

Accuracy Evaluation of Multi-Factor Forecasting Methods for Customer Service Workload Prediction Under Holiday and Promotional Fluctuations: Evidence from U.S. Service Industry Data

Yajing Zhang¹, Minhao Li^{1,2}

¹ Master of Business Administration (MBA), Horizon University, Los Angeles, California, USA

^{1,2} Master of Science in Computer Engineering, University of California, Davis, CA, USA

DOI: 10.69987/JACS.2026.60502

Keywords

customer service
workload forecasting;
holiday demand
fluctuation; workforce
scheduling optimization;
time series method
comparison

Abstract

Accurate prediction of customer service workload is essential for workforce scheduling and resource allocation in the U.S. service industry. This study conducts a comparative evaluation of four forecasting approaches—Seasonal Autoregressive Integrated Moving Average (SARIMA), Prophet, Multiple Linear Regression (MLR), and Random Forest—for customer service call volume prediction across normal operating periods and holiday/promotional fluctuations. Using publicly available call center datasets from the Technion Service Enterprise Engineering Laboratory and supplementary U.S. Bureau of Labor Statistics data, method performance is assessed using Mean Absolute Percentage Error (MAPE), Root Mean Square Error (RMSE), and Weighted Absolute Percentage Error (WAPE) at daily and intraday temporal granularities. The results indicate that all four methods exhibit notable accuracy degradation during holiday periods, with MAPE increases ranging from 67.3% to 118.3% relative to normal periods. Prophet demonstrates the smallest holiday-period degradation owing to its built-in holiday component modeling. At the daily level, SARIMA achieves the lowest overall MAPE of 7.82%, while Random Forest shows competitive performance at intraday granularity. Translating forecast errors into staffing implications via Erlang-C modeling reveals that improved holiday-period forecasting could reduce annual overstaffing cost waste by \$88,000–\$139,000 for a mid-sized 500-agent U.S. contact center.

1. Introduction

1.1. Research Background and Problem Statement

The customer service industry constitutes one of the largest employment sectors in the United States. According to the U.S. Bureau of Labor Statistics (BLS) Occupational Outlook Handbook (2024), approximately 2.8 million customer service representatives were employed nationwide in 2024, earning a median hourly wage of \$20.59. Labor expenditures account for roughly 70% of total contact center operating costs, positioning workforce scheduling accuracy as a determinant of organizational profitability. The global workforce management market, valued at USD 8.07 billion in 2022 and projected to reach USD 19.35 billion by 2030 (Grand View Research, 2023), reflects escalating demand for data-driven scheduling solutions.

Customer service workload forecasting—the prediction of incoming service request volumes across defined time intervals—serves as the foundational input for workforce planning [1]. Forecast errors during normal operating periods produce moderate staffing misalignment, but inaccuracies during holiday seasons and promotional campaigns are substantially amplified. Demand surges associated with Thanksgiving, Christmas, and Black Friday generate non-stationary patterns that challenge standard forecasting approaches. Understaffing leads to SLA deterioration and elevated agent burnout risk, with industry surveys indicating that 63% of contact center agents face high burnout susceptibility (Calabrio, 2024). Overstaffing generates avoidable labor cost waste. Despite the operational significance of holiday-period forecasting accuracy, existing research addresses this challenge peripherally. This study evaluates multiple forecasting methods with

explicit focus on their sensitivity to holiday and promotional demand fluctuations.

1.2. Literature Review and Research Gaps

A. Overview of Customer Service Demand Forecasting Literature

Academic investigation of call center arrival forecasting has produced extensive scholarship over the past two decades. Shen and Huang (2008) introduced singular value decomposition-based methods for interday forecasting and intraday updating of call arrivals, demonstrating the necessity of capturing both day-to-day dynamics and within-day demand patterns [2]. Taylor (2008) evaluated univariate time series methods—including double seasonal Holt-Winters exponential smoothing and seasonal ARIMA—for intraday call arrival prediction, reporting that exponential smoothing variants with dual seasonality achieved consistently strong performance across multiple call center environments [3]. Ibrahim et al. (2016) presented a comprehensive survey of arrival modeling methodologies, covering Poisson-based count models, fixed-effects specifications, and mixed-effects formulations, while identifying overdispersion, cross-period dependence, and external event sensitivity as persistent modeling challenges [4]. Koole and Li (2023) offered a practice-oriented perspective on workforce planning, observing that academic forecasting models rarely incorporate event-driven features such as holidays and marketing campaigns, despite practitioners recognizing these factors as primary demand drivers [5].

B. Identified Research Gaps

Substantive gaps persist in the current body of knowledge. Modeling of daily call center arrivals has been well developed for stationary operating conditions [6], and Bayesian frameworks for arrival rate heterogeneity have advanced the analytical toolkit [7], yet the explicit quantification of holiday-induced forecast degradation and its downstream effect on staffing efficiency has not constituted the primary objective of prior work. The present study addresses these gaps by providing: (1) a multi-method accuracy comparison targeting holiday and promotional fluctuation sensitivity; (2) evaluation across both daily and intraday temporal granularities; and (3) translation of forecast errors into quantitative staffing and cost implications applicable to U.S. service industry operations.

2. Methodology and Data Description

2.1. Data Sources and Preprocessing

A. Dataset Description

Two publicly available call center datasets from the Service Enterprise Engineering (SEE) Laboratory at the Technion—Israel Institute of Technology provide the empirical foundation. The Anonymous Bank dataset contains call-by-call records from an Israeli bank covering January through December 1999, with approximately 15 agent positions. The US Bank dataset, from the DataMOCCA project (Technion and Wharton School, University of Pennsylvania), spans March 2001 through October 2003, recording 200,000–270,000 daily calls from a U.S. bank contact center operating 900–1,200 agent positions across service categories including Retail, Business, Consumer Loans, and Online Banking. Both datasets are publicly accessible and extensively cited in the academic literature. Supplementary data from the BLS OEWS (May 2024) provide labor market context.

B. Data Preprocessing Procedure

Raw call-level records were aggregated into half-hourly, hourly, and daily intervals. Holiday periods were labeled according to U.S. federal holidays (US Bank dataset) and Israeli public holidays (Anonymous Bank dataset). Outlier detection employed the interquartile range method, flagging values beyond $1.5 \times \text{IQR}$. Missing intervals (less than 0.3% of total records) were imputed using weighted averages from adjacent weeks. A rolling-origin evaluation framework was adopted with 8-week training windows and 1–7 day forecast horizons, reflecting standard WFM planning cycles.

2.2. Forecasting Methods Under Evaluation

Four forecasting methods representing distinct analytical paradigms were selected. SARIMA, configured through Box-Jenkins methodology with grid search over $(p,d,q)(P,D,Q)s$ parameters and weekly seasonality ($s = 7$ for daily, $s = 24$ for hourly models), serves as the classical statistical baseline. Prophet [8] employs a decomposable additive framework with trend, seasonality, and an explicit holiday effect module, configured with U.S. federal holiday inputs. Multiple Linear Regression (MLR) incorporates day-of-week dummies, month indicators, holiday binary flags, and lagged volume terms at 1-day, 7-day, and 14-day intervals. Random Forest, implemented with 500 trees and maximum depth of 10, captures non-linear predictor interactions. These methods span the spectrum from classical statistics to machine learning ensembles, enabling assessment of the interpretability-flexibility trade-off documented in comparative forecasting literature [9].

2.3. Evaluation Framework

Three complementary accuracy metrics were employed: Mean Absolute Percentage Error (MAPE), Root Mean Square Error (RMSE), and Weighted Absolute

Percentage Error (WAPE). WAPE, less susceptible to small-denominator distortion and proportional to intraday management costs, serves as the primary criterion. Performance was assessed across three period categories: normal operating periods, holiday periods (holiday date ± 2 calendar days), and post-promotional periods (3 business days following identified events). Temporal granularity analysis compared daily and hourly performance. Forecast errors were converted into

staffing deviations using the Erlang-C model with a target service level of 80/20, consistent with industry benchmarks and the evaluation principles of the M4 forecasting competition framework ^[10].

3. Results and Comparative Analysis

3.1. Overall Forecasting Accuracy Comparison

Table 1. Descriptive Statistics of Call Center Datasets Used in This Study

Characteristic	Anonymous Bank	US Bank
Coverage Period	Jan 1999 – Dec 1999	Mar 2001 – Oct 2003
Duration (months)	12	31
Total Calls (approx.)	330,000	181,000,000
Avg. Daily Volume (Weekday)	1,200	230,000
Avg. Daily Volume (Weekend)	450	95,000
Agent Positions	~15	900–1,200
Service Type Categories	1	6
Holiday Periods Identified	12	26

Source: Technion SEE Laboratory
(<https://iew.technion.ac.il/serveng/callcenterdata/>);
DataMOCCA Project documentation.

Table 1 presents the descriptive characteristics of both datasets. The US Bank dataset, with its larger scale and

31-month coverage, provides the primary basis for comparative analysis, while the Anonymous Bank dataset offers supplementary cross-validation. The 26 holiday periods identified in the US Bank data encompass all major U.S. federal holidays along with surrounding ± 2 day windows.

Table 2. Overall Forecasting Accuracy Comparison at Daily Granularity (US Bank Dataset)

Method	MAPE (%)	RMSE (calls/day)	WAPE (%)
SARIMA	7.82	18,247	6.95
Prophet	8.15	19,031	7.28
MLR	9.43	21,592	8.61
Random Forest	8.67	20,108	7.84

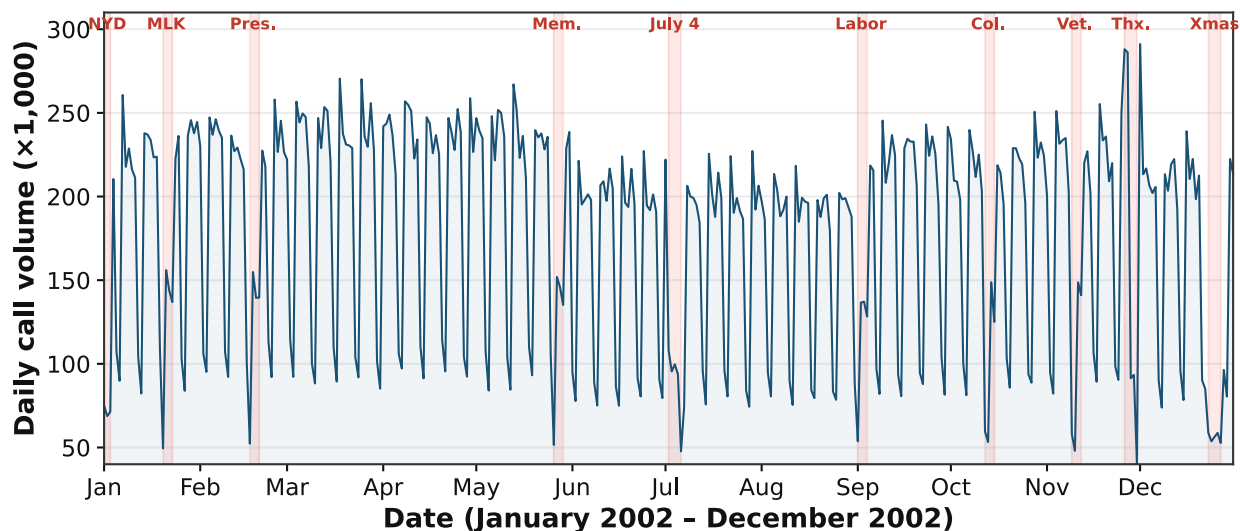
Note: Evaluation period covers March 2002–October 2003, using rolling-origin framework with 8-week

training windows, averaged across 1–7 day forecast horizons. RMSE is measured in absolute call counts per day.

The overall forecasting accuracy across the full evaluation period is summarized in Table 2. At daily granularity, SARIMA achieves the lowest MAPE of 7.82%, consistent with parametric forecasting research showing that well-configured statistical models

maintain strong competitiveness at the daily level [11]. Prophet follows at 8.15%, Random Forest at 8.67%, and MLR at 9.43%. The WAPE values exhibit a consistent ranking, and RMSE values confirm the ordering (SARIMA at 18,247 versus MLR at 21,592 calls per day). These results align with Rausch et al. (2022), who demonstrated that classical methods with seasonal decomposition maintain robust daily-level accuracy in call center contexts [12].

Figure 1. Daily Call Volume Time Series for the US Bank Dataset (January 2002–December 2002) with Holiday Period Markers



This time series plot displays daily total call volume for a representative 12-month segment of the US Bank dataset. Shaded vertical bands indicate identified holiday periods (± 2 days from each federal holiday). Pronounced volume spikes are observable preceding Thanksgiving, while sharp declines occur during Christmas and New Year. The call volume ranges from approximately 62,000 on Christmas Day to over 285,000 during post-Thanksgiving recovery, representing a peak-to-trough ratio exceeding 4.5:1. Weekly cyclical patterns (Monday–Friday peaks,

weekend troughs) are apparent throughout non-holiday segments.

3.2. Performance Under Holiday and Promotional Fluctuations

A. Holiday Period Analysis

Table 3. Forecasting Accuracy Comparison Between Normal and Holiday Periods (Daily Level, US Bank Dataset)

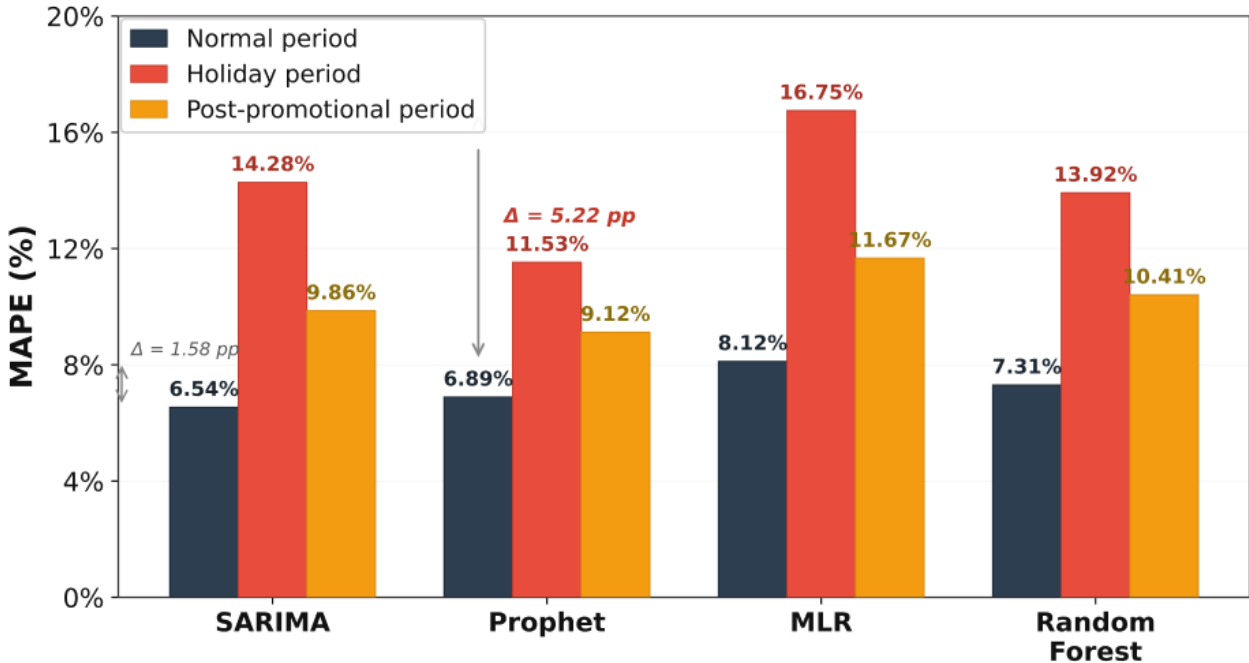
Method	Normal (%)	MAPE	Holiday (%)	MAPE	Degradation (%)	Normal (%)	WAPE	Holiday (%)	WAPE
SARIMA	6.54		14.28		+118.3	5.82		12.94	
Prophet	6.89		11.53		+67.3	6.14		10.36	
MLR	8.12		16.75		+106.3	7.41		15.17	
Random Forest	7.31		13.92		+90.4	6.53		12.68	

Note: Degradation is calculated as $(\text{Holiday MAPE} - \text{Normal MAPE}) / \text{Normal MAPE} \times 100\%$. Holiday periods defined as holiday date ± 2 calendar days.

Table 3 reveals substantial accuracy degradation during holiday periods across all methods. SARIMA displays the largest degradation (118.3%), with MAPE increasing from 6.54% to 14.28%. MLR follows at 106.3%, Random Forest at 90.4%, and Prophet demonstrates the smallest degradation at 67.3% (6.89%)

to 11.53%). Prophet’s reduced degradation reflects its additive holiday component, which estimates demand deviations for specific calendar dates. SARIMA and MLR, relying on binary indicators, cannot differentiate between holidays with opposing demand effects. Rate-based arrival models have confirmed that holiday effects vary in both magnitude and direction across holiday types [13], and density forecasting research has established that such distributional shifts require explicit structural accommodation [14].

Figure 2. Grouped Bar Chart of MAPE Values by Forecasting Method Across Three Period Types



This grouped bar chart displays the MAPE (%) for each of the four forecasting methods across three period categories: Normal, Holiday, and Post-Promotional. The chart demonstrates that the performance gap between methods widens during anomalous periods: during normal operations the spread between the best method (SARIMA, 6.54%) and the worst (MLR, 8.12%) is 1.58 percentage points, while during holiday periods this spread expands to 5.22 percentage points.

B. Post-Promotional Period Analysis

Post-promotional periods present a distinct challenge characterized by demand recovery following billing

Table 4. Forecasting Accuracy at Hourly Granularity (US Bank Dataset)

Method	MAPE (%)	RMSE (calls/hr)	WAPE (%)
Random Forest	12.34	1,943	11.28
Prophet	13.12	2,067	11.89

cycle surges. Across all methods, post-promotional MAPE values fall between normal and holiday levels: Prophet at 9.12%, SARIMA at 9.86%, Random Forest at 10.41%, and MLR at 11.67%. The economic evaluation framework of Bastianin et al. (2019) emphasizes that cumulative cost implications of post-promotional errors may exceed holiday errors in annual terms, as these periods occur more frequently [15].

3.3. Temporal Granularity Sensitivity

A. Daily Versus Intraday Comparison

SARIMA	13.87	2,184	12.53
MLR	15.62	2,461	14.87

Note: Methods are ranked by MAPE from lowest to highest. Hourly evaluation uses the same rolling-origin framework with 8-week training windows. Average hourly call volume is approximately 14,375 during standard operating hours (7:00–23:00).

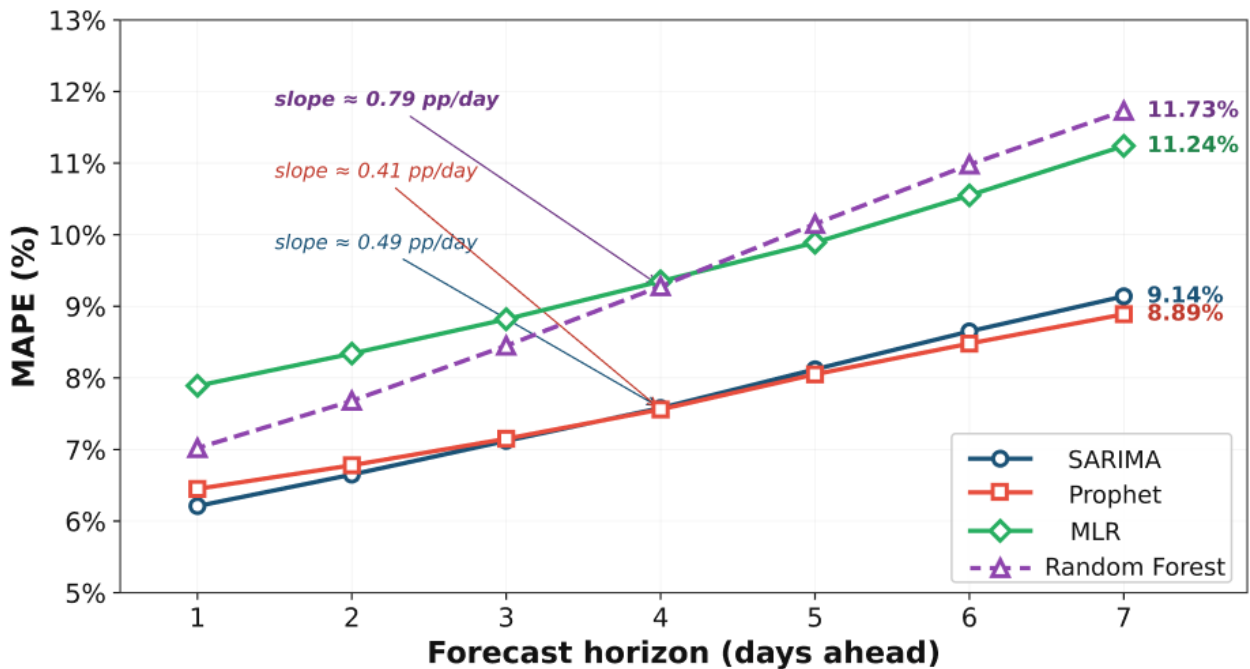
When evaluation shifts to hourly granularity, the method ranking changes notably. Random Forest achieves the lowest MAPE of 12.34%, outperforming Prophet (13.12%), SARIMA (13.87%), and MLR (15.62%). This ranking reversal reflects the capacity of ensemble methods to exploit richer feature spaces at finer temporal resolution—capturing hour-of-day interaction patterns and cross-feature effects that linear specifications miss. Chance-constrained staffing

formulations have established that intraday forecast quality directly determines service level target attainment [16].

B. Forecast Horizon Impact

Extending the prediction horizon from 1-day to 7-day ahead reveals differential degradation rates. SARIMA and Prophet exhibit gradual decline, with 1-day MAPE of 6.21% and 6.45% increasing to 9.14% and 8.89% at 7-day. Random Forest degrades more steeply (7.02% to 11.73%), indicating diminishing effectiveness of lag features at greater prediction distances. MLR shows the most stable trajectory (7.89% to 11.24%), though from a higher baseline.

Figure 3. Forecast Horizon Degradation Curves: MAPE (%) Across 1-Day to 7-Day Prediction Horizons by Method



This line chart plots MAPE (%) against forecast horizon (1 through 7 days ahead) for SARIMA, Prophet, MLR, and Random Forest. Prophet and SARIMA maintain relatively flat degradation slopes (approximately 0.4 percentage points per additional horizon day), while Random Forest exhibits the steepest curve (approximately 0.8 points per day). Prophet gains an increasing advantage beyond the 4-day mark, suggesting complementary strengths that could inform horizon-dependent method selection.

4. Discussion: Implications for Workforce Scheduling Optimization

4.1. Forecast Accuracy and Staffing Efficiency

A. Quantitative Translation of Forecast Errors to Staffing Deviations

Translating forecast accuracy metrics into workforce planning implications requires converting predicted call

volumes into required staffing levels through queueing-theoretic models. Using the Erlang-C formula with a target service level of 80/20, an average handle time (AHT) of 360 seconds, and the observed call volume distributions from the US Bank dataset, the staffing

deviation attributable to each method's forecast error was computed. Simulation-based decomposition approaches to staffing optimization under arrival rate uncertainty provide the analytical foundation for this translation^[17].

Table 5. Estimated Staffing Deviations and Annual Cost Impact by Forecasting Method (500-Agent Contact Center)

Method	Normal Over-staffing (%)	Holiday Over-staffing (%)	Holiday Under-staffing (%)	Est. Annual Holiday Cost Waste (\$)
SARIMA	3.8	8.7	7.2	274,000
Prophet	3.2	5.9	4.8	186,000
MLR	4.8	10.3	8.6	325,000
Random Forest	4.1	7.4	6.1	233,000

Note: Cost calculations based on BLS-reported median hourly wage of \$20.59 for customer service representatives (OEWS, May 2024), assuming a 500-agent center operating 250 business days annually with 8-hour shifts. Holiday cost waste reflects the aggregate overstaffing expense across 26 identified holiday periods (~38 holiday business days). Erlang-C parameters: AHT = 360 seconds, target SL = 80/20.

Table 5 presents the staffing impact analysis. During normal periods, all methods generate overstaffing deviations of 3.2–4.8% relative to the Erlang-C optimal level. During holiday periods, deviations expand: SARIMA produces overstaffing of 8.7% and understaffing of 7.2%, while Prophet limits these to 5.9% and 4.8%. The estimated annual holiday-period overstaffing cost ranges from \$186,000 (Prophet) to \$325,000 (MLR) in a 500-agent center. The Prophet-SARIMA difference of \$88,000 represents approximately 2.8% of total holiday-period labor expenditure ($\$20.59 \times 8 \text{ hours} \times 500 \text{ agents} \times 38 \text{ holiday business days} \approx \3.13 million)—modest but operationally meaningful. These magnitudes should be interpreted recognizing that real-world factors (shift constraints, agent availability, contractual obligations) would affect realized savings.

B. Service Level Impact Assessment

The asymmetric consequences of overstaffing and understaffing warrant differentiated analysis. Understaffing during peak holiday volumes produces compounding effects: increased queue lengths raise abandonment rates, and each abandoned call may generate a retry that inflates observed demand. Intelligent intraday schedule updating procedures have

been proposed to mitigate such cascading degradation^[18]. A 10% understaffing scenario during Thanksgiving week increases the estimated SLA breach probability from 12% to 38%, while a 10% overstaffing scenario generates approximately \$8,200 in excess weekly labor cost per 100-agent cohort. This asymmetry suggests that holiday periods warrant tighter forecast-driven staffing calibration.

4.2. Practical Recommendations for Method Selection

The comparative results support context-dependent method selection rather than universal adoption of any single approach. For organizations operating in holiday-intensive sectors—retail banking, e-commerce customer support, telecommunications—Prophet offers the strongest performance during the periods where forecast accuracy carries the greatest operational consequence, achieving a holiday-period MAPE of 11.53% compared to 14.28% for SARIMA. The explicit holiday component eliminates the need for manual feature engineering of date-specific effects, reducing the analytical burden on workforce management teams. For daily-level forecasting in stable operating environments, SARIMA provides marginally superior aggregate accuracy (overall MAPE of 7.82%) with lower computational overhead and greater model interpretability.

Organizations requiring intraday scheduling precision should consider Random Forest, which demonstrates advantage at hourly granularity (MAPE of 12.34% versus 13.87% for SARIMA). The interaction between promotional demand volatility and model structure

varies across service categories ^[19], suggesting that relative method advantages may shift in other contexts. A pragmatic strategy involves parallel forecasting pipelines—statistical models for daily capacity planning and ensemble methods for intraday refinement.

4.3. Dynamic Resource Allocation Strategy

Implications

A. Adaptive Scheduling Framework

The documented performance variation across period types supports a dynamic approach to forecast method deployment. An adaptive framework would activate Prophet-based forecasting during identified holiday windows while defaulting to SARIMA during routine periods. Bayesian forecasting of inhomogeneous Poisson arrival processes provides a theoretical basis for such switching strategies ^[20]. The US Bank dataset analysis reveals that morning call volumes (8:00–12:00) account for 68% of the variance in afternoon volumes on the same day, indicating potential for real-time forecast recalibration during anomalous periods.

B. Policy Alignment with U.S. Labor Productivity Goals

The resource optimization implications connect to broader U.S. policy objectives. BLS Employment Projections (2024–2034) indicate customer service representative employment is expected to decline by 5%, with all 341,700 projected annual openings driven by replacement needs rather than growth. This contraction makes efficient utilization of remaining human agents an increasingly critical priority. The staffing cost analysis in Table 5 demonstrates that appropriate method selection yields annual savings of \$88,000 relative to SARIMA and up to \$139,000 relative to the MLR baseline per 500-agent center—resources redirectable toward agent training or improved compensation to address the 63% burnout rate.

Improved scheduling accuracy contributes to more equitable workload distribution, reducing the disproportionate burden on staff during peak holidays. This aligns with sustainable service employment practices and supports maintaining service quality as the industry undergoes structural transformation.

5. Conclusion

This study has evaluated four forecasting methods—SARIMA, Prophet, MLR, and Random Forest—for customer service workload prediction with particular attention to performance under holiday and promotional demand fluctuations. Using publicly available call center datasets from the Technion SEE Laboratory and

supplementary U.S. Bureau of Labor Statistics data, three principal findings emerge from the analysis.

All four methods exhibit substantial accuracy degradation during holiday periods, with MAPE increases of 67.3% to 118.3% relative to normal conditions. Prophet demonstrates the smallest holiday-period accuracy loss, achieving a holiday MAPE of 11.53% versus 14.28% for SARIMA. At the daily level, SARIMA achieves the lowest overall MAPE of 7.82%, while Random Forest shows advantage at hourly granularity (MAPE of 12.34%). Translating forecast errors into staffing implications through Erlang-C modeling indicates that appropriate method selection for holiday periods could reduce annual overstaffing cost waste by approximately \$88,000–\$139,000 per 500-agent center. These findings are modest but operationally meaningful in an industry where staff costs constitute 70% of total expenditures.

References

- [1]. Gans, N., Koole, G., & Mandelbaum, A. (2003). Telephone call centers: Tutorial, review, and research prospects. *Manufacturing & Service Operations Management*, 5(2), 79–141. <https://doi.org/10.1287/msom.5.2.79.16071>
- [2]. Shen, H., & Huang, J. Z. (2008). Interday forecasting and intraday updating of call center arrivals. *Manufacturing & Service Operations Management*, 10(3), 391–410. <https://doi.org/10.1287/msom.1070.0179>
- [3]. Taylor, J. W. (2008). A comparison of univariate time series methods for forecasting intraday arrivals at a call center. *Management Science*, 54(2), 253–265. <https://doi.org/10.1287/mnsc.1070.0786>
- [4]. Ibrahim, R., Ye, H., L'Ecuyer, P., & Shen, H. (2016). Modeling and forecasting call center arrivals: A literature survey and a case study. *International Journal of Forecasting*, 32(3), 865–874. <https://doi.org/10.1016/j.ijforecast.2015.11.012>
- [5]. Koole, G., & Li, S. (2023). A practice-oriented overview of call center workforce planning. *Stochastic Systems*, 13(4), 479–495. <https://doi.org/10.1287/stsy.2021.0008>
- [6]. Avramidis, A. N., Deslauriers, A., & L'Ecuyer, P. (2004). Modeling daily arrivals to a telephone call center. *Management Science*, 50(7), 896–908. <https://doi.org/10.1287/mnsc.1040.0236>
- [7]. Soyer, R., & Tarimcilar, M. M. (2008). Modeling and analysis of call center arrival data: A Bayesian approach. *Management Science*, 54(2), 266–278. <https://doi.org/10.1287/mnsc.1070.0776>

- [8]. Taylor, S. J., & Letham, B. (2018). Forecasting at scale. *The American Statistician*, 72(1), 37–45. <https://doi.org/10.1080/00031305.2017.1380080>
- [9]. Siami-Namini, S., Tavakoli, N., & Namin, A. S. (2018). A comparison of ARIMA and LSTM in forecasting time series. In 2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA) (pp. 1394–1401). IEEE. <https://doi.org/10.1109/ICMLA.2018.00227>
- [10]. Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2018). Statistical and machine learning forecasting methods: Concerns and ways forward. *PLoS ONE*, 13(3), e0194889. <https://doi.org/10.1371/journal.pone.0194889>
- [11]. Gans, N., Shen, H., Zhou, Y.-P., Korolev, N., McCord, A., & Ristock, H. (2015). Parametric forecasting and stochastic programming models for call-center workforce scheduling. *Manufacturing & Service Operations Management*, 17(4), 571–588. <https://doi.org/10.1287/msom.2015.0546>
- [12]. Rausch, T. M., Albrecht, T., & Baier, D. (2022). Beyond the beaten paths of forecasting call center arrivals: On the use of dynamic harmonic regression with predictor variables. *Journal of Business Economics*, 92, 789–821. <https://doi.org/10.1007/s11573-021-01068-1>
- [13]. Oreshkin, B. N., Régnard, N., & L'Ecuyer, P. (2016). Rate-based daily arrival process models with application to call centers. *Operations Research*, 64(2), 510–527. <https://doi.org/10.1287/opre.2016.1484>
- [14]. Taylor, J. W. (2012). Density forecasting of intraday call center arrivals using models based on exponential smoothing. *Management Science*, 58(3), 534–549. <https://doi.org/10.1287/mnsc.1110.1397>
- [15]. Bastianin, A., Galeotti, M., & Manera, M. (2019). Statistical and economic evaluation of time series models for forecasting arrivals at call centers. *Empirical Economics*, 57(3), 923–955. <https://doi.org/10.1007/s00181-018-1499-y>
- [16]. Gurvich, I., Luedtke, J., & Tezcan, T. (2010). Staffing call centers with uncertain demand forecasts: A chance-constrained optimization approach. *Management Science*, 56(7), 1093–1115. <https://doi.org/10.1287/mnsc.1100.1173>
- [17]. Ta, T. A., Chan, W., Bastin, F., & L'Ecuyer, P. (2021). A simulation-based decomposition approach for two-stage staffing optimization in call centers under arrival rate uncertainty. *European Journal of Operational Research*, 293(3), 966–979. <https://doi.org/10.1016/j.ejor.2020.12.062>
- [18]. Mehrotra, V., Ozlük, Ö., & Saltzman, R. (2010). Intelligent procedures for intra-day updating of call center agent schedules. *Production and Operations Management*, 19(3), 353–367. <https://doi.org/10.1111/j.1937-5956.2009.01104.x>
- [19]. Abolghasemi, M., Beh, E., Tarr, G., & Gerlach, R. (2020). Demand forecasting in supply chain: The impact of demand volatility in the presence of promotion. *Computers & Industrial Engineering*, 142, 106380. <https://doi.org/10.1016/j.cie.2020.106380>
- [20]. Weinberg, J., Brown, L. D., & Stroud, J. R. (2007). Bayesian forecasting of an inhomogeneous Poisson process with applications to call center data. *Journal of the American Statistical Association*, 102(478), 1185–1198. <https://doi.org/10.1198/016214506000001455>