Rental Price Dynamics in Inner Melbourne: A Suburb-Level Analysis Using SA2 Regions, and Year Fixed Effects (2007–2025) with Forecasts to 2030 of 1-bedroom apartments

Sanjana Bhavin Lakhpatwala

January 2026

Monash Statistical Analytics Club

Table of Contents

1. ACKNOWLEDGEMENT 3

2. ABSTRACT 4

3. INTRODUCTION 5

4. LITERATURE REVIEW 6

5. CONTEXT AND DATA 7

5.1 DATA CLEANING 7

5.2 GEOGRAPHIC ALIGNMENT AND AGGREGATION 7

6. EMPIRICAL METHODOLOGY 9

6.1 MODEL 1: SEASONAL NAÏVE 10

6.1.1 Results 11

6.2 MODEL 2: SEASONAL REGRESSION 12

6.2.1 Results 13

6.3 MODEL 3: SEASONAL ARIMA 15

6.3.1 Results 16

6.4 MODEL 4: ETS MODEL 17

6.4.1 Results 18

6.5 MODELS ANALYSIS 19

6.5.1 Aggregate model performance 19

6.5.2 Suburb-level performance 19

6.6 FINAL MODEL SELECTION 20

7. SARIMA PREDICTIVE ANALYSIS (2026- 2030) 21

7.1 AVERAGE INNER MELBOURNE RENTAL FORECAST 21

7.2 CARLTON NORTH FORECAST 22

RESULTS 23

8. LIMITATIONS 24

9. CONCLUSION 25

10. BIBLIOGRAPHY 26

11. APPENDIX 27

1. Acknowledgement

The author would like to thank the Ms. Anika Jain for her help and contribution as Research Program Director at Monash Student's Analytics Club, throughout this project. The author also extends sincere thanks to Mr. Dustin Keene for his support and assistance in the completion of this paper. The views expressed in this paper are solely those of the author, and any remaining errors are their own.

2. Abstract

This paper examines rental price dynamics across Inner Melbourne using quarterly median rental data for one-bedroom apartments at the Statistical Area 2 (SA2) level[1], compiled from Victorian rental market reports published by the Victorian State Government Department of Families, Fairness and Housing (Victorian State Government Department of Families, Fairness and Housing, 2025). The dataset covers 22 Inner Melbourne SA2 regions over the period 2007Q1 to 2025Q1, capturing both long-run trends and substantial short-run fluctuations in rental prices.

The analysis applies a range of time-series forecasting models, progressing from simple benchmark specifications to more flexible stochastic frameworks. Models are estimated separately for each suburb and evaluated using rolling one-step-ahead out-of-sample forecasts to ensure realistic predictive assessment. Forecast accuracy is measured using root mean squared error (RMSE), mean absolute error (MAE), and mean absolute percentage error (MAPE). Based on aggregate and suburb-level performance, a Seasonal Autoregressive Integrated Moving Average (SARIMA) model is selected for final projection.

Using the selected model, rental prices are forecast through to 2030 at both the individual suburb level and for the average Inner Melbourne rental market, with 95% prediction intervals used to quantify forecast uncertainty. The results indicate considerable heterogeneity in rental dynamics across suburbs, alongside a stabilisation in aggregate rental growth accompanied by widening uncertainty at longer horizons. Overall, the findings highlight the importance of flexible time-series models in capturing the complex and location-specific behaviour of urban rental markets.

[1] Statistical Areas Level 2 (SA2s) are medium-sized general purpose areas built up from whole Statistical Areas Level 1 (SA1s). Their purpose is to represent a community that interacts together socially and economically. (Australian Bureau of Statistics, 2021)

3. Introduction

Rental affordability has become an increasingly prominent economic issue in Australian cities (Lam, 2025), particularly in inner-city areas where population pressures, constrained housing supply, and shifting demand patterns place sustained upward pressure on rents. In Melbourne, these concerns are especially acute in Inner Melbourne, where rental prices differ markedly across nearby locations and fluctuate substantially over time. Such variation complicates affordability assessments and makes it difficult for households, policymakers, and planners to form expectations about future rental market conditions.

Understanding rental price dynamics in Inner Melbourne is therefore important for evaluating housing stress, anticipating affordability challenges, and informing urban policy discussions. While much of the existing debate focuses on structural drivers such as income growth, housing supply, and demographic change, rental markets are also characterised by strong temporal patterns, including persistence, seasonality, and cyclical movements. These features suggest that rental prices may exhibit predictable behaviour over time even in the absence of detailed information on underlying market fundamentals.

This paper contributes to the rental affordability literature by examining how rental prices evolve across Inner Melbourne suburbs over time and by assessing the extent to which historical rental patterns can inform expectations about future price movements. By focusing on rental dynamics at a fine geographic scale, the analysis highlights both the common trends shared across the inner-city rental market and the substantial heterogeneity that exists between individual suburbs. In doing so, the paper provides a forward-looking perspective on rental price behaviour that is directly relevant to affordability monitoring and medium-term housing policy considerations.

4. Literature review

Melbourne’s rental market has attracted growing attention due to persistent affordability pressures and substantial spatial variation in rental prices. Existing research consistently highlights that rental outcomes are strongly shaped by location-specific characteristics, reinforcing the importance of accounting for regional heterogeneity when analysing rent dynamics.

Palm, Raynor, and Warren-Myers (Palm, Raynor, & Warren-Myers, 2021) examine rental affordability in Melbourne through the lens of housing stock characteristics and price filtering. Their findings challenge the assumption that increased housing supply necessarily improves affordability, particularly in inner-city areas where older dwellings may continue to command high rents due to amenity access and neighbourhood desirability. This evidence suggests that rental prices are deeply embedded in local housing market structures and that unobserved regional characteristics play a central role in shaping rent levels.

The importance of regional dynamics is further supported by Otto and Stapledon (Otto & Stapledon, 2017), who analyse rent growth and rental yields across Sydney and Melbourne. Using long-horizon regression techniques, they find that rent growth in Melbourne exhibits limited and uneven predictability compared to Sydney, indicating that rental dynamics vary substantially across locations and time. Their results imply that rental prices are not driven by uniform city-wide relationships alone but instead reflect local market conditions that persist over time.

Recent data-driven studies reinforce the relevance of spatial heterogeneity in Melbourne’s housing market. Phan (Phan, 2018) applies machine learning techniques to predict Melbourne housing prices and documents large disparities between inner and outer suburbs, with location emerging as a key determinant of predictive performance. Similarly, Nahar and Yun (Nahar & Yun, 2023) employ advanced machine learning models to forecast housing prices, highlighting the complex and non-uniform nature of Melbourne’s housing market. Although these studies focus on sale prices rather than rents, their findings underscore the importance of modelling location explicitly.

Finally, Jafary et al. (Jafary, Shojaei, Pishgahi, Rajabifard, & Ngo, 2024) incorporate rental outcomes into a data-driven framework aimed at improving housing affordability in Melbourne. Their results show that rental prices and yields vary systematically across suburbs, particularly within inner metropolitan areas. Collectively, this literature supports modelling Melbourne rental prices using regional fixed effects alongside time-based variation.

5. Context and Data

This paper draws on a single primary data source: quarterly rental market reports published by the Victorian Government (Victorian State Government Department of Families, Fairness and Housing, 2025). The dataset provides median rental prices for one-bedroom apartments at the Statistical Area 2 (SA2) level across Inner Melbourne suburbs, covering the period from 2007Q1 to 2025Q1. This historical series is used to estimate time-series models and generate rental price forecasts through to 2030Q1.

5.1 Data Cleaning

Here we collected median rental values for each of the 22 areas[2] for each quarter from 2007 Quarter 1 to 2025 Quarter 1. For periods from 2016 onwards, rental information is reported by lease commencement period rather than as a single quarterly median. To construct comparable quarterly series, weighted averages are calculated by first computing weighted sums of rents across lease commencement categories and then dividing by the total number of observations within each quarter. For earlier years (2007–2016), rental reports provide only a single quarterly median and total observation counts, with no further disaggregation.

Several data limitations arise from reporting inconsistencies in the rental series. No usable rental data are available for 2013, and the 2014 rental reports do not contain valid information for the regions considered in this study. Additionally, 2007Q2 rental values are excluded because the source material is mislabelled and duplicates 2007Q4 data. In 2007Q1, regional definitions differ slightly, with Southbank and Docklands reported jointly; these are subsequently separated to align with the regional structure used in later years. Finally, with the exception of a small number of quarters (2011Q1, 2011Q3, 2015Q1, 2015Q3, 2016Q1, and 2016Q2), rental data prior to 2016 do not consistently report three-bedroom apartment rents, and the analysis therefore focuses on one-bedroom dwellings.

5.2 Geographic Alignment and Aggregation

Several SA2 regions used in the rental data do not map directly to a single census SA2. In these cases, proxy regions or weighted aggregations are employed to maintain consistent spatial units over time. For Carlton–Parkville, Albert Park–Middle Park–West St Kilda, Collingwood–Abbotsford, Fitzroy North–

[2] Areas can be found in the Appendix.

Clifton Hill, Flemington–Kensington, and North Melbourne–West Melbourne, weighted averages of census variables are constructed using population counts.

Where direct SA2 matches are unavailable, the closest geographic substitutes are used. For example, West St Kilda is not available as a standalone SA2, and St Kilda is used as a proxy. Richmond–Burnley does not exist as an SA2, so Richmond (South)–Cremorne is used instead, while Clifton Hill–Alphington is used for Clifton Hill. Because the analysis relies primarily on proportional variables rather than raw counts, these substitutions are unlikely to materially distort relative demographic comparisons across regions.

In the 2021 Census, South Yarra is divided into South Yarra North and South Yarra South, requiring a different aggregation approach compared to earlier years. Additionally, some overlap arises where St Kilda data are used both directly and as part of aggregated regions, implying limited double-counting in descriptive population measures. These issues are documented to maintain transparency but do not affect the rental fixed-effects estimation, which relies solely on rental price data.

6. Empirical Methodology

This paper applies a time-series forecasting approach to analyse rental price dynamics across Inner Melbourne suburbs. Four models of increasing flexibility are estimated separately for each suburb using quarterly rental data: a Seasonal Naïve benchmark, a Seasonal Regression model with a linear time trend, a Seasonal Autoregressive Integrated Moving Average (SARIMA) model, and an Exponential Smoothing (ETS) model. All specifications rely solely on time, seasonality, and past rental prices. Model performance is evaluated using rolling one-step-ahead out-of-sample forecasts and standard accuracy measures, including RMSE, MAE, and MAPE. Based on aggregate and suburb-level forecasting performance, the best model is selected and used to generate rental price projections through to 2030, with prediction intervals used to quantify forecast uncertainty.

6.1 Model 1: Seasonal Naïve

The seasonal naïve model is a benchmark forecasting approach in which the rental price for a given quarter is predicted using the observed rental price from the same quarter in the previous year. Formally, the forecast for quarter t is given by

$$\hat{y}_t = y_{t-4},$$

where the data are measured quarterly. The model contains no estimated parameters and relies solely on the persistence and seasonal structure of the rental price series.

Forecast accuracy is evaluated using a one-year (four-quarter) hold-out sample for each suburb. The model is estimated using data up to the start of the hold-out period, and forecasts are generated for the subsequent four quarters.

The seasonal naïve model is used as a baseline benchmark against which the upcoming more complex forecasting models are evaluated. Rental prices are known to exhibit strong persistence and seasonal patterns due to lease structures, academic calendars, and regular market cycles. As a result, a simple seasonal benchmark often performs surprisingly well in practice.

Figure 1: Seasonal Naïve: Actual vs Forecast — Carlton North

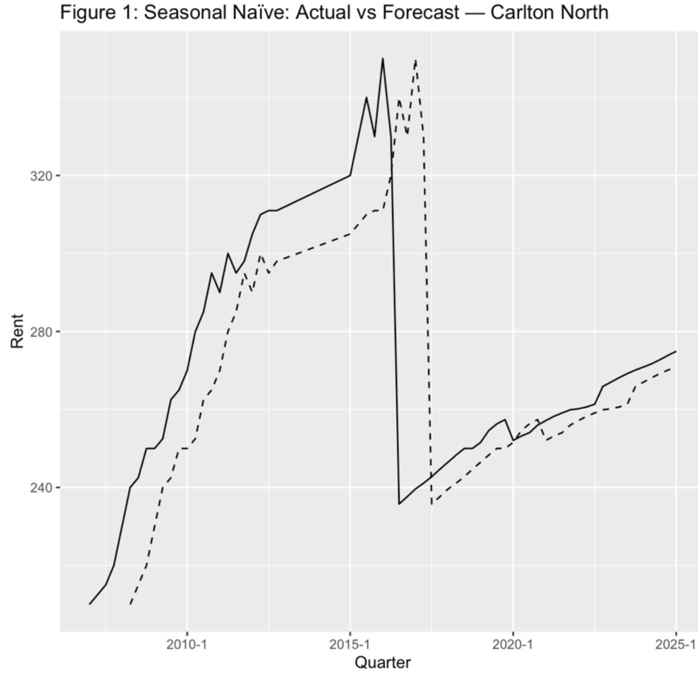

Including a seasonal naïve specification ensures that any improvement observed from more advanced models reflects genuine predictive gains rather than overfitting. If a more complex model fails to outperform this benchmark, it cannot be considered to add meaningful forecasting value.

Figure 1 compares the observed quarterly rental prices in Carlton North (solid line) with forecasts generated by the seasonal naïve model (dashed line). The seasonal naïve forecast for each quarter is equal to the observed rental price from the same quarter in the previous year. As shown, the model closely tracks the overall level and seasonal pattern of rents, reflecting the high persistence in rental prices over time.

Periods where the dashed line lags behind the solid line correspond to episodes of rapid rent growth, which the seasonal naïve model cannot anticipate because it relies solely on historical seasonal repetition. Overall, the plot illustrates why the seasonal naïve specification provides a strong benchmark while also highlighting its limitation in adapting to structural changes in rental dynamics.

6.1.1 Results

The seasonal naïve model provides a strong benchmark for forecasting quarterly rental prices across inner Melbourne suburbs. The model achieves mean absolute percentage errors ranging from approximately 2.8% to 7.1% across suburbs (Table 1). Lower forecast errors are observed in high-density inner-city locations such as the CBD–St Kilda Road precinct and Southbank, where rental prices exhibit relatively stable seasonal patterns.

Suburbs with higher forecast errors, including Fitzroy, Carlton North, and Richmond–Burnley, tend to experience more pronounced rent growth and short-term volatility, which the seasonal naïve model cannot anticipate due to its reliance on fixed seasonal repetition.

Aggregating across all 22 suburbs, the seasonal naïve model achieves a mean RMSE of 8.78, a mean MAE of 8.62, and a mean MAPE of 2.58% over the one-year hold-out period. These results suggest that, on average, the model's forecasts deviate from actual rental prices by less than 3% of the rent level, highlighting a high degree of stability in inner Melbourne rental markets.

However, variation in forecast accuracy across suburbs indicates that the seasonal naïve model may struggle to fully capture periods of accelerated rent growth or structural change. This motivates the consideration of more flexible models that allow for trends and dynamic adjustments while being evaluated against this strong benchmark.

6.2 Model 2: Seasonal Regression

The seasonal regression model represents quarterly rental prices as the sum of a long-run trend and recurring seasonal effects. Formally, for each suburb, quarterly rental prices are modelled as a linear time trend plus quarter-specific seasonal effects:

$$\text{Rent}_t = \beta_0 + \beta_1 t \; + \; \gamma_2 D_{2,t} + \gamma_3 D_{3,t} + \gamma_4 D_{4,t} + \varepsilon_t$$

where:

- Rent_t is the observed rent in quarter t
- t is a time index increasing by 1 each quarter (captures the long-run trend)
- $D_{2,t}, D_{3,t}, D_{4,t}$ are dummy variables for quarters Q2, Q3, and Q4
- β_0 is the baseline rent level in Q1 at $t = 0$(intercept) which is equivalent to 2007Q1
- β_1 is the average change in rent per quarter (trend slope)
- $\gamma_2, \gamma_3, \gamma_4$ measure how rents in Q2–Q4 differ from Q1, holding the trend constant
- ε_t is the error term capturing unexplained deviations

The time trend captures gradual structural changes in rental levels over the sample period, while the quarterly indicators allow rents to systematically differ across quarters within a year.

Each quarter is treated as a categorical variable, ensuring that seasonal effects are estimated flexibly rather than imposed mechanically. This allows the model to account for predictable within-year patterns in rental prices that may arise from lease cycles, academic calendars, and regular demand fluctuations. The inclusion of a time trend ensures that these seasonal effects are estimated around an evolving rental level rather than a fixed mean.

The model is estimated separately for each suburb using historical rental data and evaluated using rolling one-step-ahead forecasts. At each point in time, parameter estimates are obtained using only information available up to the previous quarter, and forecasts are generated for the subsequent quarter. This approach ensures that forecast evaluation is strictly out-of-sample and reflects realistic predictive conditions.

The seasonal regression model is included to assess whether explicitly modelling trend and seasonality improves the ability to forecast rental prices. Unlike purely heuristic approaches, this specification imposes a structured relationship between time, seasonality, and rental levels, allowing rents to evolve smoothly over time while accommodating recurring quarterly patterns. As such, the model provides an interpretable and parsimonious framework for capturing predictable components of rental dynamics at the regional level.

Figure 2: Seasonal Regression: Actual vs Forecast — Carlton North

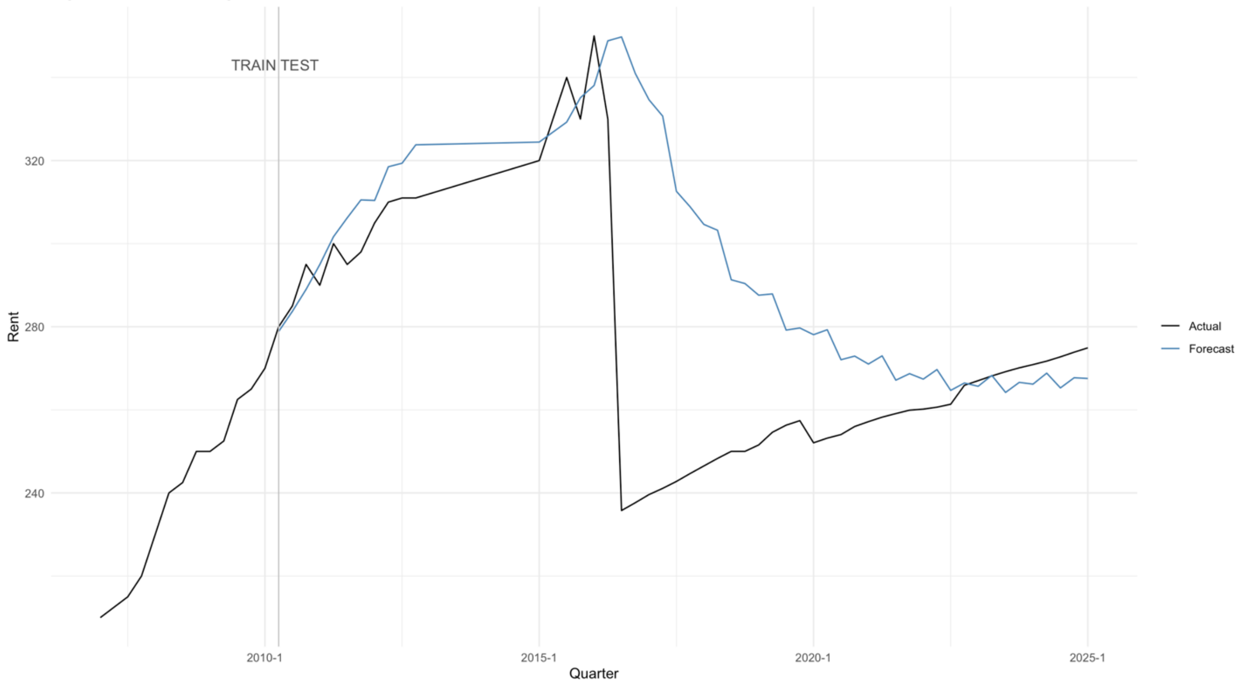

Figure 2 illustrates how the seasonal regression model decomposes rental prices into a long-run trend and recurring seasonal effects using the example of Carlton North. The smooth evolution of the forecast line reflects the role of the linear time trend in capturing gradual changes in rental levels over time, while deviations between quarters are governed by the estimated quarterly dummy variables. During the training period, the model learns a stable relationship between time, seasonality, and rents, which is then extrapolated forward in the test period using rolling one-step-ahead forecasts. Because the model imposes a single linear trend and fixed seasonal effects, predicted rents evolve smoothly and adjust only gradually when underlying rental dynamics change. The divergence observed in the test period highlights the model's reliance on historically estimated trend and seasonality, demonstrating how forecasts are driven by structural components rather than short-term fluctuations.

6.2.1 Results

The seasonal regression model exhibits substantial variation in forecast accuracy across inner Melbourne suburbs. Mean absolute percentage errors range from approximately 4% in centrally located, high-density areas such as the CBD–St Kilda Road precinct and Southbank to over 9% in suburbs including Carlton

North and Fitzroy (as seen in Table 2). Across all suburbs, the model produces an average RMSE of 28.6, an average MAE of 19.8, and a mean MAPE of 6.81%, based on 51 rolling out-of-sample forecasts per suburb.

These results indicate that while the model captures broad trends and seasonal patterns in rental prices, its performance differs markedly across locations, reflecting heterogeneity in rental dynamics and the presence of localised fluctuations that are not fully explained by a single linear trend and fixed seasonal effects. The dispersion in forecast errors suggests that rental price movements in some suburbs deviate considerably from smooth trend-seasonal behaviour, motivating the consideration of more flexible time-series specifications.

6.3 Model 3: Seasonal ARIMA

The Seasonal Autoregressive Integrated Moving Average (SARIMA) model is included to capture dynamic temporal dependencies in rental prices that extend beyond fixed trend and seasonal patterns. Unlike regression-based approaches that impose a deterministic trend and constant seasonal effects, SARIMA allows rental prices to evolve according to their own past values and past shocks, while explicitly accounting for quarterly seasonality making the model well suited to quarterly rental data.

The SARIMA specification accommodates both short-run dynamics and recurring seasonal behaviour by modelling autoregressive and moving average components at both the quarterly and annual frequency. This flexibility allows the model to adapt to changing rental conditions over time without imposing a fixed functional form on the underlying trend.

The model is estimated separately for each suburb using the quarterly rental data. For each series, an initial training window of 12 quarters is used to select the model specification via automatic model selection. Forecasts are then generated using a rolling one-step-ahead procedure, where the model is updated sequentially as new observations become available. This approach ensures that all forecast evaluation is strictly out-of-sample and reflects realistic forecasting conditions.

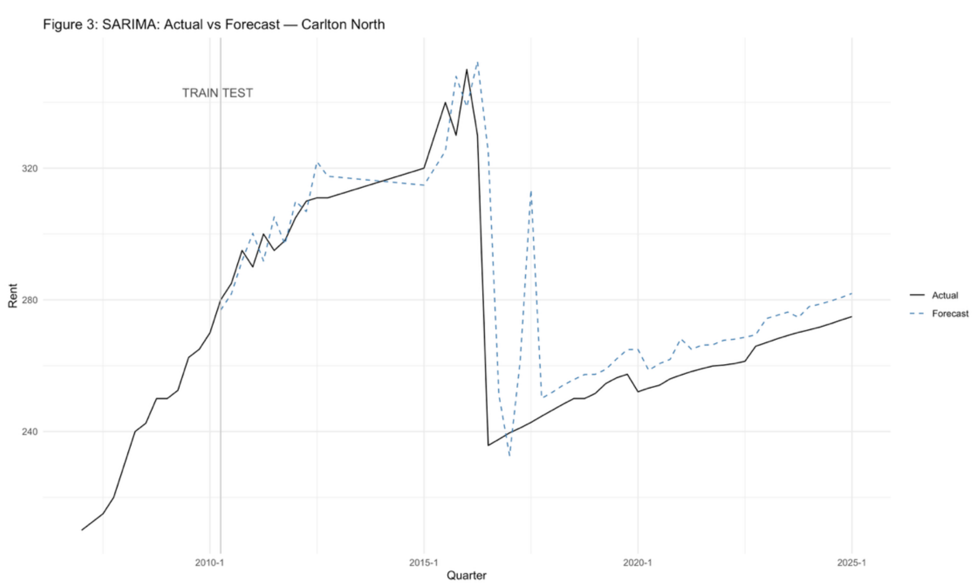

Figure 3 presents observed and forecast rental prices for Carlton North as an illustrative example of the SARIMA model's behaviour. The forecast series closely follows the underlying trajectory of rental prices, reflecting the model's ability to capture persistent movements and seasonal regularities in the data. The predicted path evolves smoothly over time, indicating that forecasts are driven by accumulated information from past observations rather than abrupt quarter-to-quarter changes.

Following the structural break in rental levels, forecasts adjust gradually as new data are incorporated, highlighting the model's reliance on learned temporal dynamics. This behaviour is consistent with the SARIMA framework, which responds to changes through updated autoregressive and moving average components rather than imposing instantaneous level shifts. Figure 3 therefore illustrates how the SARIMA model generates forecasts based on the internal dynamics of the rental series while maintaining quarterly seasonal structure.

6.3.1 Results

The SARIMA model also exhibits notable regional variation in forecast accuracy across inner Melbourne suburbs. Forecast errors are lowest in several centrally located, high-density areas, including Docklands, the CBD–St Kilda Road precinct, and Southbank, where mean absolute percentage errors are generally below 2% (as shown in Table 3). In contrast, higher forecast errors are observed in suburbs such as Fitzroy, Carlton–Parkville, and South Melbourne, where MAPE values exceed 6%, indicating greater volatility and less predictable rental dynamics in these locations.

Across all suburbs, the SARIMA model achieves an average RMSE of 16.3, an average MAE of 9.53, and a mean MAPE of 3.23%**,** based on 51 rolling one-step-ahead forecasts per suburb**.** These results suggest that while the SARIMA specification effectively captures persistent temporal dependence and quarterly seasonality in rental prices, its forecasting performance varies across regions. The observed dispersion in errors reflects heterogeneity in local rental market conditions, with some suburbs exhibiting more stable and predictable dynamics than others. This spatial variation highlights the importance of allowing rental prices to evolve according to suburb-specific time-series behaviour rather than assuming uniform dynamics across locations.

6.4 Model 4: ETS Model

The exponential smoothing (ETS) model is included to capture evolving trends and seasonal patterns in rental prices without imposing fixed coefficients or deterministic structures. Unlike regression-based approaches, ETS models allow the level and seasonal components of the series to update gradually over time, placing greater weight on more recent observations. This makes the model particularly suitable for rental markets where underlying conditions may change slowly in response to shifting demand, policy settings, or broader economic conditions. By modelling rental prices within a state-space framework, ETS provides a flexible yet parsimonious representation of rental dynamics that adapts naturally as new information becomes available.

Figure 4: Exponential Smoothing (ETS): Actual vs Forecast — Carlton North

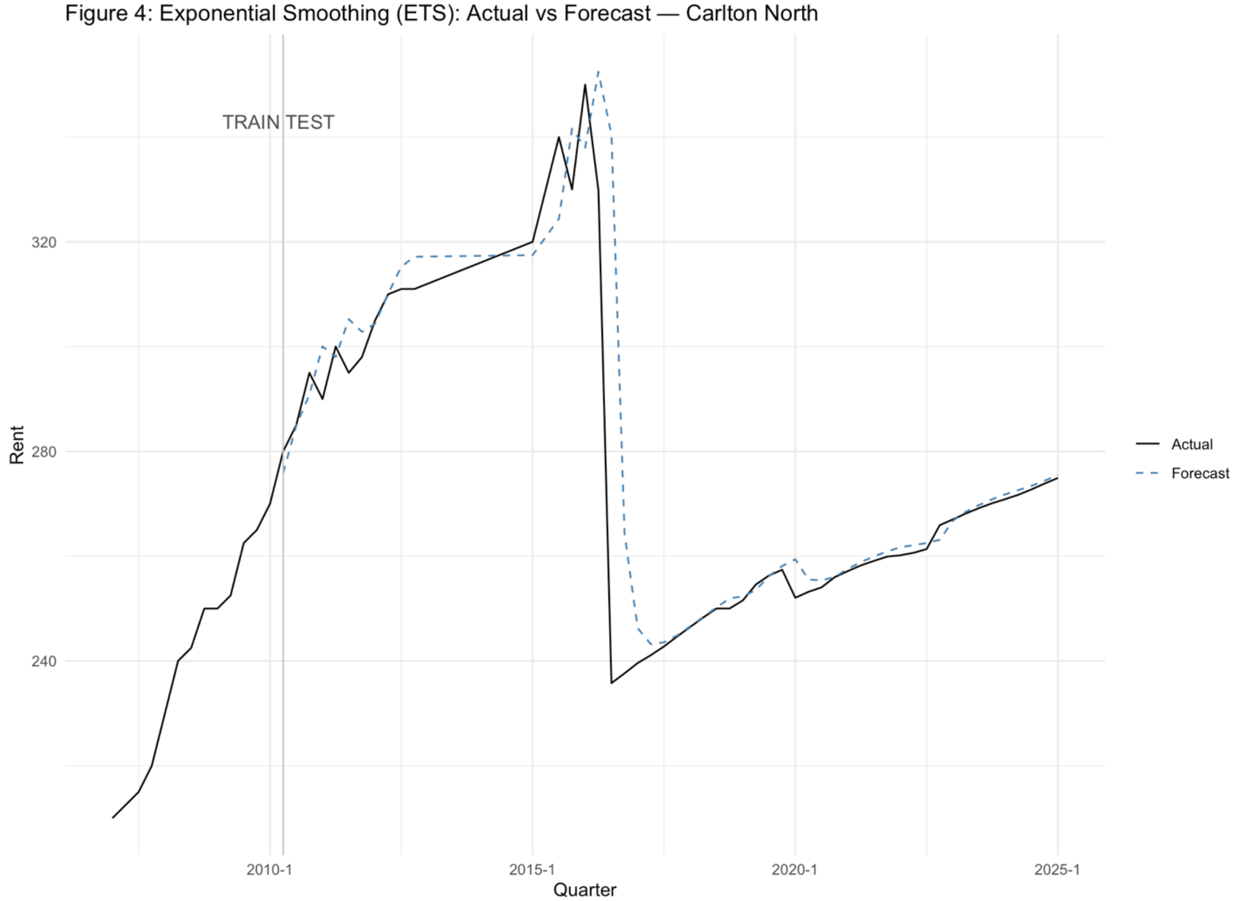

Figure 4 illustrates the behaviour of the ETS model for Carlton North by plotting observed rental prices against rolling one-step-ahead forecasts. The vertical line separates the training and test periods, after which all forecasts are generated using only information available up to the previous quarter.

The forecast path closely tracks observed rental prices throughout most of the sample, reflecting the ETS model's ability to adapt to evolving rental dynamics through gradual updates to the level and seasonal components. Because exponential smoothing places greater weight on recent observations, forecasts adjust smoothly as rental conditions change rather than remaining fixed around a long-run average. Following the sharp decline in rents, predicted values adjust progressively as new data are incorporated, highlighting the model's reliance on incremental updating rather than abrupt regime shifts. Overall, Figure 4 demonstrates how the ETS framework captures persistent movements and stable quarterly seasonality while allowing rental levels to evolve over time.

6.4.1 Results

Across all inner Melbourne suburbs, the ETS model demonstrates strong forecasting performance under rolling one-step-ahead evaluation. On average, the model achieves a root mean squared error (RMSE) of 17.6, a mean absolute error (MAE) of 9.84, and a mean absolute percentage error (MAPE) of 3.32%, based on 51 out-of-sample forecasts per suburb. Forecast accuracy varies across locations, with lower errors observed in several high-density inner-city suburbs and higher errors in areas exhibiting greater volatility in rental movements. These results indicate that the ETS model effectively captures persistent trends and evolving seasonal patterns in quarterly rental prices while allowing for regional heterogeneity in rental dynamics.

The ETS results reveal clear regional heterogeneity in rental price dynamics across inner Melbourne suburbs (Table 4). Forecast accuracy is highest in several centrally located, high-density areas such as Carlton–Parkville, Docklands, and Port Melbourne, where rental prices appear to evolve smoothly and exhibit stable seasonal patterns. In contrast, larger forecast errors are observed in suburbs including South Yarra, Armadale, and East St Kilda, suggesting greater volatility and less predictable short-term movements in these markets. This spatial variation indicates that rental dynamics differ meaningfully across locations, reflecting differences in housing stock composition, demand pressures, and local market conditions. The results underscore the importance of modelling rental prices at a granular regional level, as aggregate or uniform specifications may mask substantial localised variation in rental behaviour.

6.5 Models Analysis

Table 5: Forecast Accuracy Metrics by Model

Model	Mean RMSE	Mean MAE	Mean MAPE (%)
Seasonal Naïve	8.78	8.62	2.58
Seasonal Regression	28.60	19.80	6.81
SARIMA	16.30	9.53	3.23
Exponential Smoothing (ETS)	17.60	9.84	3.32

6.5.1 Aggregate model performance

Aggregate forecast accuracy across models is summarised in Table 5, which reports mean RMSE, MAE, and MAPE values averaged across all inner Melbourne suburbs. The results indicate substantial differences in predictive performance depending on model specification. The seasonal naïve model performs well as a benchmark, reflecting the strong and persistent quarterly seasonality present in rental prices. In contrast, the seasonal regression model exhibits considerably higher forecast errors, suggesting that a specification based on a single linear trend with fixed seasonal effects is insufficient to capture the complexity of rental price dynamics over time. Models that allow rental prices to evolve dynamically, namely SARIMA and exponential smoothing, achieve substantially lower aggregate errors. This highlights the importance of modelling temporal dependence and adaptive behaviour in quarterly rental series. Among these, the SARIMA model records the lowest average RMSE and MAE, while the ETS model delivers comparably strong performance with marginally higher aggregate errors.

6.5.2 Suburb-level performance

At the suburb level, forecast accuracy varies markedly across locations for all models, underscoring pronounced regional heterogeneity in rental market dynamics. Detailed suburb-level results are reported in Tables 1, 2, 3 and 4. Across models, centrally located and high-density suburbs such as Docklands, the CBD–St Kilda Road, and Southbank tend to exhibit lower forecast errors, indicating smoother and more predictable rental price evolution. In contrast, higher errors are observed in suburbs including Fitzroy, Carlton North, and South Yarra, where rental prices appear more volatile and subject to localised fluctuations. These patterns suggest that rental dynamics differ meaningfully across inner Melbourne, reflecting variation in housing stock composition, demand pressures, and neighbourhood-specific factors. Models that permit greater flexibility in capturing temporal dependence, such as SARIMA and ETS, display more stable performance across suburbs than specifications that impose fixed trend and seasonal structures.

6.6 Final Model Selection

Based on forecast accuracy across suburbs and in aggregate, as well as model interpretability, the SARIMA model is selected as the final forecasting specification.

The SARIMA framework offers a flexible yet structured representation of rental price dynamics by explicitly modelling persistence, seasonal autocorrelation, and stochastic trends. It consistently delivers strong forecast performance across a wide range of suburbs while remaining transparent and well-established within the time-series literature. Importantly, SARIMA adapts to evolving rental conditions through its autoregressive and moving average components without imposing fixed deterministic structures.

Accordingly, the SARIMA model is used to generate final rental price forecasts for future quarters.

7. SARIMA Predictive Analysis (2026- 2030)

The SARIMA model was estimated separately for each suburb using quarterly rental data from 2007Q1 to 2025Q1. After model selection via the automated SARIMA procedures took place, the final fitted models were used to generate forecasts from 2025Q2 through 2030Q4. Forecasts were produced on a quarterly basis and include 95% prediction intervals, which reflect both parameter uncertainty and the accumulation of forecast error over time.

To obtain an aggregate forecast for Inner Melbourne, suburb-level SARIMA forecasts were first generated individually and then **a**veraged across all suburbs at each forecast horizon. This ensures that the aggregate projection reflects the underlying heterogeneity in suburban rental dynamics rather than imposing a single pooled time-series model. The same approach was applied to the historical data, where observed rents were averaged across suburbs up to 2025Q1, allowing a direct comparison between historical trends and projected outcomes.

7.1 Average Inner Melbourne rental forecast

Figure 5 (Average Inner Melbourne Rent: Actual and SARIMA Forecast) presents the historical average rent across Inner Melbourne suburbs alongside the SARIMA forecast and its 95% prediction interval.

Figure 5: Average Inner Melbourne Rent: Actual and SARIMA Forecast (95% Interval)

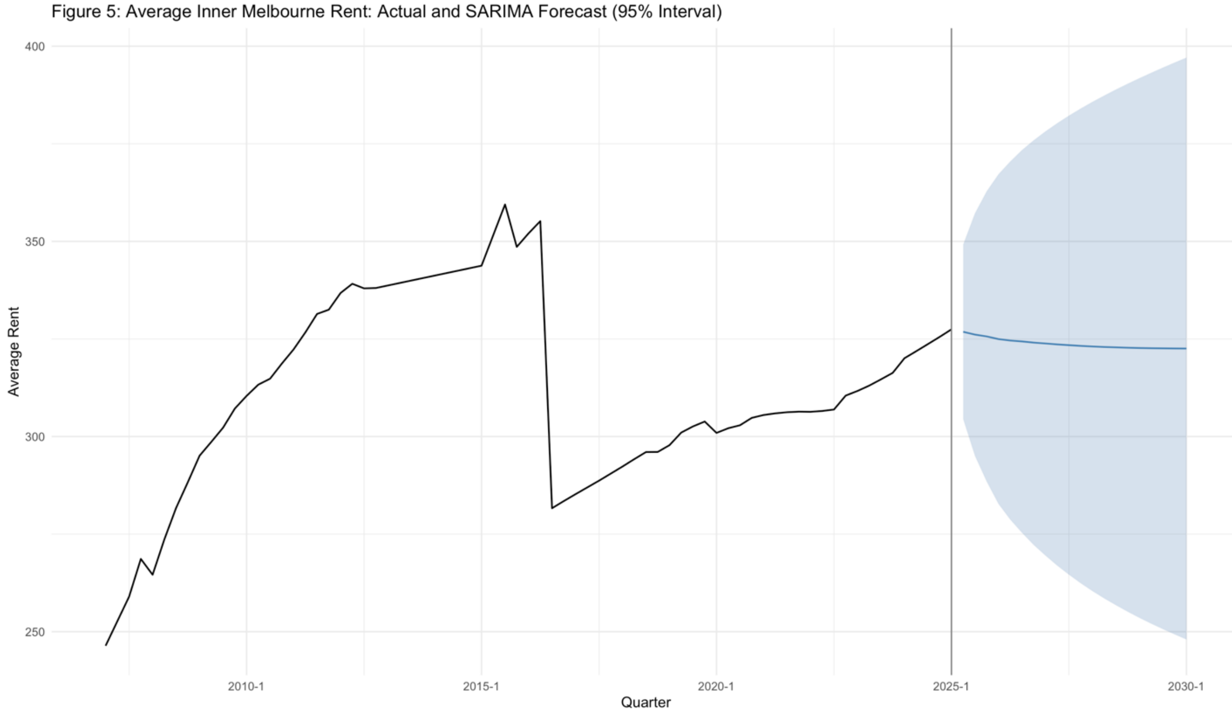

The historical series shows a strong upward trajectory from 2007 through 2016, followed by a sharp structural decline and a gradual recovery thereafter. The SARIMA forecast extends this into the projection period, with the point forecast indicating a relatively stable to mildly decreasing average rental level through 2030. However, the widening prediction band highlights increasing uncertainty at longer horizons.

By 2030, the forecast interval spans a broad range, indicating that while the model captures the central tendency of rental movements, future rental outcomes are sensitive to shocks and structural changes not explicitly modelled.

7.2 Carlton North forecast

Figure 6 focuses on suburb-level dynamics for Carlton North.

Figure 6: Carlton North: SARIMA Forecast with 95% Prediction Interval

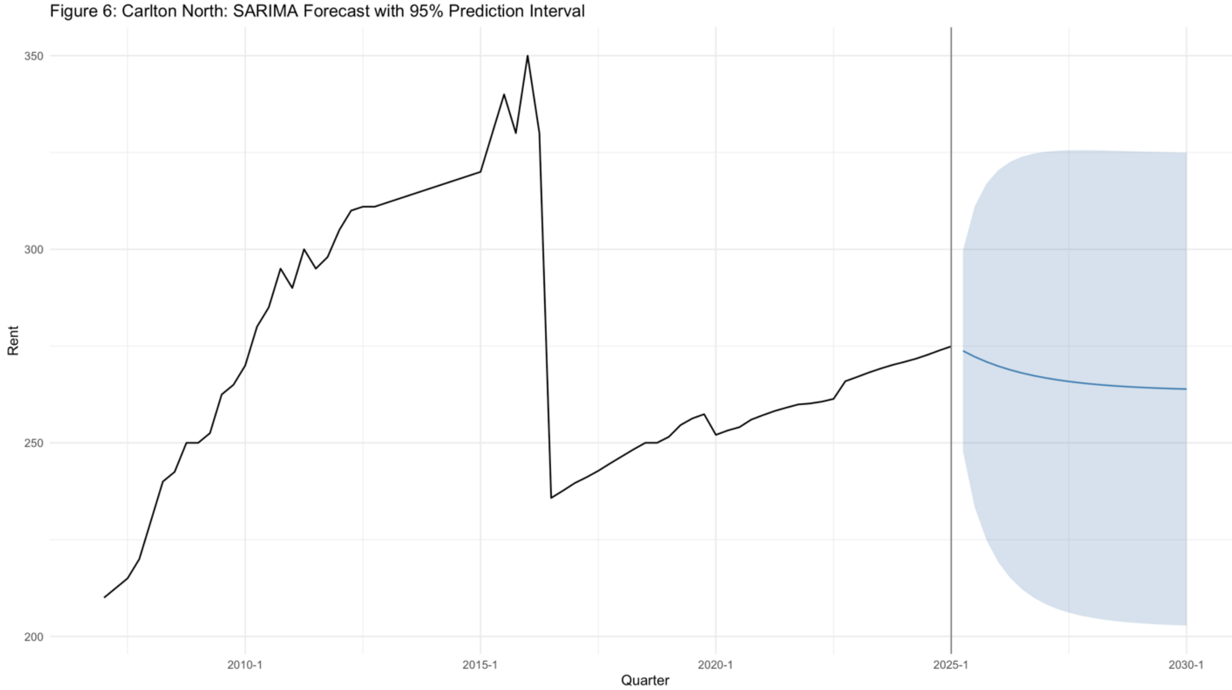

The historical pattern mirrors the broader Inner Melbourne trend but with greater volatility, particularly around the mid-2010s peak and subsequent correction. The SARIMA model projects a modest decline followed by stabilisation in the point forecast for Carlton North rents between 2025 and 2030.

The prediction interval for Carlton North is notably wide relative to the point forecast, especially at longer horizons. This reflects higher historical variability in this suburb's rental series and indicates substantial uncertainty around future rental paths. While the central forecast suggests relative stability, the lower and upper bounds imply that both sustained softness and renewed upward pressure remain plausible over the medium term.

Results

Taken together, the SARIMA forecasts provide a structured, data-driven projection of rental dynamics under historical patterns of trend, seasonality, and persistence. The point forecasts summarise the most likely rental path conditional on past behaviour, while the prediction intervals quantify uncertainty and highlight the limits of long-horizon forecasting.

At the aggregate level, Inner Melbourne rents are projected to stabilise with moderate decline potential, whereas suburb-specific forecasts such as Carlton North reveal greater dispersion and sensitivity to localised dynamics. These results underscore the importance of combining aggregate indicators with suburb-level analysis when assessing future rental market conditions.

Overall, the SARIMA framework offers a flexible and empirically grounded approach to medium-term rental forecasting, while the widening forecast intervals appropriately caution against over-interpretation of long-run point estimates.

8. Limitations

Several limitations should be acknowledged when interpreting the results of this study. First, all models rely exclusively on historical rental price data and therefore assume that past patterns of trend, seasonality, and persistence continue into the future. Structural changes in the housing market such as shifts in migration, housing supply, regulatory interventions, or macroeconomic shocks are not explicitly modelled and may lead to forecast deviations.

Second, while SARIMA and related time-series models effectively capture temporal dependence and seasonal structure, they do not incorporate exogenous explanatory variables. Extreme changes in factors such as income growth, interest rates, housing construction, or demographic change may play an important role in shaping rental dynamics but are omitted from the forecasting framework. As a result, forecasts should be interpreted as conditional on historical rental behaviour rather than as causal predictions.

Third, forecast uncertainty increases substantially over longer horizons, as reflected in the widening prediction intervals. Although these intervals appropriately account for stochastic variation, they may understate uncertainty associated with rare but impactful events. This limitation is particularly relevant for suburb-level forecasts, where localised shocks can have outsized effects.

Finally, averaging forecasts across suburbs masks heterogeneity in rental dynamics. While the aggregate Inner Melbourne forecast provides a useful summary indicator, individual suburbs may experience divergent trends that are not fully captured by the average trajectory.

9. Conclusion

This paper analysed quarterly rental price dynamics across Inner Melbourne suburbs using a sequence of increasingly flexible time-series models. After data collection and cleaning, multiple benchmark and structured models were estimated and evaluated using rolling out-of-sample forecasts, with performance assessed via RMSE, MAE, and MAPE.

Among the models considered, SARIMA was selected for final projection due to its ability to jointly model trend, seasonality, and serial dependence while delivering strong and stable forecast performance across suburbs. Using SARIMA, rental prices were forecast through 2030 at both the aggregate Inner Melbourne level and for individual suburbs such as Carlton North, with prediction intervals highlighting increasing uncertainty over time.

Overall, the results indicate a stabilisation of rental prices at the aggregate level, alongside continued heterogeneity across suburbs. The modelling framework provides a transparent and empirically grounded approach to medium-term rental forecasting, while emphasising the importance of uncertainty and local variation in interpreting future rental market outcomes.

10. Bibliography

Australian burea of Statistics. (2025). *Census All persons QuickStats*. Retrieved from Australian burea of Statistics: https://www.abs.gov.au/census/find-census-data/quickstats/2021/206051512

Australian Bureau of Statistics. (2021, July 20). *Statistical Area Level 2*. Retrieved from Australian Bureau of Statistics: https://www.abs.gov.au/statistics/standards/australian-statistical-geography-standard-asgs-edition-3/jul2021-jun2026/main-structure-and-greater-capital-city-statistical-areas/statistical-area-level-2

Jafary, P., Shojaei, D., Pishgahi, S., Rajabifard, A., & Ngo, T. (2024). Data-driven strategies for affordable housing: A hybrid genetic algorithm–machine learning optimization model in the Melbourne metropolitan area. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 175–182.

Lam, L. (2025, June 11). *Average Australian home passes A$1m amid housing crisis*. Retrieved from BBC News: https://www.bbc.com/news/articles/c9vg923vkdko

Nahar, J., & Yun, C. J. (2023). A Data-Driven Approach to Predicting Melbourne Housing Prices Using Advanced Machine Learning Models. *Proceedings of the International Conference on Smart Technologies for Energy, Environment and Resources (I-COSTE 2023)*.

Otto, G., & Stapledon, N. (2017). How Predictable? Rent Growth and Returns in Sydney and Melbourne Housing Markets. *SSRN Electronic Journal* .

Palm, M., Raynor, K., & Warren-Myers, G. (2021). Examining building age, rental housing and price filtering for affordability in Melbourne, Australia. *Urban Studies*, 2881-2900.

Phan, T. D. (2018). Housing Price Prediction Using Machine Learning Algorithms: The Case of Melbourne City, Australia. *Proceedings of the 2018 International Conference on Machine Learning and Data Engineering (iCMLDE)*, 117-122.

Victorian State Government Department of Families, Fairness and Housing. (2025, November 14). *Past Rental Reports*. Retrieved from Past Rental Reports: https://www.dffh.vic.gov.au/publications/past-rental-reports

11. Appendix

a) SA2 Level regions used in analysis:

Albert Park-Middle Park-West St Kilda
Armadale
Carlton North
Carlton-Parkville
CBD-St Kilda Rd
Collingwood-Abbotsford
Docklands
East Melbourne
East St Kilda
Elwood
Fitzroy
Fitzroy North-Clifton Hill
Flemington-Kensington
North Melbourne-West Melbourne
Port Melbourne
Prahran-Windsor
Richmond-Burnley
South Melbourne
South Yarra
Southbank
St Kilda
Toorak

Table 1: Seasonal Naïve Forecast Performance by Suburb

Rank	Suburb	RMSE	MAE	MAPE (%)
1	CBD–St Kilda Rd	14.35	10.20	2.87
2	Southbank	15.70	11.61	2.88
3	Carlton–Parkville	16.38	11.64	3.87
4	Docklands	16.69	12.27	3.00
5	Port Melbourne	17.90	10.80	2.82
6	South Melbourne	19.89	12.80	3.57
7	Prahran–Windsor	20.38	13.55	4.87
8	Albert Park–Middle Park–West St Kilda	21.24	12.10	4.36
9	Elwood	21.38	11.86	4.55
10	Toorak	22.18	12.44	4.84
11	St Kilda	23.22	13.19	4.67
12	East St Kilda	23.24	13.73	5.69
13	South Yarra	24.16	15.51	5.16
14	Armadale	24.29	14.82	5.84
15	North Melbourne–West Melbourne	24.80	18.10	6.13
16	East Melbourne	25.99	15.39	4.84
17	Fitzroy North–Clifton Hill	27.34	17.21	6.58
18	Flemington–Kensington	28.36	17.49	6.71
19	Collingwood–Abbotsford	28.42	20.00	6.33
20	Carlton North	29.52	17.29	6.61
21	Richmond–Burnley	29.69	18.31	6.18
22	Fitzroy	34.22	21.23	7.08

Notes: Errors are computed using rolling seasonal naïve forecasts, where each quarter's rent is predicted using the observed rent from the same quarter in the previous year. Forecast accuracy is evaluated over 59 out-of-sample quarters per suburb.

Table 2: Seasonal Regression Forecast Performance by Suburb

Rank	Suburb	RMSE	MAE	MAPE (%)
1	CBD–St Kilda Rd	19.67	15.11	4.21
2	Southbank	20.17	16.18	4.02
3	Docklands	21.11	17.37	4.23
4	Prahran–Windsor	21.61	14.31	5.22
5	Carlton–Parkville	22.25	19.32	6.23
6	North Melbourne–West Melbourne	22.50	18.26	6.11
7	Port Melbourne	23.98	16.42	4.33
8	Elwood	27.09	17.78	6.98
9	Collingwood–Abbotsford	28.21	21.54	6.73
10	St Kilda	28.41	17.74	6.48
11	South Melbourne	28.77	21.91	6.05
12	Albert Park–Middle Park–West St Kilda	28.87	19.45	7.24
13	Toorak	29.27	19.80	7.86
14	East St Kilda	30.55	19.70	8.37
15	East Melbourne	31.82	21.04	6.82
16	South Yarra	31.86	20.26	6.82
17	Armadale	32.53	21.12	8.39
18	Flemington–Kensington	33.37	22.20	8.62
19	Fitzroy North–Clifton Hill	33.89	21.66	8.46
20	Carlton North	36.28	22.78	9.03
21	Richmond–Burnley	37.08	24.28	8.33
22	Fitzroy	38.97	26.79	9.23

Notes: Errors are computed using rolling one-step-ahead forecasts from the seasonal regression model, where quarterly rents are predicted using a linear time trend and quarter-specific dummy variables. At each forecast origin, model parameters are estimated using only data available up to the previous quarter. Forecast accuracy is evaluated over 51 out-of-sample quarters per suburb.

Table 3: SARIMA Forecast Performance by Suburb

Rank	Suburb	RMSE	MAE	MAPE (%)
1	Docklands	6.28	3.04	0.74
2	CBD–St Kilda Rd	7.34	3.21	0.90
3	Southbank	9.33	6.56	1.63
4	Port Melbourne	10.53	5.00	1.28
5	East Melbourne	13.32	4.32	1.36
6	Armadale	13.54	6.04	2.35
7	Collingwood–Abbotsford	14.07	7.52	2.33
8	Richmond–Burnley	14.28	5.23	1.75
9	Toorak	14.33	7.95	3.10
10	North Melbourne–West Melbourne	14.47	9.21	3.04
11	East St Kilda	14.66	8.83	3.65
12	Prahran–Windsor	14.86	8.08	2.89
13	St Kilda	17.16	10.31	3.68
14	Carlton North	18.14	10.58	4.02
15	Fitzroy North–Clifton Hill	18.43	9.36	3.53
16	South Yarra	18.88	12.72	4.21
17	Fitzroy	19.24	8.18	2.63
18	Elwood	21.15	9.87	3.85
19	Albert Park–Middle Park–West St Kilda	21.29	14.52	5.34
20	Flemington–Kensington	25.05	11.83	4.68
21	South Melbourne	25.30	22.49	6.19
22	Carlton–Parkville	26.72	24.77	7.95

Notes: Errors are computed using rolling one-step-ahead forecasts from a Seasonal Autoregressive Integrated Moving Average (SARIMA) model. For each suburb, an initial SARIMA specification is selected using historical data and subsequently updated recursively as new observations become available. Forecast accuracy is evaluated over 51 out-of-sample quarters per suburb.

Table 4: Exponential Smoothing (ETS) Forecast Accuracy

Rank	Suburb	RMSE	MAE	MAPE (%)
1	Carlton–Parkville	5.43	2.64	0.85
2	Docklands	6.28	3.04	0.74
3	Port Melbourne	8.72	3.31	0.85
4	Albert Park–Middle Park–West St Kilda	10.81	3.64	1.32
5	Toorak	11.38	3.94	1.49
6	Elwood	12.02	3.88	1.46
7	East Melbourne	13.32	4.32	1.36
8	Richmond–Burnley	13.78	4.80	1.60
9	North Melbourne–West Melbourne	13.95	6.51	2.12
10	Collingwood–Abbotsford	14.00	6.78	2.10
11	Flemington–Kensington	14.47	5.97	2.25
12	Fitzroy North–Clifton Hill	16.04	5.90	2.17
13	Carlton North	16.07	5.56	2.09
14	CBD–St Kilda Rd	18.79	14.51	4.05
15	Southbank	20.01	16.08	3.99
16	Prahran–Windsor	20.65	13.65	4.98
17	Fitzroy	22.34	12.72	4.16
18	St Kilda	27.20	17.02	6.21
19	South Melbourne	27.46	20.95	5.79
20	East St Kilda	29.25	18.90	8.03
21	Armadale	31.17	20.26	8.05
22	South Yarra	33.16	22.14	7.44

Notes: Errors are computed using rolling one-step-ahead forecasts from an Exponential Smoothing (ETS) model, which allows for time-varying level, trend, and seasonal components. Model structure and smoothing parameters are re-estimated using data available up to each forecast origin. Forecast accuracy is evaluated over 51 out-of-sample quarters per suburb.