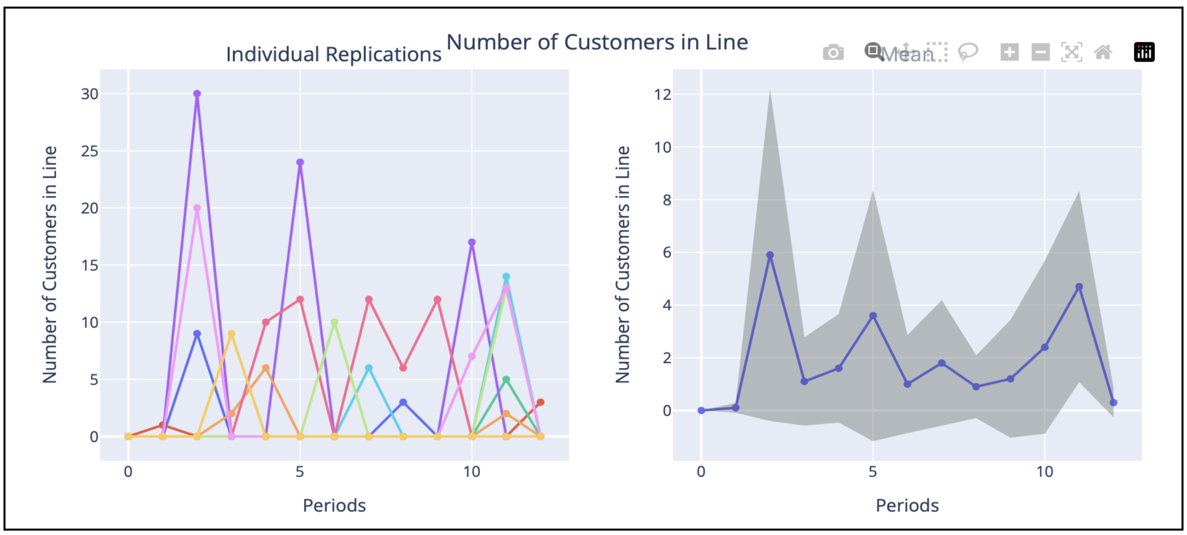

After minimizing the size of the error as much as possible, we found that the trends between individual replications were quite variable.

Figure 2. Michigan Medicine Metric: Graph of Number of Weeks Patients Wait

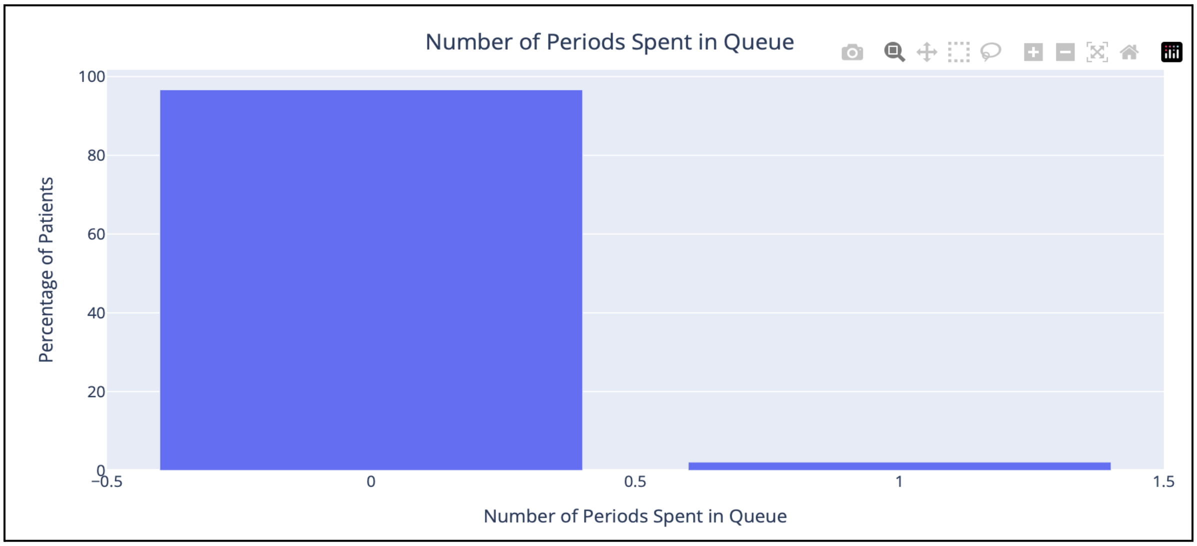

In the simulation pictured (Figure 2), 97.3 percent of patients were seen within one week. In various other simulations, it was always over 95 percent, meeting the Michigan Medicine target metric.

Comparing Clinic-Specific vs. Michigan Medicine Metrics

The clinic-specific target is more appropriate to follow than the Michigan Medicine metric because it better reflects the resource constraints and operational goals of a smaller clinic. While Michigan Medicine's target of seeing nearly all patients within one week is appropriate for a large center, a smaller clinic benefits from a more flexible and realistic goal.

Clinic-Specific Target: 95% of patients scheduled within three weeks
Michigan Medicine Target: 95% of patients scheduled within one week

The three-week window supports sustainable operations and personalized care without overcommitting resources. Moreover, when examining the graph for individual replications for client-specific vs Michigan Medicine metrics, the client-specific replications follow a clear upward trend while Michigan Medicine does not, meaning the client-specific solution is more reliable in future replications and testing of the parameters.

Determining Minimum Slots for Client-Specific Target

We began with the Michigan Medicine solution (114 slots, Table 1) and observed that 100% of patients were seen within three weeks. Since the client's target is 95% within three weeks, we reduced the number of slots to test lower limits. We continued reducing the number of weekly slots incrementally, simulating each configuration to test whether it still met the client's target. Ultimately, we found that 85 slots per week was the minimum supply required to ensure at least 95 percent of patients were seen within three weeks across all replications. This offered a more resource-efficient solution while still aligning with the client's performance goals. By combining simulation data with a systematic reduction in supply, we ensured our recommendation balanced both access to care and operational efficiency. Using a timeframe of 12 weeks worked well for this as well, although we tried other values such as 4, 10, and 20 weeks. We found that the same number of replications, being ten, was successful in keeping errors and costs as low as possible.

After simulation-based testing with reduced weekly slots, we determined that 85 slots per week (Table 1) was the minimum required to meet the target while balancing performance and efficiency.

Table 1. Summary of Total Cost and Performance For Client-Specific and Michigan Medicine Parameters

Metric	Client-Specific Metric	Michigan Medicine Metric
Slots per week	85	114
Operational cost	85 × \$200 = \$17,000	114 × \$200 = \$22,800
Late patients/week	~1	~1
Late fees	1 × \$100 = \$100	1 × \$100 = \$100
Total weekly cost	\$17,100	\$22,900

Using the provided prices for operational and late costs, this calculates the cost per week for both metrics' optimized solutions. The cost savings between these configurations is \$5,800 per week, and performance still meets client goals.

Final Optimization: Balancing Cost and Access

To further minimize the total weekly cost of the client-specific model, we considered the trade-off between operational cost and late patient fees. Since the cost of one late patient (\$100) is significantly less than the cost of one additional slot (\$200), it is more economical to allow a slight increase in late patients rather than oversupply.

By reducing slots further and testing each new parameter, we found that 81 slots per week (Table 2) is the minimum number needed to still meet the client's 95% access metric across all replications.

Table 2. Final Parameters and Cost Post-Cost Optimization

Metric	Value
Slots per week	81
Operational cost	81 × \$200 = \$16,200
Late patients/week	~4
Late fees	4 × \$100 = \$400
Total weekly cost	\$16,600

This results in an additional \$500 cost reduction per week compared to the 85-slot configuration.

Figure 3. Client-Specific Metric: Graphs of Queue, Individual Replications (left) and Mean with Error (right)

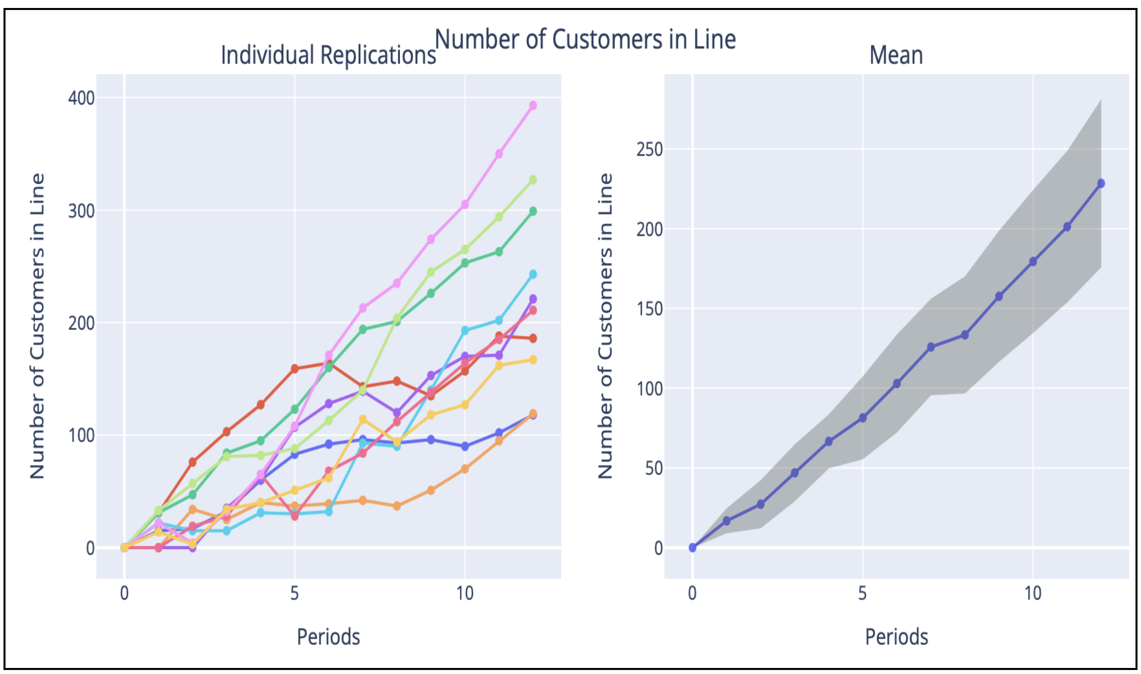

This graph, after reducing error as much as possible, has a much clearer and more consistent trend between individual replications.

Figure 4. Client-Specific Metric: Graph of Number of Weeks Patients Wait

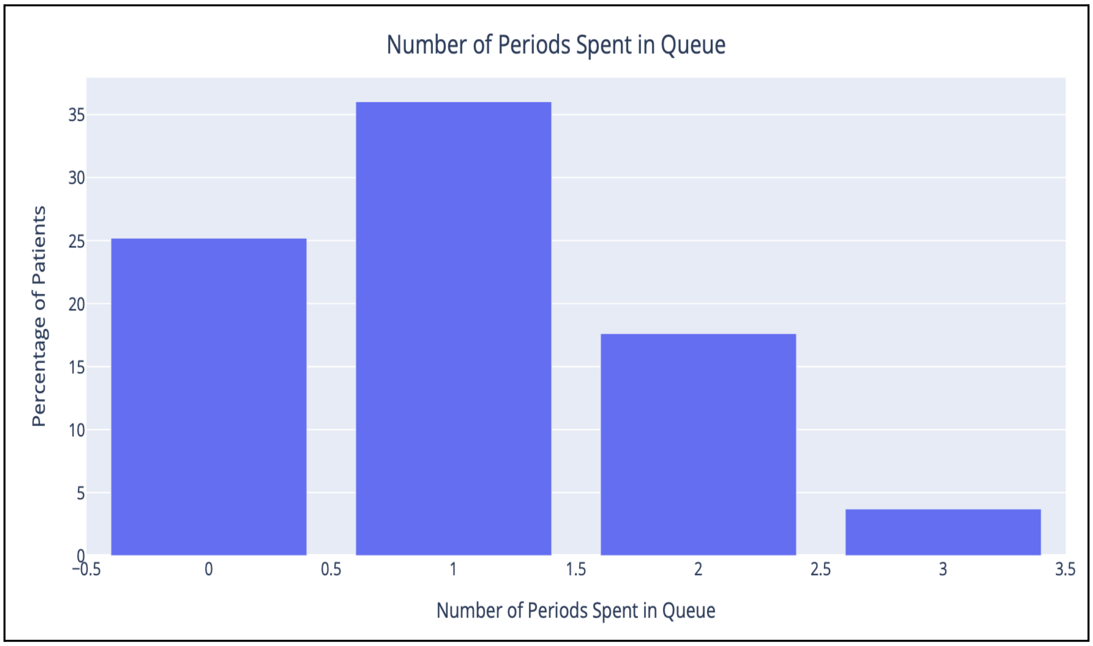

Since late costs are less than the cost per slot, making the percentage of patients seen after three weeks as close to five percent as possible (but still below) saves costs. Additionally, this meets the client-specific target metric with approximately 96.1 percent of patients being seen in three weeks(Figure 4).

FINANCIAL SUMMARY

In the 120 weeks that the simulation is run, the total runtime cost for the optimized solution is $1,992,000 (Table 3). The average clinic earns $250–$300 per patient (Dojo Business). For every 100 patients, which is the mean number of patients seen in the clinic weekly, our client's clinic is estimated to earn between $3,000,000 and $3,600,000 throughout the total runtime of the simulation (Table 4). The 81-slot configuration allows for only one main provider, as the average physician sees about 93 patients per week (Washington Post, 2014). This contributes to lower staffing costs, and late fees remain under 3% of total cost, indicating strong access performance. Well below the total revenue, incurring this cost for approximately two years is a great benefit to the clinic in the long term, since our discrete event simulation and optimizing the slots will decrease costs and increase access substantially. The total profit of the clinic (assuming the operational costs and revenue include materials, staffing, etc) would range between $1,008,000 and $1,608,000 during the simulation runtime.

Table 3. Summary of Weekly, Annual, and Total Costs Per Client Metric

Configuration	Total Runtime	Weekly Slots	Weekly Cost	Annual Cost	Total Runtime Cost
Michigan Medicine Metric	120 weeks	114	$22,900	$1,190,800	$2,748,000
Original Client-Specific Metric	120 weeks	85	$17,100	$889,200	$2,052,000
Optimized Client-Specific Metric	120 weeks	81	$16,600	$863,200	$1,992,000

This shows the weekly, annual, and total runtime cost (120 weeks) for each solution, with the optimized client-specific metric having the lowest cost.

Table 4. Estimated Minimum and Maximum Revenue Per Patient, Weekly Patients, Annual Patients, and Total Patients During Runtime

Type of Estimate	Revenue per Patient	Weekly Revenue	Annual Revenue	Total Runtime Revenue (120 weeks)
Minimum	$250	$25,000	$1,300,000	$3,000,000
Maximum	$300	$30,000	$1,560,000	$3,600,000

This visualizes the revenue ranges of a primary care clinic of our size, used for comparison.

POTENTIAL CHANGES IN VARIANCE

Decreasing Variance

Primary care clinics attract many different patients with a great variety of ailments – from annual check-ups and health screenings to treating illnesses, wounds, and managing chronic illnesses. The diversity in length and types of appointments leads to this clinic having a high variance in patients per week. Because the standard deviations are high, a large number of slots must be included. Artificially lowering the variance would slightly increase the mean number of patients attending the clinic per week, but it would drastically lower the standard deviation of patients, so fewer slots must be accounted for, and thus, less will have to be paid. The primary care clinic would be willing to pay to decrease variance – it would save a large portion of the budget per month.

Increasing Variance

Since this clinic already has a high variance, it would not be beneficial to increase it in order to receive discounts on the slots. Since increasing variance means more slots have to be included, as explained earlier, it would not be worth it for the clinic unless the discount given is extremely high. More slots at a lesser cost would be around the same cost as fewer slots at a greater price, so increasing the variance and uncertainty would be stressful on the business.

CONCLUSION

Primary care is the cornerstone of healthcare access, both for individual and community well-being. Primary care is the crucial first point of the healthcare system, and the variety of services primary care provides for its patients is the key to improving health outcomes, promoting health equity, and potentially lowering overall healthcare costs. KVVJ's ability to develop promising results for widespread patient access to healthcare using computational modeling and systems operations is thus the determining factor in ensuring the highest possible impact on the healthcare system. Our results have modeled patient access difficulties, addressed by evaluating the number of slots required to maintain adequate performance in a Michigan Medicine primary care clinic while satisfying specific clinic goals. In the development of statistical modeling, to address our problem, KVVJ minimized the total weekly cost of operation by reducing slots from 85 to 81 slots per week, while still configuring our data to retain 95% of patients within metric between all replications and including a comprehensively integrated review of differing variance within our statistical analysis. In our minimization, KVVJ developed an 81-slot configuration for the main provider, producing a total weekly revenue of $25,000-30,000. In the 120 weeks that the simulation is run, the contribution of lowering staffing costs and late fees remains under 3% of total cost, indicating strong access performance even though the estimated total revenue is between $3,000,000 and $3,600,000 (Table 4). Moreover, KVVJ, the use of substantial funds within 120 weeks for our overall cost balances out due to our optimization of the number of slots, thus decreasing the overall cost in the final analysis. The result of KVVJ's consultation not only ensures timely, effective care for every patient while maintaining sustainability and cost-effectiveness, but also designs an outline for our client to efficiently and effectively care for a variety of patients.