Migration Letters

Volume: 21, No: S13 (2024), pp. 1901-1924

ISSN: 1741-8984 (Print) ISSN: 1741-8992 (Online)

www.migrationletters.com

Predictive Analytics For AI-Assisted Patient No-Show Management And Clinic Revenue Optimization: A Simulation-Based Research

Ahmad Jamal[1], Fatima Tauseef[2], Zeeshan Akbar[3]

1 Abstract

Patient non-attendance at planned appointments for outpatient care is a constant operational and financial challenge facing healthcare systems. Missed appointments lead to less capacity utilisation, less efficient workflow, and direct revenue loss. While predictive models have increasingly been used in no-show forecasting, many studies focus more on the statistical performance of models without integrating the predictions into operational decision-making and financial evaluation. This Research develops a simulation-based predictive analytics framework relating machine learning, threshold optimisation and revenue modelling in an outpatient clinic context. A synthetic data set of 10 000 records of appointments was created to model behavioural, demographic and operational determinants of attendance. Logistic Regression and Random Forest models were evaluated based on ROC-AUC, Precision, Recall and Cross validation. Although predictive discrimination was moderate (AUC approx. 0.58), threshold sensitivity analysis showed that lowering the classification threshold to 0.30 allowed for an increase in recall to 60%, which would allow for effective identification of high-risk patients in a cost-sensitive environment. A financial simulation was performed to compare baseline scheduling against an AI assisted risk-based intervention strategy. Under normal circumstances, no-shows incurred significant revenue loss. The combined use of predictive risk scoring and targeted reminder interventions decreased revenue loss, and produced positive net financial improvement, even when conservative assumptions of intervention effects were used. Sensitivity analysis confirmed the robustness of financial gain in multi effectiveness scenarios. The results show that predictive analytics can help to add efficiency to outpatient clinics as long as the findings are linked to economic decision logic. Rather than prioritizing statistical accuracy, value arises along the lines of combining model outputs with operational intervention strategies. This Research presents contribution of a decision-oriented framework for AI assisted no-show management and emphasizes the need for cost sensitive threshold calibration in healthcare analytics.

Keywords: *Non-Attendance, Predictive, Simulation, Threshold, Revenue, Intervention*

[1]MBA- Business Analytics, King Graduate School, Monroe University 434 Main St, New Rochelle, NY, 10801, ajamal7503@monroeu.edu

[2]MBA-IT, Department of Business Administration and Management, Washington University of Science and Technology 2900 Eisenhower Ave, Alexandria, VA 22314, ftauseef.student@wust.edu

[3]Raymond A Mason School of Business, William and Mary University, ukropway 101, Williamsburg, Virginia, zakbar@wm.edu

2 Introduction

Healthcare systems still struggle with inefficiencies of operation affecting both service delivery and financial viability (Brikci et al., 2024). Among these challenges, non-attendance of patients at scheduled outpatient appointments is a persistent and expensive problem. Missed appointments disrupt clinical workflows, affects capacity utilisation, waiting times and generates direct revenue losses. For outpatient clinics that are constrained in staffing and/or infrastructure, even moderate no-show rates can seriously impact their operational efficiency and financial performance.

Some traditional mitigation strategies have mostly been based on uniform reminder systems (e.g., SMS notifications, emails, and phone calls). While these interventions may have some impact on non-attendance, it is often done in a non-discriminatory way across patient populations without taking into account the variation in risk (Opon et al., 2020). This blanket approach is likely to lead to inefficient use of resources, where low-risk patients are unnecessarily reminded, whereas high-risk people are not adequately targeted. With the growth of electronic health records and data access, predictive analytics provides a way to move away from reactive scheduling practices toward proactive, risk-based scheduling management.

Machine learning models have been applied to no-show prediction to a greater and greater degree, and techniques have extended from logistic regression to ensemble methods such as Random Forest (Fan et al., 2021). However, available research often assesses predictive performance using mostly statistical measures of accuracy or AUC without fully integrating predictions into operational frameworks for decision-making. Furthermore, classification thresholds are often chosen with default settings instead of cost-sensitive reasons, although the financial impacts of missed appointment and reminder interventions are asymmetric (McLean et al., 2016).

This Research fills these gaps by developing a simulation-based predictive analytics framework that combines the methodologies of machine learning, threshold optimisation and financial modelling in the context of an outpatient clinic. Rather than only concentrating on the issue of predictive accuracy, the Research considers how interventions aimed at reducing the risk of losses can reduce revenue losses, if they are consistent with economic decision logic.

To guide the investigation, the following research questions are proposed:

RQ1: To what extent are machine learning models able to forecast outpatient no-show behaviour using behavioural and operational variables?

RQ2: What effect does classification threshold optimisation have on the operational effectiveness of predictive models in a cost-sensitive healthcare setting?

RQ3: What is the financial impact of the integration of AI-based prediction with targeted intervention strategies in the outpatient clinic operations?

The rest of this dissertation discusses the literature on no-show prediction and operational analytics, the simulation methodology, predictive and financial results, and managerial implications and limitations.

3 Literature review

3.1 Patient No-Show Behaviour in Outpatient Settings

Patient non-attendance has been widely recognised as a persistent challenge to operation in outpatient healthcare systems (Silva et al., 2024). Missed appointments create idle clinical capacity, waiting times, interfere with resource allocation and depress revenue stability. Previous studies have identified a variety of factors that determine no-show behaviour, such as demographic factors, socioeconomic status, distance travelled and lead time to the appointment (Chen, 2023). However, behavioural history is again consistently found as one of the strongest predictors. Studies have shown that those patients with a history of missing appointments have a much higher chance of missing their next appointment, indicating behavioural reinforcement effects.

Lead time between scheduling and appointment is also commonly linked with no-show probability. Longer waiting times have been associated with longer waiting times, forgotten appointments and lower psychological commitment (Xie & Or, 2017). Accessibility barriers like transportation difficulty only add to non-attendance risk (Duerstock et al., 2023). These results provide evidence favouring the consideration of behavioural and structural variables in predictive modelling schemes.

However, much of the prior literature has been devoted to find correlates of non-attendance rather than incorporating these findings into decision systems in operation. Consequently, there still exists a gap between behavioural prediction research and strategies of implementation in management.

3.2 Machine Learning for No-Show Prediction

With the development of digital health records, machine learning approaches have gained wider application in the field of no-show prediction. Logistic regression has traditionally been the overwhelming modelling method because of its interpretability and simplicity. While effective in identifying the linear relations, logistic models have often difficulties in identifying the complex interaction effects of the behavioural and operational variables.

More recent studies have been focused on ensemble algorithms like Random Forest or gradient boosting algorithms. These methods are more advanced in terms of predictive performance, as they can model non-linear relationships and interactions between variables, without making strong parametric assumptions (Conzuelo Rodriguez et al., 2021). Comparisons of linear and ensemble models often find that tree-based methods have better discrimination, although results differ from one data set to another.

Despite such methodological improvements, predictive discrimination in no-show studies is often moderate as opposed to high (Deina et al., 2024). Reported values for the AUC are commonly between 0.60 and 0.75 because of the natural uncertainty of human behaviour. This would indicate that predictive systems should be tested not only against statistical correctness but also against the practical utility of decisions.

3.3 Cost-Sensitive Classification and Threshold Optimization

A major limitation of many studies of no-show prediction is the use of default classification thresholds (Liu et al., 2022). In the case of classification (either binary or not), the default threshold of 0.50 may not reflect operational cost structures. In the case of healthcare settings, the cost of a high-risk no-show for which we fail to identify (false negative) is usually higher than the cost of unnecessary intervention (false positive).

Cost sensitive classification theory focuses on the importance of calibrating decision boundaries based on asymmetric loss functions. Precision- Recall trade-offs need to be considered in the economic context of the decision environment. However, in many predictive studies the accuracy and the AUC are reported without linking the threshold selection to financial results.

The present Research addresses this gap by taking into account threshold sensitivity analysis and matching threshold choice with revenue exposure and cost asymmetry of intervention. This integration is more than simply predictive modelling and is driven towards decision-oriented analytics.

3.4 Appointment Scheduling and Operational Analytics

Operations research literature has a long history of Researching the issues of scheduling optimisation and overbooking in outpatient clinics (Ala & Chen, 2022). Traditional overbooking models try to account for expected no-shows by scheduling extra patients in the same slot of time. While effective under some assumptions, overbooking has the risks of overcrowding and service delays when attendance exceeds expectations.

More recent methods call for dynamic scheduling systems based on predictive risk scoring. By pre-empting high risk patients, clinics can selectively overbook or assign extra reminders without having to make uniform policy changes across all clinic appointments. This risk-based approach is in line with current operational analytics frameworks, which focus on data-driven resource allocation.

However, only a relatively few studies explicitly quantify the economics of the integration of predictive models with scheduling decisions (Wang et al., 2024). Revenue simulation and return-on-investment analysis are poorly developed fields of no-show prediction research.

3.5 Revenue Optimization in Healthcare Contexts

In the case of outpatient care, revenue optimisation is a matter of optimising the utilisation of capacity and minimising idle time (Humphreys et al., 2022). Missed appointments directly result in reduced billable activity and inefficiency in staffs and infrastructure utilisation. Whilst reminder systems have been shown to reduce no-show rates, blanket reminder strategies can be resource-intensive and inefficient.

Risk-based intervention allocation is a more focused solution. By targeting outreach efforts to patients whose predicted risk is higher, clinics can target intervention resources more effectively and cause less revenue leakage (Gouveia et al., 2023). However, empirical evidence of net financial gain from predictive targeting is scarce in the literature.

The present Research is a contribution to this domain, since it integrates predictive modelling, threshold calibration and financial simulation in a unified framework. And instead of measuring the performance of models in isolation, it measures revenue impact under different scenarios of intervention effectiveness.

3.6 Identified Research Gap

The literature shows large advances in terms of finding determinants or no-show behaviour and the applications of machine learning techniques for predictive tasks (Issah et al., 2023). However, three gaps are still obvious. First, predictive models are frequently assessed in terms of statistical measures without operations. Second, threshold selection is not often based on cost-sensitive analysis. Third, financial implications of predictive deployment are rarely measured using simulation.

This Research makes a contribution to these gaps by incorporating behavioural modelling, threshold optimisation and revenue simulation in a unified analytical framework. By connecting predictive outputs directly to intervention and economic evaluation in the operating room, it helps advance a decision-oriented approach to AI-assisted clinic management.

4 Methodology

4.1 Research Design

This Research uses a quantitative, simulation-based research design to determine the effectiveness of the use of predictive analytics for patient no-show management and clinic revenue optimization (Zhou et al., 2023). Due to the lack of access to datasets from real-world hospital data and the ethical issues that surround patient data privacy, a dataset was created to model realistic patient behaviour in outpatient clinic appointments (Chiruvella & Guddati, 2021). Simulation-based modelling has been extensively used in healthcare operations research applications where controlled experiment is needed to test decision frameworks within a structured set of assumptions.

The research is structured around a pipeline of analysing the data:

(1) data generation,

(2) exploratory data analysis,

(3) predictive model development,

(4) threshold optimization, and

(5) financial impact simulation.
The goal is not only to evaluate predictive accuracy but also combine risk prediction with operational intervention and economic evaluation.

4.2 Dataset Gathering
A data set with 10,000 appointment records was collected in an attempt to reproduce normal outpatient clinic activity. Each record was for an individual scheduled appointment. The dataset contained 31 predictive features after encoding representing demographic, behavioral, and operational characteristics.
The variables were grouped into three categories:

4.2.1 Patient Characteristics:
Age, gender, distance from clinic (km), chronic condition status, socioeconomic band.

4.2.2 Appointment Characteristics:
Appointment type, department, booking channel, lead time (days between booking and appointment), day of week, time slot.

4.2.3 Behavioural History:
Previous no-shows, previous attended appointments, days since last attendance, reminder type.
The binary target variable, **no_show**, indicated whether the patient failed to attend the appointment.
A baseline no_show probability of 12% was defined, and adjusted by rule-based increments relevant for realistic behavioral patterns. For example, longer lead time, longer travel distance and prior no-show history increased no show probability while reminder interventions decreased no show probability. The resulting data set generated an overall no-show rate of 28.56%, which is consistent with upper range outpatient estimates documented in healthcare operations literature.
The probability column used to generate the outcome variable was removed before model training to prevent data leakage.

4.3 Data Preparation
Data preprocessing was done in Python in Google Colab. Categorical variables were converted to a machine learning compatible format (one-hot encoding). The final feature matrix had 10000 observations and 31 independent variables.
The data set was split into a train and test dataset in an 80/20 ratio. Stratified sampling was used to maintain the 28.56% no-show rate in the two subsets to maintain class distribution consistency. The training data set had 8000 observations and the testing data set had 2000 observations.

4.4 Model Development
Two classification models were developed and compared:
1. Logistic Regression (baseline linear model)
2. Random Forest Classifier (ensemble non-linear model)

Logistic Regression was chosen as a benchmark because it is popularly used in healthcare prediction studies and is interpretable (Levy & O'Malley, 2020). However, the linear models may have difficulty with complex interaction effects between behavioral and operational variables.
Random Forest was chosen to model non-linear relationships as well as feature interactions without much parameter tuning (Ao et al., 2019). The model was trained with 200 decision trees of default depth parameters and fixed random seed to make it reproducible.

4.5 Model Evaluation Metrics
Model performance was assessed using multiple evaluation metrics to provide a comprehensive assessment:

- Accuracy
- Precision
- Recall
- F1-score
- Receiver Operating Characteristic (ROC) curve
- Area Under the Curve (AUC)
- Precision–Recall curve

Because there was a moderate class imbalance in the data (i.e., about 29% no-shows) recall and precision were given priority over accuracy. In the case of healthcare operational applications, missing those high-risk patients (false negatives) could lead to significant revenue losses whereas false positives generally have low intervention costs.

Five-fold cross-validation was performed on the entire data set and the scoring measure was ROC-AUC. This gave an estimate of the model stability and generalisability in the simulated environment.

4.6 Threshold Optimization

Classification models have a default to a probability of 0.50. However, this cut-off is not ideal in cost-sensitive healthcare scenarios.

To account for this, a threshold sensitivity analysis was done with thresholds set at 0.20 to 0.50. For each threshold, the values of precision and recall were computed. A threshold of 0.30 was chosen on the basis of operational trade-off analysis, where approximately a 60% recall is achieved with an acceptable precision (approximately 34%).

This adjustment is based on cost sensitive decision-making logic, in which finding high risk no-shows is more valuable than achieving high classification accuracy.

4.7 Financial Simulation Model

To evaluate the operational impact of predictive analytics, a financial simulation framework was developed.

The following assumptions were applied:

- Average appointment revenue: $120
- Reminder intervention cost (phone call): $5 per high-risk patient
- Intervention effectiveness: 40% reduction in predicted no-shows

Two scenarios were compared:

Baseline Scenario (No AI):

Revenue loss calculated as the number of no-shows multiplied by appointment revenue.

AI-Assisted Scenario:

High-risk patients (predicted probability ≥ 0.30) received targeted intervention. Among actual no-shows within this group, 40% were assumed to convert to attendance.

Net financial improvement was computed as:

$$\text{Baseline revenue loss} - \text{AI} - \text{assisted revenue loss} - \text{intervention cost}$$

Sensitivity analysis was conducted by varying intervention effectiveness from 20% to 60% to evaluate robustness of financial outcomes.

4.8 Ethical and Methodological Considerations

As the data collected for the research was used, patient privacy risks were avoided. However, ethical considerations are still relevant in the real world. Predictive systems can cause bias if demographic variables have a greater influence on classification results (Dominguez-Catena et al., 2024). Operational decisions such as overbooking have to balance efficiency with fairness for the patient.

Furthermore, the simulation is based on assumed intervention effectiveness and simplified cost structures. Therefore, results should be taken as indicative rather than definitive.

5 Results

5.1 Exploratory Data Analysis

5.1.1 Overall No-Show Distribution

The dataset contains 10,000 outpatient appointments records out of which 28.56% were non attended by the patient. This means that close to one in three scheduled appointments result in underutilization of resources into the simulated clinic environment. Such a prevalence level presents a meaningful predictive challenge and also a substantial operational burden. From a financial standpoint, this baseline non-attendance rate suggests substantial revenue exposure and inefficiency in the utilization of clinic capacity. The class distribution also makes sense of moderate imbalance which has implications to the model evaluation metrics and threshold choice in the later predictive aspects.

Figure 1: Overall No-Show Rate

```
Overall No-Show Rate: 28.56 %
         count
no_show
   0      7144
   1      2856

dtype: int64
```

Figure 2: Attended vs No-Show distribution

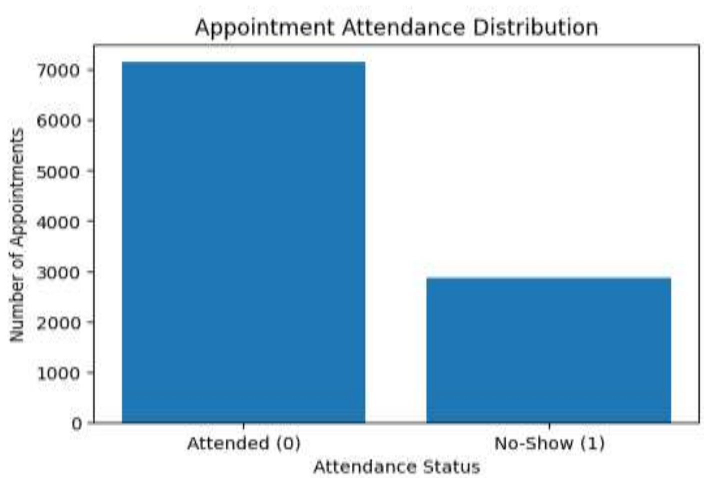

5.1.2 Behavioural Persistence: Effect of Previous No-Shows

Analysis of previous attendance history shows a clear escalation in no show behaviour. Patients with no history of missing appointments have a much lower likelihood of not attending the appointment than those with one or more prior no-shows. Increasing levels of probability are associated progressively with the more the number of historical no-shows, with the highest levels of risk being shown by individuals who are associated with repeated non-attendance. This pattern supports the concept of behavioural persistence; whereby past attendance behaviour is a good predictor of future compliance. The finding emphasises the importance of including historical attendance variables in predictive modelling frameworks.

Figure 3: No-Show by Previous No-Shows

prev_no_shows	no_show
0	0.183468
1	0.244908
2	0.280119
3	0.361658
4	0.357977

dtype: float64

Figure 4: Escalation trend of no-show rate by previous no-shows

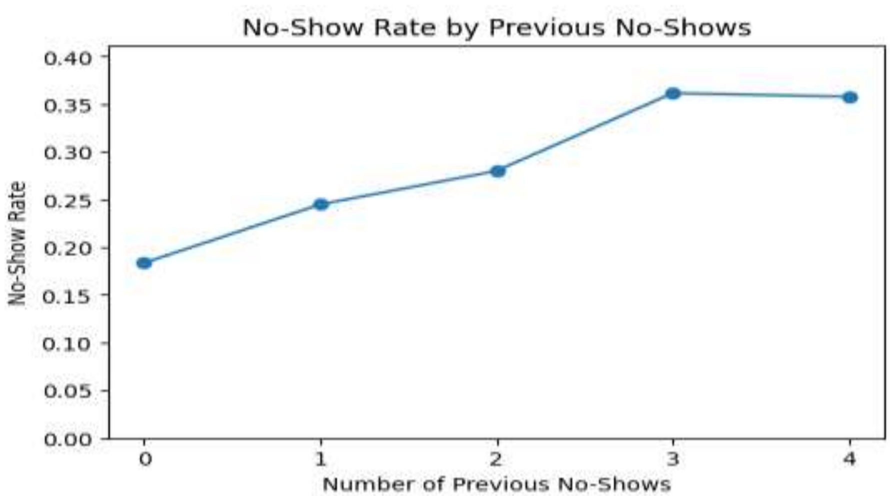

5.1.3 Temporal Commitment: Lead Time Analysis

The relationship between scheduling horizon and adherence to the appointment shows a noticeable trend. Appointments scheduled with shorter lead times have lower no-show rates than appointments scheduled further in advance. As the lead time becomes longer, the likelihood of not attending increases more and more. This pattern implies that long time between booking and appointment might decrease the commitment of patients or make them less sure-footed. From an operational standpoint, the need for longer lead times may necessitate improved strategies to remind patients, or overbooking policies to reduce attendance risk. The results support the significance of time-based variables in predictive healthcare analytics.

Figure 5: No-Show Rate by Lead Time

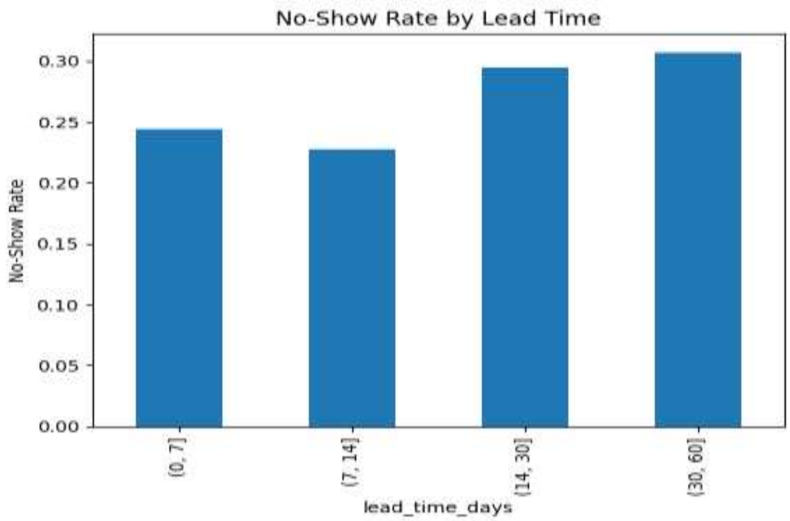

5.1.4 Reminder Type and Attendance Behaviour

Reminder interventions show a moderate variation in effectiveness between communication channels. Appointments preceded by direct phone call reminders have relatively lower levels of non-attendance than SMS or email reminders. While differences are not extreme, the results suggest that more personalized mechanisms of reminders may enhance patient accountability. However, reminder strategy alone does not eliminate the risk of non-attendance, suggesting that intervention based on predicted risk may be more efficient than blanket strategies based on reminders.

Figure 6: Impact of Reminder Type

	no_show
reminder_type	
Call	0.256817
SMS	0.272127
Email	0.299764

dtype: float64

Figure 7: No-show rates across reminder types

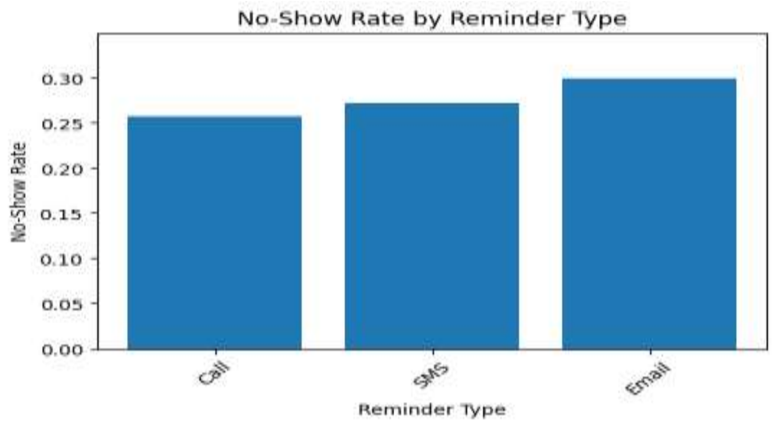

5.1.5 Accessibility Factors: Distance from Clinic

Accessibility becomes a significant influence in attendance behaviour. Patients who live farther away from the clinic have higher probabilities of non-attendance compared to those who live closer by. This finding suggests that barriers to transportation or inconvenience of traveling may contribute to missed appointments. In operational terms, distance-based risk profiling could be used to identify and target intervention strategies such as telehealth substitution, or reminders prioritisation. The significance of access is consistent with other work in healthcare service delivery that focuses on the structural determinants of attendance. As shown below, non-attendance rates rise more as the distance becomes further from the clinic, indicating that access constraints play an important role in adherence to appointments.

Figure 8: No-show rate by distance category

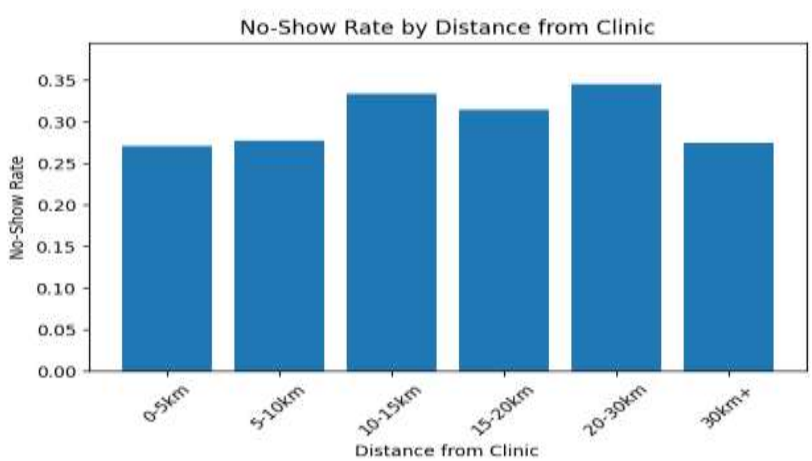

5.1.6 Correlation Structure of Numerical Variables

The correlation analysis of numerical variables gives an insight into multivariate relations in the data set. Behavioural variables, like past no-shows, recency of the last attendance, and lead time have a greater relationship with the no-show outcome than static demographic characteristics have. Although the magnitude of correlations is moderate in magnitude, their joint action justifies the use of non-linear ensemble models that have the ability to capture interaction effects. The correlation structure also emphasises the limitations associated with relying on linear models only for tasks of complex behavioural prediction.

Figure 9: Heatmap of correlation matrix

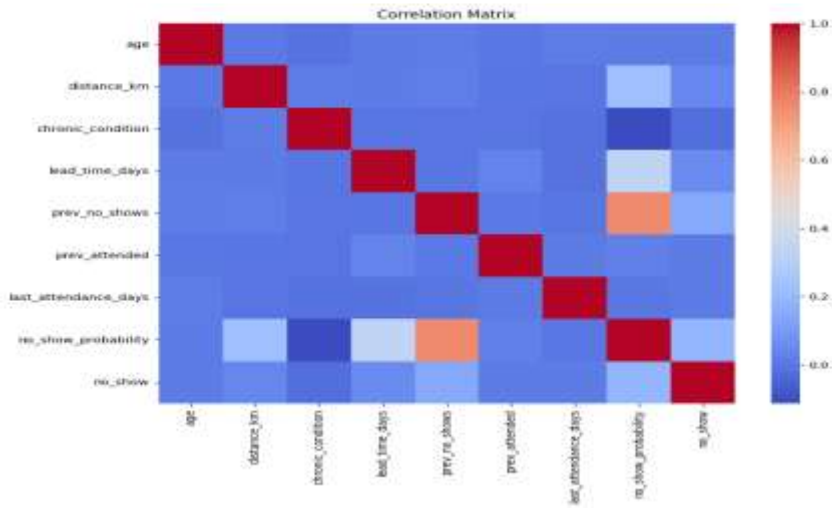

5.2 Model Comparison

Two classification models were evaluated: Logistic Regression and Random Forest.

5.2.1 Model 1: Logistic Regression

Logistic Regression gave an ROC-AUC of around 0.593, which is to say that the discrimination capability is not very good. However, using the 0.50 default probability threshold the model provided extremely low recall for no-shows (approx. 1%). This means the model predicted that almost all patients would attend and this leads to a high false negative rate. Although overall accuracy seemed to be ok, this measure was misleading with the class imbalance.

Figure 10: Logistic Regression Results

Logistic Regression Results

	precision	recall	f1-score	support
0	0.72	0.99	0.83	1429
1	0.38	0.01	0.03	571
accuracy			0.71	2000
macro avg	0.55	0.50	0.43	2000
weighted avg	0.62	0.71	0.60	2000

ROC-AUC Score: 0.5931768630531681

5.2.2 Model 2: Random Forest

A non-linear ensemble technique, Random Forest, resulted in a similar ROC-AUC of about 0.579 on the test set. While discrimination was still moderate, the ensemble model gave more flexibility in the capturing of interaction effects between behavioral and operational variables.

Figure 11: Random Forest Results

```
Random Forest Results
              precision    recall  f1-score   support

           0       0.72      0.99      0.83      1429
           1       0.44      0.01      0.03       571

    accuracy                           0.71      2000
   macro avg       0.58      0.50      0.43      2000
weighted avg       0.64      0.71      0.60      2000

ROC-AUC Score: 0.5785333576809619
```

Five-fold cross-validation yielded a mean ROC-AUC of around 0.574 with a standard deviation of 0.017 indicating stable model performance across folds and low overfitting in the simulated environment.

Figure 12: 5-Fold Cross-Validation

```
Cross-Validation AUC Scores: [0.60550969 0.56105393 0.55887183 0.56807327 0.5756478 ]
Mean CV AUC: 0.5738313023568005
Std Dev: 0.0168939143982461
```

These results indicate that predictive discrimination is moderate, rather than strong. However, predictive strength alone is not the determining factor of operational value because threshold of classification is immensely important in cost-sensitive applications.

5.3 Threshold Optimization

Using the 0.50 threshold, which is the default, both models had extremely low sensitivity to no-shows. Given the operative goal of identifying high-risk patients for intervention, threshold optimization was conducted.

Threshold sensitivity analysis was carried out for values of 0.20 to 0.50. At a threshold of 0.20, recall was around 91.9%, but the precision was lower, around 29.6%, which means an excessive number of false positives, and therefore interventions that are not necessary. At a threshold of 0.35, while there was an increase in precision by almost 34.7% there was a huge decrease in recall to about 34.3%.

Figure 13: Random Forest (Threshold = 0.30) Results

```
Random Forest (Threshold = 0.30)
              precision    recall  f1-score   support

           0       0.77      0.54      0.63      1429
           1       0.34      0.60      0.44       571

    accuracy                           0.55      2000
   macro avg       0.56      0.57      0.53      2000
weighted avg       0.65      0.55      0.58      2000
```

A threshold of 0.30 provided the most balanced trade-off, achieving:

- Precision ≈ 34.1%
- Recall ≈ 60.0%

This threshold allowed to identify 60% of actual no-shows and with a manageable intervention volume. The corresponding precision-recall curve verified this trade off pattern, showing the expected inverse relationship between sensitivity and specificity in imbalanced classification situations.

These results highlight the need for threshold adjustment in healthcare decision systems. Default thresholds that are optimized for statistical precision may not be consistent with operational goals.

5.4 Feature Importance Analysis

Random Forest feature importance analysis identified the most influential predictors of no-show behaviour.

Figure 14: Feature Importance Analysis (Graph)

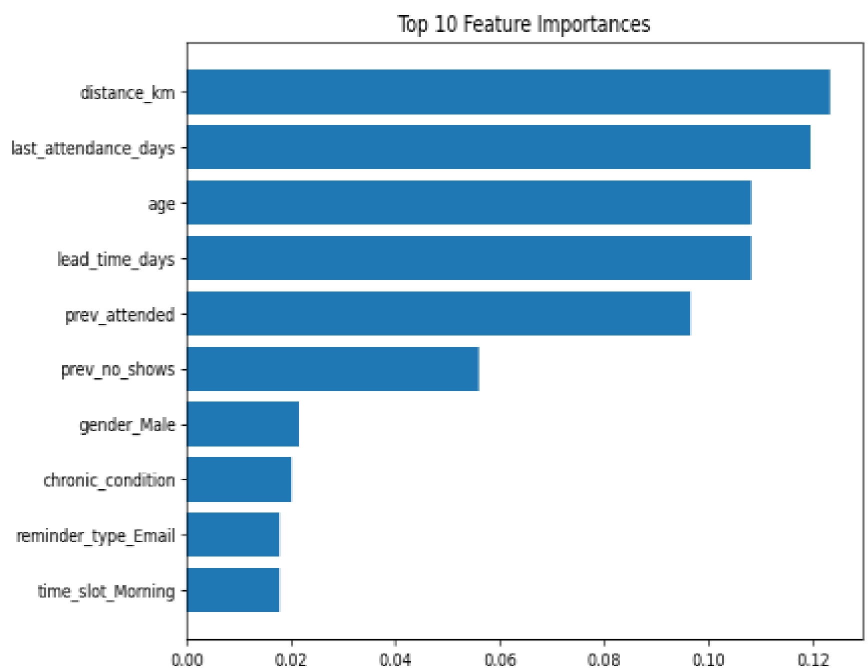

The top predictors were:

1. Distance from clinic (**distance_km**)
2. Days since last attendance (**last_attendance_days**)
3. Age
4. Lead time (**lead_time_days**)
5. Previous attended appointments
6. Previous no-shows

Figure 15: Feature Importance Analysis (Tabular)

	Feature	Importance
1	distance_km	0.123357
6	last_attendance_days	0.119546
0	age	0.108225
3	lead_time_days	0.108129
5	prev_attended	0.096555
4	prev_no_shows	0.056002
7	gender_Male	0.021503
2	chronic_condition	0.019995
29	reminder_type_Email	0.017754
28	time_slot_Morning	0.017695

Distance from clinic proved to be the most powerful predictor indicating that transportation burden has a significant effect on attendance likelihood. Days since last attendance also correlated highly, suggesting an increment of risk over time of disengagement. Lead time reinforced earlier exploratory findings as it reinforces the operational impact of scheduling horizon.
Interestingly, reminder type variables were not one of the strongest predictors. This implies that blanket reminder strategies may have a lesser impact than targeted intervention based on behavioral risk profiling.

5.5 Financial Simulation Outcomes
The financial simulation compared baseline scheduling with AI-assisted intervention.

- Under the baseline scenario (no predictive system), the dataset that was used for testing showed revenue losses of $68,520 because of no-shows.
- Under the AI assisted scenario, high-risk patients (predicted probability >= 0.30) were offered targeted intervention. Assuming that a reduction of 40% of no-shows among high-risk patients, AI-assisted revenue loss was reduced to $49,440.
- After considering costs of interventions ($5 for high-risk patient), the overall net financial improvement was $14,050 for the test data set of 2,000 appointments.

Figure 16: Revenue Loss Simulation Results

```
Baseline Revenue Loss: 68520
AI-Assisted Revenue Loss: 49440
Intervention Cost: 5030
Net Financial Improvement: 14050
```

Scaling this improvement proportionally suggests substantial annualized impact for larger clinics.

5.6 Revenue Sensitivity Analysis
To assess robustness, intervention effectiveness was varied from 20% to 60%.
Even under conservative assumptions of 20% effectiveness of intervention, the system generated positive net financial improvement ($2,890). Financial gains rose by larger amounts as effectiveness levels increased and reached $18,730 at 60% effectiveness.

Figure 17: Revenue Sensitivity Analysis

	Intervention Effectiveness	Net Financial Improvement
0	0.2	2890
1	0.3	7450
2	0.4	12130
3	0.5	15970
4	0.6	18730

This demonstrates that the economic viability of the predictive framework is not highly sensitive to optimistic assumptions. The financial case remains positive across multiple scenarios.

6 Operational & Financial Impact

6.1 Baseline Revenue Exposure

The baseline simulation scenario is a traditional clinic scheduling system with no integration of predictive analytics. Under this scenario, scheduling of an appointment takes place without consideration of risk stratification and all patients are treated the same when it comes to the reminder and intervention strategy.

Within the test dataset containing 2000 appointments, 571 appointments were no-shows. 28.56% non-attendance rate observed during exploratory analysis at an average of $120 per appointment, the $68,520 total loss of revenue due to no-shows under the baseline scenario can be calculated.

This figure is direct revenue loss from empty clinical capacity. It does not take into account secondary costs such as idle time of staff, opportunity cost of unused time slots, or inefficiency of administration. Therefore, the true economic burden in an actual physical environment could be even higher. The baseline scenario shows the associated financial vulnerability with unmanaged risk of attendance.

6.2 AI-Assisted Risk-Based Intervention

The AI-assisted scenario integrates predictive risk scoring into the workflow of scheduling appointments (Toker et al., 2024). Using the Random Forest classifier and an optimized probability threshold of 0.30, patients who were considered high-risk were identified before the execution of the appointment.

The choice of threshold was economically justified using precision-recall sensitivity analysis, with a recall of 60% and a resulting acceptable intervention volume. High-risk patients were assumed to get targeted intervention in the form of personalized phone call reminders at a cost per patient of $5.

Figure 18: Revenue Loss Comparison (Baseline vs AI-Assisted)

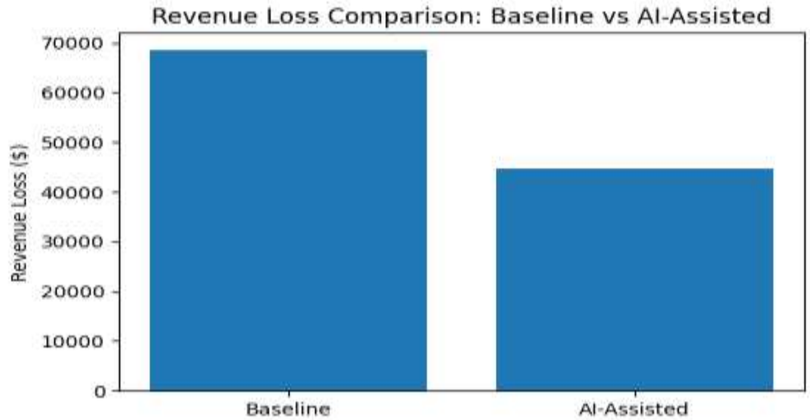

In order to simulate the effectiveness of intervention, it was conservatively assumed 40% of high-risk patients that would not otherwise attend were converted to attendance through targeted outreach. This assumption is based on moderate effectiveness of intervention and does not tend to over project.

Under this AI-assisted scenario, revenue loss decreased from $68,520 to $49,440 within the test dataset.

6.3 Net Financial Improvement

Although AI-assisted intervention reduced revenue loss by $19,080, intervention costs must be considered. A total of $5,030 was spent on reminder calls for high-risk patients.

After accounting for intervention costs, the net financial improvement was calculated as:

Baseline loss − AI-assisted loss − intervention cost

= $68,520 − $49,440 − $5,030

= $14,050

This result demonstrates that predictive analytics combined with targeted intervention yields positive net financial returns even under moderate predictive performance (AUC ≈ 0.58).

Figure 19: Net Financial Improvement After Intervention Costs

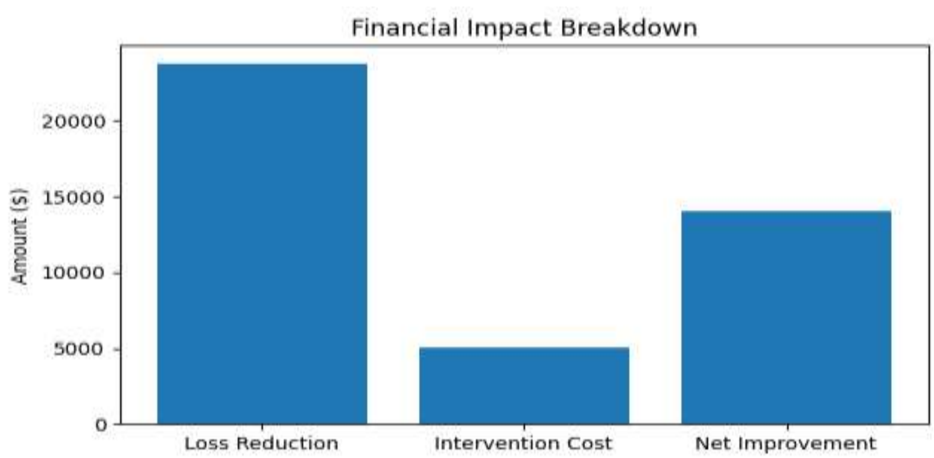

Importantly, this improvement was achieved without having to increase patient volume, hire additional staff, or increase the infrastructure of the clinic. The gains were obtained solely through the increase in scheduling efficiency and risk-based intervention.
If scaled up proportionately, a clinic with 10,000 appointments per year could expect much greater economic gains, going on assumptions of similar distributions of behaviour.

6.4 Sensitivity Analysis of Intervention Effectiveness

To assess the robustness, the effectiveness of interventions was varied between 20 to 60%. Results show that even at a conservative level of effectiveness (20%), the system created a positive net improvement of $2,890.

- As the effectiveness of intervention increased to 30%, 40%, 50% and 60%, net financial gains increased progressively to $7,450, $12,130, $15,970, and $18,730 respectively.
- This monotonic increase supports the conclusion that financial viability does not depend on optimistic assumptions. Even with relatively low intervention effectiveness, predictive targeting is still economically rational.
- The importance of the sensitivity analysis is to strengthen the managerial argument that investment in predictive analytics is financially defensible under multiple operational conditions.

Figure 20: Sensitivity Analysis (Effectiveness vs Net Improvement)

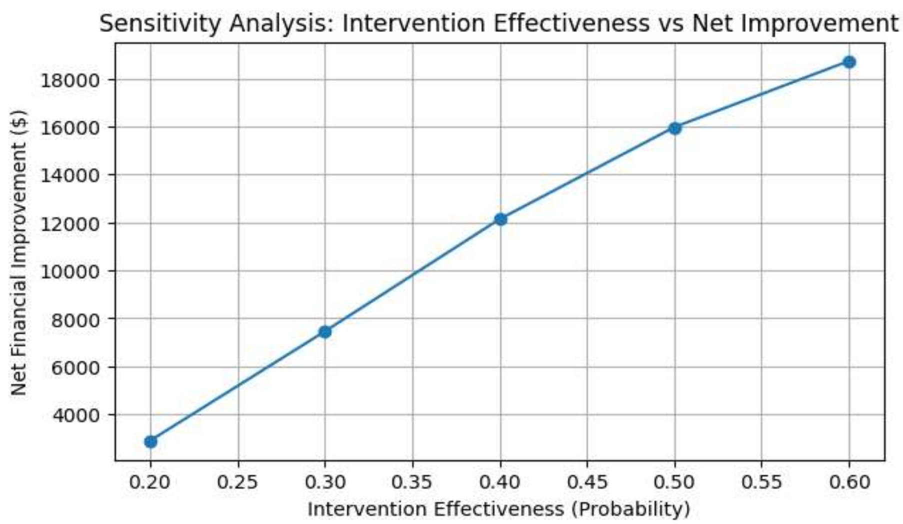

6.5 Operational Trade-Off Analysis

The economic logic of the AI assisted framework is based on cost sensitive classification. Missing a no-show (false negative) has a cost of $120, and falsely identifying a low-risk patient (false positive) has an intervention cost of only $5.
This asymmetric cost structure explains the importance of favouring recall over precision. At the chosen threshold of 0.30, the model was able to predict 60% of actual no-shows. Although the precision was about 34%, the fact that the intervention cost was low means that false positives were not reducing the financial gains.

Figure 21: Threshold Trade-Off (Precision vs Recall)

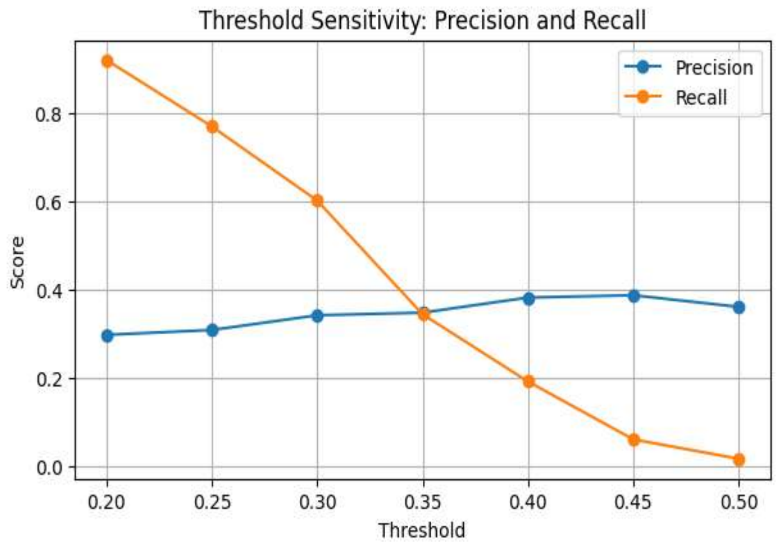

The results therefore demonstrate that predictive value does not depend solely on high statistical discrimination. Instead, value emerges from aligning model thresholds with economic objectives.

7 Strategic Implications for Clinic Management

The findings suggest several operational implications:

- First, blanket reminder strategies are less efficient than risk-targeted intervention. By targeting outreach among high-risk patients, clinics will be able to invest their resources more efficiently.
- Second, predictive analytics can be used for dynamic scheduling policies. For example, high-risk appointments might be strategically double-booked or given an alternative of telehealth.
- Third, financial simulation shows that the predictive systems can drive a measurable ROI even with a moderate level of predictive performance.
- Finally, incorporation of predictive analytics in appointment systems may improve resource utilization without the need for expanding the clinical capacity structurally.

8 Discussions

8.1 Behavioural and Structural Determinants of No-Show Risk

The exploratory analysis provides strong empirical evidence showing that patient no-show behaviour is not random but patterned and structurally influenced. As shown in Figure 4, no-show probability rises successively with past missed appointments, reaching about 36% in the case of patients with history of repeated past non-attendance. This trend of escalation is consistent with behavioural persistence theory; whereby previous adherence behaviour reinforces future attendance patterns. The steep gradient in the escalation plot suggests that the prior no-show history is not simply correlated with future absence but is perhaps indicative of an embedded behavioural trajectory.

Similarly, Figure 2 shows a definite upward trend of non-attendance rates with increasing distance from the clinic. Distance turned out to be the most important predictor in the Random Forest feature importance analysis, which seemed to indicate that accessibility constraints have a greater impact than many demographic variables (Tan et al., 2021). This supports the interpretation that there may be a critical role played by structural barriers, such as transportation burden or travel inconvenience, in appointment adherence. The combined effect

of behavioural history and accessibility suggests that predictive systems have to include both individual level and structural level signals to produce meaningful risk stratification.
This interpretation is strengthened by lead time analysis. As has been noted earlier in the lead time bar chart, the higher non-attendance rates of appointments scheduled 30-60 days in advance, as compared to short-term bookings. This temporal effect suggests that psychological commitment will be weaker the longer the time between scheduling and appointment. From an operational perspective, this helps to support differentiated reminder intensity based on scheduling horizon.

8.2 Predictive Performance in Context

The Random Forest model was able to achieve a ROC-AUC of approximately 0.58 on the test set, while the 5-fold cross validation resulted in a mean AUC of approximately 0.57 and low standard deviation. While these values are only moderately discriminatory power, they should be interpreted in the operational framework of the Research.
The ROC curve shows that the model is better than random classification but does not show high separation between classes. However, predictive discrimination alone is not the determining factor in the utility of an operation. Under the default threshold of 0.50 recall for no-shows was very low (around 1%), making the model essentially useless for intervention targeting. This underscores the shortfalls of using default classification settings in cost sensitive healthcare settings.

Figure 22: ROC Curve

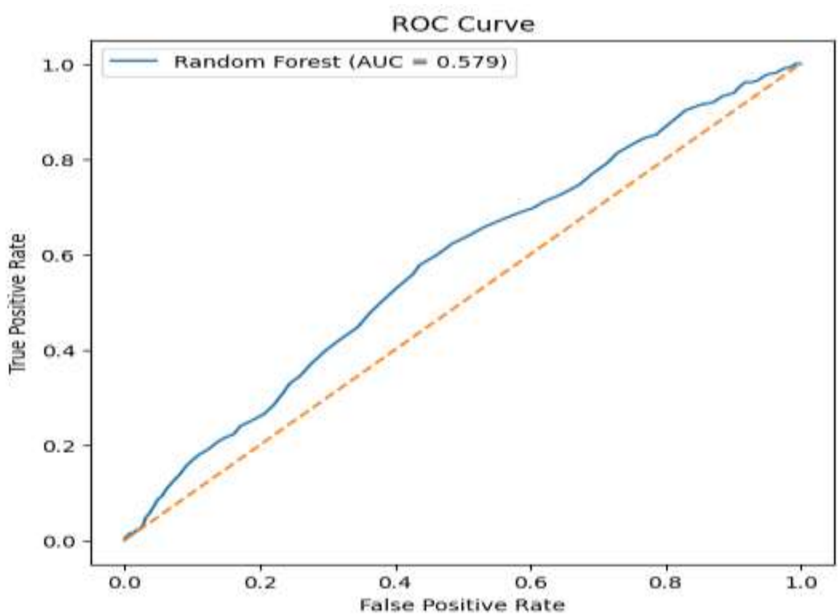

The precision-recall curve and threshold sensitivity analysis show that model utility is improved quite considerably if the threshold for classification is changed. At a level of 0.30, recall was closer to 60% allowing for a majority of high-risk patients to be identified. Although this precision declined to about 34%, this is an economically defensible trade-off given the asymmetry between revenue lost ($120 per missed appointment) and intervention cost ($5 per reminder call).
Figure 23: Precision–Recall Curve

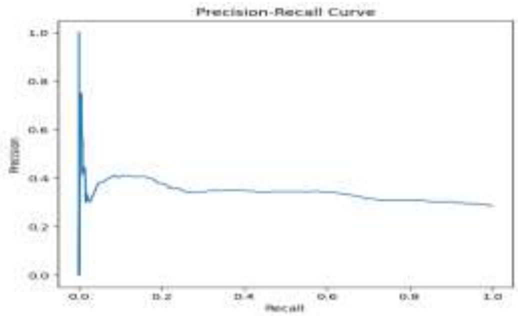

Thus, the moderate AUC does not affect the value of the model. Instead, the results show that economic alignment can bring up to par imperfect discrimination with appropriately calibrated decision thresholds.

8.3 Cost-Sensitive Threshold Optimization and Decision Logic

The threshold sensitivity table offers important information about the trade-off between intervention intensity and predictive selectivity. At a threshold of 0.20, the recall was greater than 90%, but the precision fell significantly so that excess of intervention was allocated. On the other hand, at thresholds greater than 0.35, there was a sharp deficit in recall which made missing no-shows more likely.

The chosen threshold of 0.30 is a decision point where the value is economically balanced. This threshold results in maximal recall with minimal false positive inflation. The choice has thus not been based on statistical optimization only, but on cost-sensitive decision logic.

8.4 Financial Impact and Revenue Optimization

The financial comparison of the baseline versus AI-assisted scenarios gives the greatest validation of the predictive framework. As can be seen in the revenue comparison figure, baseline revenue loss was $68,520 in the test dataset. The AI assisted approach trimmed this down to $49,440, resulting in a gross reduction of $19,080. After considering the cost of intervention, the net financial improvement was $14,050.

This further supported by the sensitivity analysis. Even under conservative assumptions of the effectiveness of the interventions, the system had positive financial returns. The monotonicity of net improvement as intervention effectiveness increases is a confirmation of the structural robustness of the economic model.

These results suggest that predictive systems may not be able to produce ROI through perfect prediction, but through targeted, low-cost interventions to reduce high-cost events (missed appointments).

8.5 Strategic and Implementation Considerations

The findings suggest that uniform reminder strategies may be less than optimal as compared to risk-based allocation. Reminder type analysis indicated moderate effects; however, results of feature importance analysis suggest that reminder strategy is not sufficient if it is simply not predictive in targeting. Therefore, the optimal operational strategy is the combination of risk scoring with selective intervention.

Nonetheless, there are still challenges in implementation. Integration into real-world clinical systems would require compatibility with electronic health records, training of staff, and compliance with data protection regulations. Additionally, issues of ethics involve using demographic or accessibility variables for risk classification. Clinics will need to be mindful of disproportionate disadvantage in vulnerable populations based on predictive systems.

Despite these difficulties, the Research shows that predictive analytics can be used to improve operational efficiency when integrated into a decision framework that is linked to economic inputs.

9 Limitations

Despite showing promising operational and financial results, there are some limitations that should be considered when interpreting results from this Research.

First, the research is based entirely on a generally gathered dataset as opposed to real-world clinical data. Although the dataset was built on realistic behavioural assumptions and probability rules, it is not possible to completely simulate the complexity of real outpatient environments. Real patient behaviour may be affected by other variables that are not accounted for in the model, such as job schedules, family commitments, reliability of transportation, weather, unexpected changes in health, etc. Consequently, the predictive performance as seen in this simulated environment may not be the same as results seen in actual clinical environments.

Second, the probability generation logic that is used to build the no-show outcome has inherent effects on model learnability. Because the labels were created from structured increments of probability, the relationships may be more regular and stable than those in real-world datasets, which usually contain noise, missing values and unobserved confounding variables (Giamattei et al., 2024). Although the model was able to predict with moderate levels of discrimination this could be over or underestimating predictive capability in the real world.

Third, the financial simulation framework is anchored on simplified economic assumptions. The fixed assumption of $120 for appointment revenue and $5 for the cost of intervention per patient were used in this analysis. In reality, appointment revenue depends on the type of service, duration of consultation, and the payer structure. Similarly, intervention costs may vary based on staffing time, communication platform, and administrative overhead. While sensitivity analysis was undertaken to change intervention effectiveness, cost variability itself wasn't modelled to any great extent. Therefore, the financial estimates should be taken as indicative and not definitive.

Fourth, the effectiveness of intervention was assumed, instead of being empirically validated. The simulation assumed that targeted outreach would convert 40% of no-shows predicted to show up for appointment (Carreras-García et al., 2020). Although sensitivity analysis examined levels of effectiveness between 20% and 60%, so the actual behavioural impact of reminders in the real world of clinical settings may vary considerably. A randomised controlled trial would be needed to provide an empirical estimate of intervention efficacy.

Fifth, while threshold optimization was performed to match predictive outputs with cost-sensitive decision logic, more advanced model tuning techniques were not used. Hyperparameter optimization, feature engineering or other ensemble techniques such as gradient boosting might have worked better in predictive discrimination (Si et al., 2024). The objective of the Research, however, was to assess the operational integration and not to achieve the highest possible predictive performance.

Sixth, no empirical assessment was done on ethical and fairness considerations. Such variables as socioeconomic band and distance from clinic may correlate with structural inequalities (Laajaj et al., 2022). In a real-world deployment situation, predictive systems might inadvertently perpetuate disparities if they are not carefully audited. Future research should include measures of fairness and the analysis of bias to ensure fair implementation.

Finally, the simulation only considers revenue loss, and it does not include broader impacts on the healthcare system (Elendu et al., 2024). For example, decreased no-shows could impact patient health outcomes, waiting times and staff productivity. These secondary benefits in the current framework were not monetized.

10 Conclusions

This Research set out to evaluate the feasibility of predictive analytics in the integration into the operations of outpatient clinics to mitigate patient no-shows and optimise revenues performance. Using a simulation-based framework, the research showed that machine learning models - especially Random Forest models - are able to detect behavioural and structural determinants of non-attendance, such as prior no-show history, lead time and distance from clinic. Although predictive discrimination was not extremely high, the results support the notion that the utility of models should not be evaluated based on statistical measures alone, such as AUC.

A major contribution of the Research is a combination of threshold optimisation and cost-sensitive decision logic. By tuning the classification threshold to capture the asymmetric financial costs of missed appointments and interventions to send reminders, this model had much better operational relevance. At the chosen threshold, 60% of high-risk patients were identified, and appropriate intervention strategies could be implemented.

Most importantly, the financial simulation proved that AI assisted scheduling can help to create real economic benefits. The combination of predictive risk scoring and targeted outreach prevented loss of revenue, and created positive financial improvement on net, even with conservative assumptions of intervention effectiveness. Sensitivity analysis provided further evidence that these gains were robust to varying levels of effectiveness.

Overall, the Research progresses a decision-oriented approach to predictive healthcare analytics, which ties together behavioural modelling, operational intervention and financial evaluation in a unified framework. While the direct generalisability has been limited by the simulation-based design, the results actually provide strong evidence that predictive systems, when aligned with economic objectives, can improve clinic efficiency without increasing capacity.

11 References

1 Ala, A., & Chen, F. (2022). Appointment Scheduling Problem in Complexity Systems of the Healthcare Services: A Comprehensive Review. Journal of Healthcare Engineering, 2022(5819813), e5819813. https://doi.org/10.1155/2022/5819813

2 Ao, Y., Li, H., Zhu, L., Ali, S., & Yang, Z. (2019). The linear random forest algorithm and its advantages in machine learning assisted logging regression modeling. Journal of Petroleum Science and Engineering, 174, 776–789. https://doi.org/10.1016/j.petrol.2018.11.067

3 Brikci, N., Alao, R., Wang, H., Wang, H., Erlangga, D., & Hanson, K. (2024). Improving the Efficiency in Spending for health: a Systematic Review of Evidence. SSM. Health Systems, 2, 100008–100008. https://doi.org/10.1016/j.ssmhs.2024.100008

4 Carreras-García, D., Delgado-Gómez, D., Llorente-Fernández, F., & Arribas-Gil, A. (2020). Patient No-Show Prediction: A Systematic Literature Review. Entropy, 22(6), 675. https://doi.org/10.3390/e22060675

5 Chen, A. M. (2023). Socioeconomic and demographic factors predictive of missed appointments in outpatient radiation oncology: An evaluation of access. Frontiers in Health Services, 3, 1288329. https://doi.org/10.3389/frhs.2023.1288329

6 Chiruvella, V., & Guddati, A. K. (2021). Ethical Issues in Patient Data Ownership. Interactive Journal of Medical Research, 10(2), 1–9. https://doi.org/10.2196/22269

7 Conzuelo Rodriguez, G., Bodnar, L. M., Brooks, M. M., Wahed, A., Kennedy, E. H., Schisterman, E., & Naimi, A. I. (2021). Performance Evaluation of Parametric and Nonparametric Methods When Assessing Effect Measure Modification. American Journal of Epidemiology, 191(1), 198–207. https://doi.org/10.1093/aje/kwab220

8 Deina, C., Fogliatto, F. S., da, C., & Anzanello, M. J. (2024). Decision analysis framework for predicting no-shows to appointments using machine learning algorithms. BMC Health Services Research, 24(1). https://doi.org/10.1186/s12913-023-10418-6

9 Dominguez-Catena, I., Paternain, D., & Galar, M. (2024). DSAP: Analyzing bias through demographic comparison of datasets. Information Fusion, 115, 102760–102760. https://doi.org/10.1016/j.inffus.2024.102760

10 Duerstock, B., Mesa, J., Aldoukhi, M., Angel, M., Bare, L., Blacher, D., Cawi, J., Chabanel, T., Ciccone, M., Clapper, K., Jai, S., Ferrer, A., Figueroa, P., Fraticelli, K., Freudenthal, B., Glime, M., Gupta, M., Henley, I., Hill, J., & Huber, E. (2023). Access of Persons with Disabilities to Public Ground Transportation and Roadways. https://docs.lib.purdue.edu/cgi/viewcontent.cgi?article=1005&context=ugcw

11 Elendu, C., Amaechi, D. C., Okatta, A. U., Amaechi, E. C., Elendu, T. C., Ezeh, C. P., & Elendu, I. D. (2024). The impact of simulation-based training in medical education: a review. Medicine, 103(27), 1–14. https://doi.org/10.1097/MD.0000000000038813

12 Fan, G., Deng, Z., Ye, Q., & Wang, B. (2021). Machine learning-based prediction models for patients no-show in online outpatient appointments. Data Science and Management, 2, 45–52. https://doi.org/10.1016/j.dsm.2021.06.002

13 Giamattei, L., Guerriero, A., Pietrantuono, R., & Russo, S. (2024). Causal reasoning in Software Quality Assurance: A systematic review. Information and Software Technology, 178, 107599. https://doi.org/10.1016/j.infsof.2024.107599

14 Gouveia, A., Mancinetti, M., Genné, D., Méan, M., John, G. R., Bütikofer, L., Aujesky, D., Schnipper, J. L., & Donzé, J. (2023). Effectiveness of Transition Care Intervention Targeted to High-Risk Patients to Reduce Readmissions: Study Protocol for the TARGET-READ Multicenter Randomized-Controlled Trial. 11(6), 886–886. https://doi.org/10.3390/healthcare11060886

15 Humphreys, P., Spratt, B., Tariverdi, M., Burdett, R. L., Cook, D., Yarlagadda, P. K. D. V., & Corry, P. (2022). An Overview of Hospital Capacity Planning and Optimisation. Healthcare, 10(5), 826. ncbi. https://doi.org/10.3390/healthcare10050826

16 Issah, I., Appiah, O., Appiahene, P., & Inusah, F. (2023). A systematic review of the literature on machine learning application of determining the attributes influencing academic performance. Decision Analytics Journal, 7, 100204. https://doi.org/10.1016/j.dajour.2023.100204
17 Laajaj, R., Webb, D., Aristizabal, D., Behrentz, E., Bernal, R., Buitrago, G., Cucunubá, Z., de la Hoz, F., Gaviria, A., Hernández, L. J., De Los Rios, C., Ramírez Varela, A., Restrepo, S., Schady, N., & Vives, M. (2022). Understanding how socioeconomic inequalities drive inequalities in COVID-19 infections. Scientific Reports, 12(1), 8269. https://doi.org/10.1038/s41598-022-11706-7
18 Levy, J. J., & O'Malley, A. J. (2020). Don't dismiss logistic regression: the case for sensible extraction of interactions in the era of machine learning. BMC Medical Research Methodology, 20(1). https://doi.org/10.1186/s12874-020-01046-3
19 Liu, D., Shin, W.-Y., Sprecher, E., Conroy, K., Santiago, O., Wachtel, G., & Santillana, M. (2022). Machine learning approaches to predicting no-shows in pediatric medical appointment. Npj Digital Medicine, 5(1). https://doi.org/10.1038/s41746-022-00594-w
20 McLean, S., Booth, A., Gee, M., Salway, S., Cobb, M., Bhanbhro, S., & Nancarrow, S. (2016). Appointment reminder systems are effective but not optimal: results of a systematic review and evidence synthesis employing realist principles. Patient Preference and Adherence, 10, 479. https://doi.org/10.2147/ppa.s93046
21 Opon, S. O., Tenambergen, W. M., & Njoroge, K. M. (2020). The effect of patient reminders in reducing missed appointment in medical settings: a systematic review. PAMJ - One Health, 2(9). https://doi.org/10.11604/pamj-oh.2020.2.9.21839
22 Si, B., Ni, Z., Xu, J., Li, Y., & Liu, F. (2024). Interactive effects of hyperparameter optimization techniques and data characteristics on the performance of machine learning algorithms for building energy metamodeling. Case Studies in Thermal Engineering, 104124–104124. https://doi.org/10.1016/j.csite.2024.104124
23 Silva, J., Paulo, Moreira, M., Prata, D., Trevisan, D., & Santos, O. dos. (2024). Patients' Nonattendance in Outpatient Specialist Consultations: A National Cohort Analysis of a Health System. Risk Management and Healthcare Policy, Volume 17, 2705–2716. https://doi.org/10.2147/rmhp.s468455
24 Tan, L., Main, J. B., & Darolia, R. (2021). Using random forest analysis to identify student demographic and high school-level factors that predict college engineering major choice. Journal of Engineering Education, 110(3), 572–593. https://doi.org/10.1002/jee.20393
25 Toker, K., Ataş, K., Mayadağlı, A., Görmezoğlu, Z., Tuncay, I., & Kazancıoğlu, R. (2024). A Solution to Reduce the Impact of Patients' No-Show Behavior on Hospital Operating Costs: Artificial Intelligence-Based Appointment System. Healthcare, 12(21), 2161–2161. https://doi.org/10.3390/healthcare12212161
26 Wang, X., Zhu, X., Anwar, M. K., Meng, Q., & Zhong, N. (2024). Evaluating the Role of AI and Empirical Models for Predicting Regional Economic Growth and Transportation Dynamics: An Application of Advanced AI Approaches. International Journal of Transportation Science and Technology. https://doi.org/10.1016/j.ijtst.2024.08.007
27 Xie, Z., & Or, C. (2017). Associations Between Waiting Times, Service Times, and Patient Satisfaction in an Endocrinology Outpatient Department: A Time Study and Questionnaire Survey. INQUIRY: The Journal of Health Care Organization, Provision, and Financing, 54, 004695801773952. https://doi.org/10.1177/0046958017739527
28 Zhou, Y., Viswanatha, A., Motaleb, A. A., Lamichhane, P., Chen, K.-Y., Young, R., Gurses, A. P., & Xiao, Y. (2023). A predictive decision analytics approach for primary care operations management: A case study of double-booking strategy design and evaluation. Computers & Industrial Engineering, 177, 109069. https://doi.org/10.1016/j.cie.2023.109069