Why Homelessness Persists in Halifax

Much of the discussion around homelessness in the Halifax Regional Municipality follows a familiar pattern. The implicit message is that homelessness is largely the result of recent pressures in the market and that it increases when conditions worsen and should recede when they improve. Rising rents, insufficient income supports, population growth, and policy decisions are each presented as primary drivers.

The problem with this narrative is that it assumes the system resets. When we examine homelessness over time, what emerges is not a series of short-term fluctuations, but a system that keeps building on itself. The level observed today is not simply a function of current conditions, rather, it is heavily shaped by what came before.

Homelessness is best understood as a system with underlying momentum operating under constraint. How it is framed matters, because it changes both the interpretation of its drivers and the implications for policy. In this context, the "system" refers to the interaction between the housing market and households at risk of homelessness. Specifically, it refers to how individuals move into and out of homelessness given available housing, income supports, and broader economic conditions.

With the help of our partners at the Affordable Housing Association of Nova Scotia, we were able to assemble a consistent monthly dataset on homelessness in Halifax spanning from October 2019 through March 2026. This provides a rare opportunity to examine how homelessness has evolved over time using a continuous and relatively high-frequency series. Rather than focusing on individual points or short-term changes, the data allows us to step back and look at the broader time-series picture.

Using this dataset, we examined which factors are most consistently associated with changes in homelessness over time, while accounting for persistence in the series. We developed and tested a range of econometric models across multiple specifications, incorporating housing market conditions, income supports, labour market indicators, and policy variables. Each model was evaluated for statistical validity, including tests for underlying momentum, autocorrelation, and overall stability. The goal was not simply to fit the data, but to isolate relationships that are both empirically robust and economically meaningful.

A detailed discussion of the data, model development, alternative specifications, and diagnostic testing can be found in the Appendix.

The results point to a consistent set of relationships. Based on the preferred specification, homelessness in Halifax is primarily shaped by its own persistence, housing capacity, and broader economic pressures, with income supports playing a meaningful preventative role. These relationships are consistent across specifications and form the basis of the analysis that follows.

Active Homelessness in HRM

Before examining the drivers of homelessness, it is important to establish the scale of the change. Active homelessness in HRM increased from 220 individuals in October 2019 to 1,860 by March 2026 (see Chart 1). This represents a net increase of 1,640 people, or 745%. Put differently, active homelessness was more than eight times higher at the end of the period than at the beginning.

Homelessness first rose sharply through 2020, increasing from 220 in October 2019 to 491 in December 2020, a gain of 271 people, or 123%. This period overlapped with the first year of the COVID-19 pandemic, when emergency accommodation and public-health responses reshaped the homelessness system.

This was followed by a partial decline in early 2021. Active homelessness fell from 487 in January 2021 to 324 in June 2021, a decline of 163 people, or 33%. However, even at this low point, homelessness remained above its pre-pandemic level.

From mid-2021 onward, homelessness began rising again, reaching 738 by December 2022 and 1,309 by July 2024. This period coincided with rapid population growth, mainly due to immigration from other parts of Canada because of economic growth in HRM, rising rents, and an extremely tight rental market. In other words, the number of people at risk was growing while the housing system had limited capacity to absorb pressure.

The series then eased from 1,309 in July 2024 to 987 in August 2025, a decline of 322 people, or 25%. This decline coincided with some rental-market easing and expanded temporary or supportive housing responses. Still, homelessness remained more than four times higher than in October 2019.

The final months of the series show a sharp renewed increase. From August 2025 to March 2026, active homelessness rose from 987 to 1,860, an increase of 873 people, or 88%, in just seven months. This pushed homelessness to its highest observed level.

Overall, the pattern points to a system that does not reset. Declines occur, but they are partial and temporary. Each renewed increase begins from a higher base, leaving homelessness progressively more elevated over time. This descriptive pattern is consistent with the central finding of the model: homelessness in HRM behaves less like a short-term flow and more like a persistent stock operating under housing constraint.

Chart 1: Active Homelessness in HRM

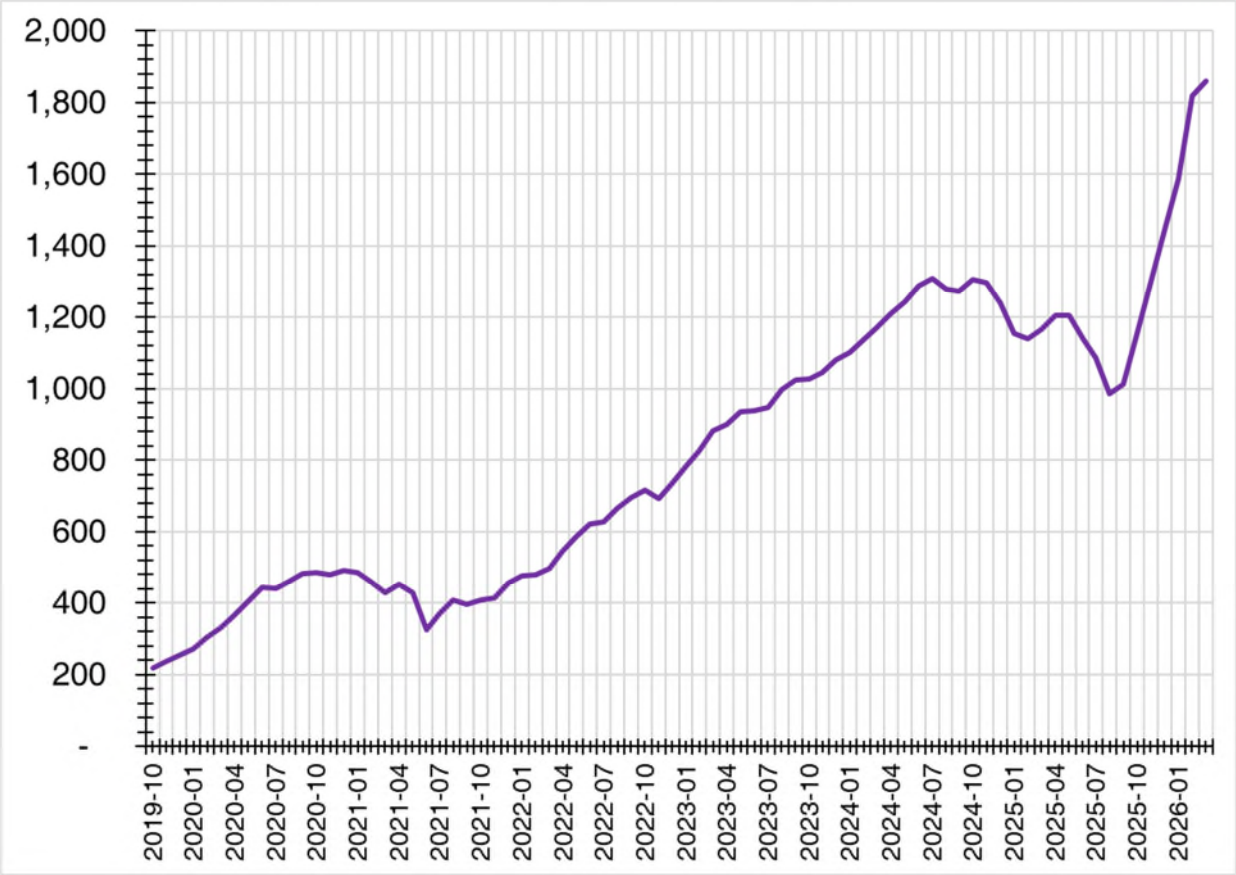

Source: AHANS & TDP EIU

Persistence

At first glance, homelessness appears to move in response to economic and housing conditions. There are periods of increase, stabilization, and occasional decline. But this view is not holistic enough. What matters is not just the direction of change in the level of homelessness, but the size of the force pushing those changes.

When we model homelessness dynamically, we see clearly that the past strongly determines the present. **Approximately 85% of the previous month's level carries forward into the next.** This represents a high degree of persistence. In practical terms, this means that homelessness behaves less like a flow and more like a stock.

A temporary increase does not dissipate; it accumulates. Each month builds on the last, and the system adjusts only gradually. Even when underlying conditions improve, the level remains elevated because it is anchored by its own history. This helps explain why homelessness can continue rising even as certain indicators stabilize.

The persistence in homelessness can also be illustrated by how slowly the system adjusts over time. As shown in Chart 2, even if no new individuals enter homelessness, a large share carries forward from one month to the next. The decline is gradual rather than immediate, with a significant portion remaining even after several months.

Chart 2: Persistence in HRM Homelessness

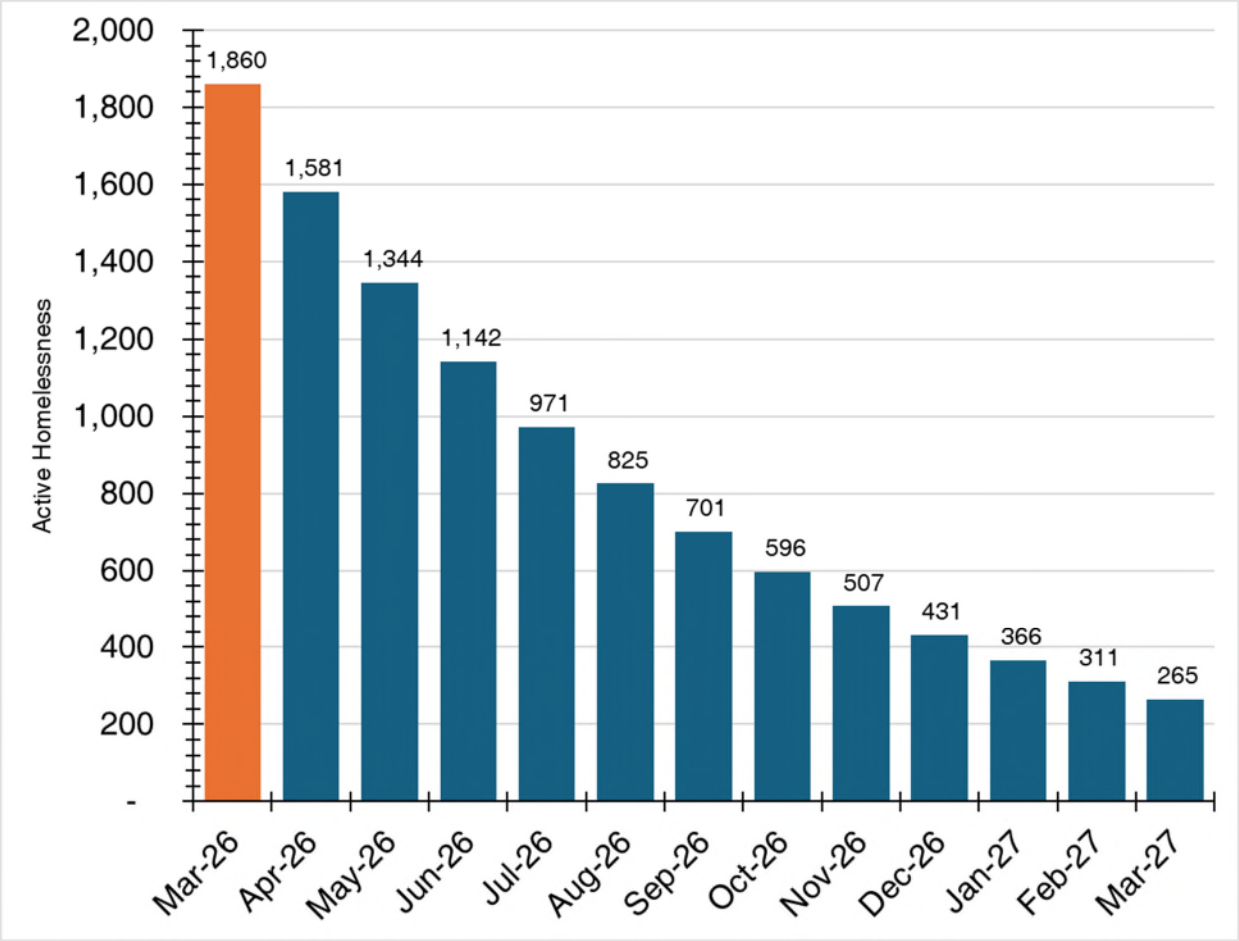

Source: AHANS & TDP EIU
Note: This chart is illustrative in nature, showing decline in homelessness assuming 85% of homelessness carries forward each month.

Housing Capacity

If persistence explains why homelessness accumulates, housing availability explains why it rises.

Across all model specifications, the vacancy rate, with a month's lag, emerges as the most important structural variable. The relationship is both statistically strong and economically meaningful. **A one percentage point increase in vacancy is associated with, approximately, a 5% reduction in homelessness, within the modelled period**.

As shown in the Chart 3, movements in vacancy rates tend to precede changes in homelessness, consistent with the lag structure in the model. Periods of tightening vacancy are followed by increases in homelessness, while improvements in availability take time to translate into declines.

To understand why, it is helpful to reframe the role of the housing market. Rather than being only a measure of housing supply, the vacancy rate is also an indicator of housing capacity.

Chart 3: Homelessness and Vacancy Rates, (lagged 1 month), November 2019 to March 2026

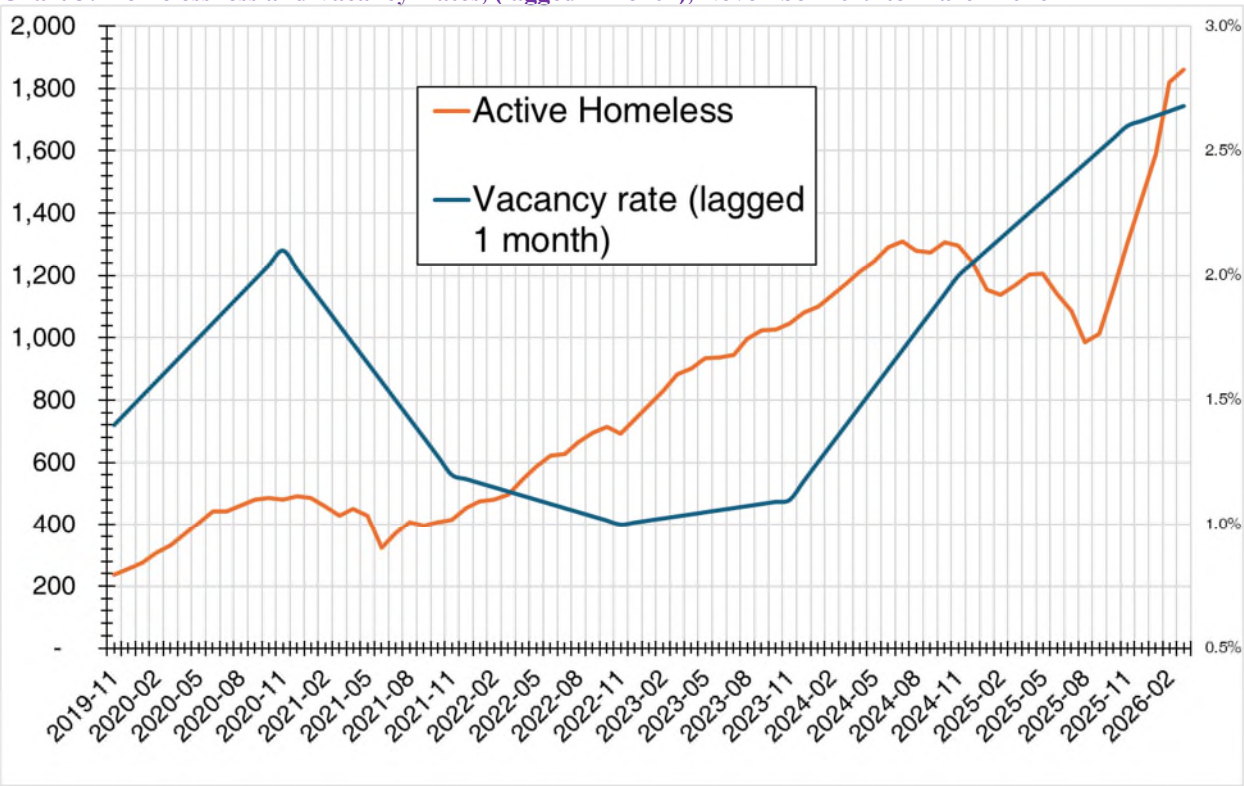

Source: AHANS, CMHC, Rentals.ca, & TDP EIU

When vacancy is low, the market has little flexibility. There are fewer units available, fewer alternatives for households under pressure, and limited ability to absorb shocks. Every sort of disruption, be it through job loss, illness, or rent increases, can lead to displacement. As vacancy increases, that capacity expands. Households have more options, transitions are easier, and fewer individuals are pushed into homelessness.

But not all vacant units are equally relevant. For households at risk of homelessness, what matters is the availability of units that are financially attainable and immediately accessible. In that sense, vacancy acts as a proxy for the availability of affordable housing at the margin; the segment of the market where displacement either occurs or is avoided. When vacancy is low, this segment can become functionally unavailable.

This is why vacancy matters so consistently in the model. It captures not just supply, but the functional availability of affordable housing within the system. Homelessness is not only a function of how much housing is built, but whether enough of that housing is accessible to those at risk.

The dynamics of this relationship become clearer when viewed differently. Rather than forming a single downward-sloping line, the relationship traces a loop over time as shown in Chart 4. The same vacancy rate is associated with very different levels of homelessness depending on when it is observed.

Changes in vacancy do not translate into immediate changes in homelessness. Instead, they work through the system with a lag, while existing levels are carried forward by strong underlying momentum. As a result, homelessness adjusts gradually, moving along a path rather than jumping to a new equilibrium.

In other words, the relationship is path-dependent. What matters is not just where vacancy is, but how it has been changing. The system retains memory. That is why the same market conditions can produce very different outcomes, and why improvements in housing availability, while effective, do not generate immediate reversals in homelessness.

Chart 4: Homelessness and Vacancy: A Lagged and Path-Dependent Relationship

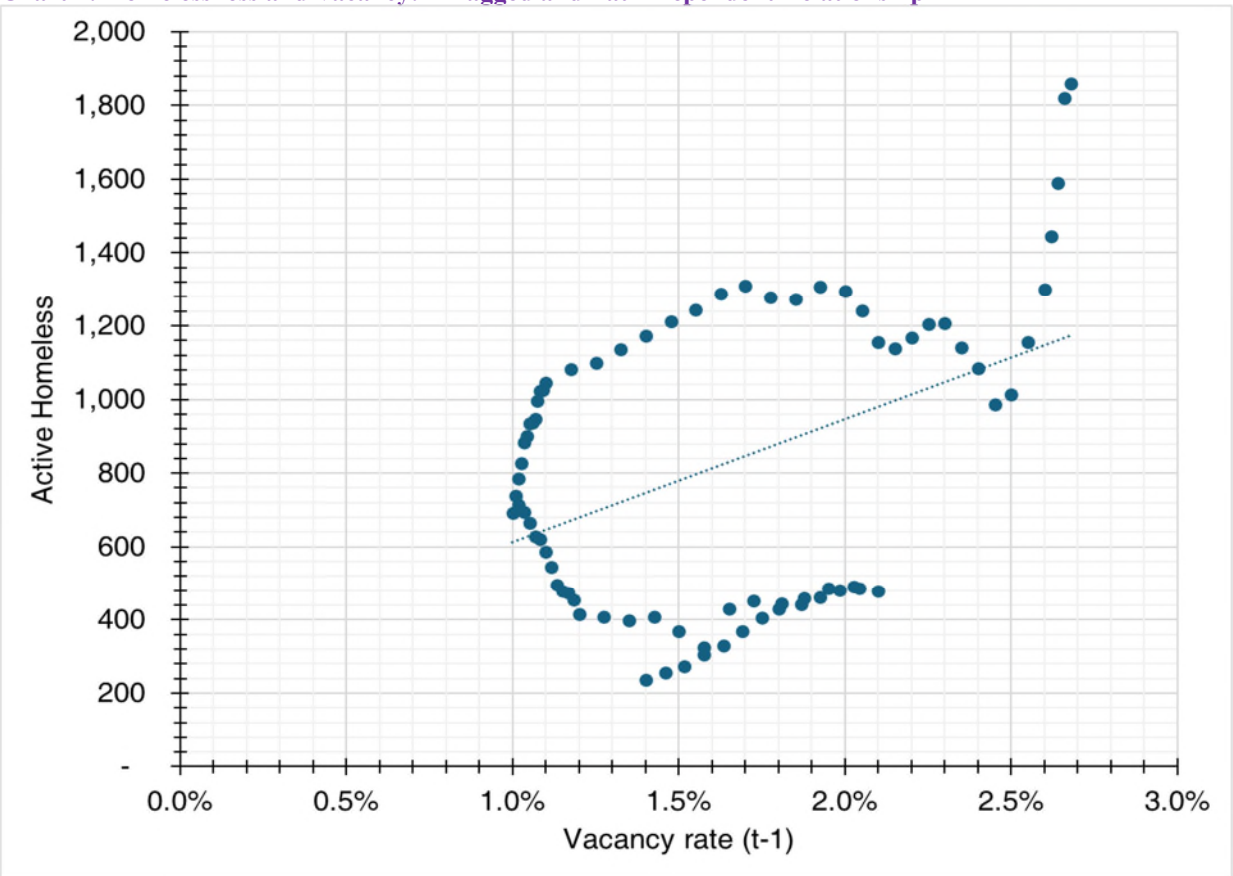

Source: AHANS, CMHC, Rentals.ca, & TDP EIU

Income Supports as a Preventative Measure

Using sickness benefit recipients i.e. those unemployed due to sickness, our model finds that **an increase of 1,000 recipients is associated with roughly a 1.9% reduction in homelessness, with the effect appearing after a lag.**

Unlike housing capacity, which affects the system with a month's lag, income supports operate more gradually over a four-month period. They stabilize households before displacement occurs, reducing the number of individuals entering homelessness. In other words, they influence the flow into the system, not the stock already within it.

In a system characterized by high persistence, reducing inflows is necessary but not sufficient to reduce overall levels. Even if fewer people enter homelessness, the existing population remains, and declines occur only slowly.

Income supports are therefore best understood as a preventative mechanism; they mitigate risks, but do not rapidly reverse existing conditions.

Chart 5: Homelessness and Sickness Benefits Recipients (per 1000, lagged 4 months), November 2019 to March 2026

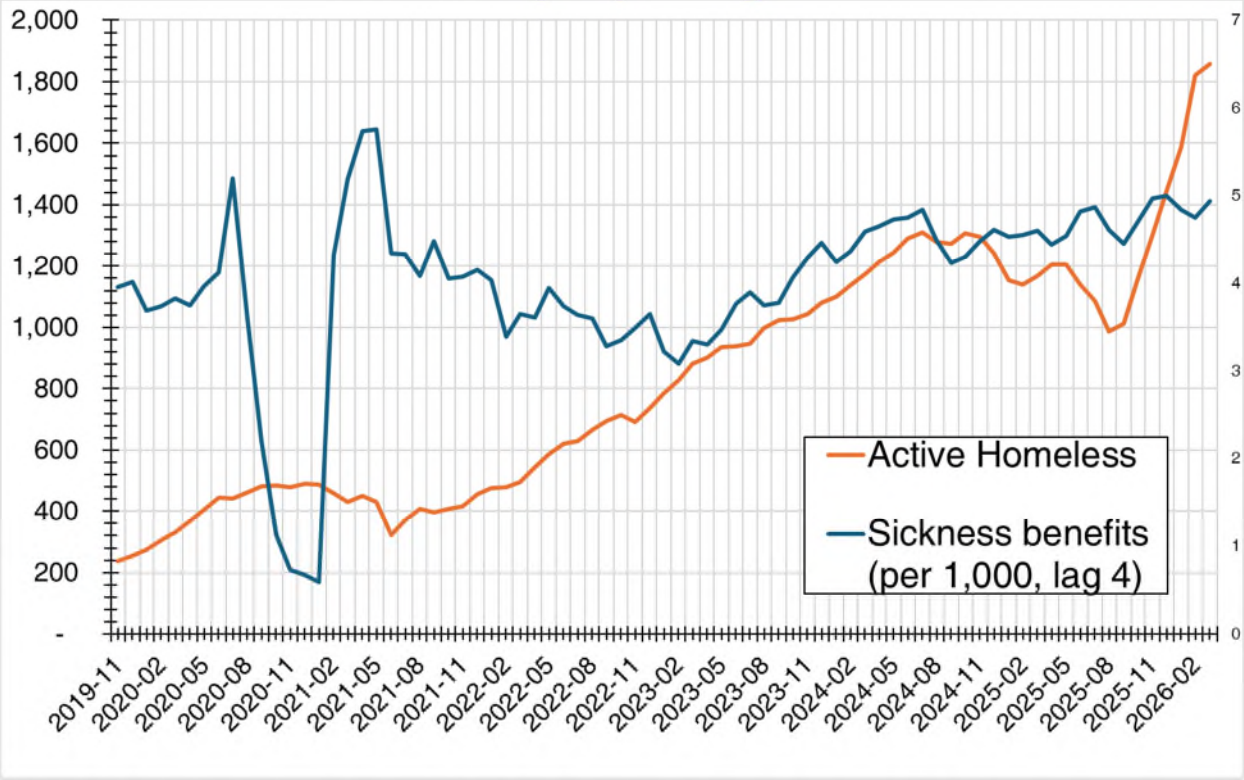

Source: AHANS and Statistics Canada

The Cost of Economic Growth

One of the more counterintuitive findings in the analysis is the role of wages. **Higher wages are associated with higher levels of homelessness.**

This is not because rising incomes make households worse off. Rather, wages act as a proxy for broader economic conditions. When wages increase, it typically reflects stronger labour markets, increased employment, and in the context of Halifax, greater in-migration. These dynamics increase demand for housing.

If supply does not expand at the same pace, that demand translates into tighter vacancy and increased pressure at the lower end of the market. The result is a system in which economic growth, while beneficial overall, can contribute to rising homelessness when housing constraints are binding. In our study, wages are capturing demand pressure, not affordability improvements.

Chart 6: Homelessness and Wage Rate, November 2019 to March 2026

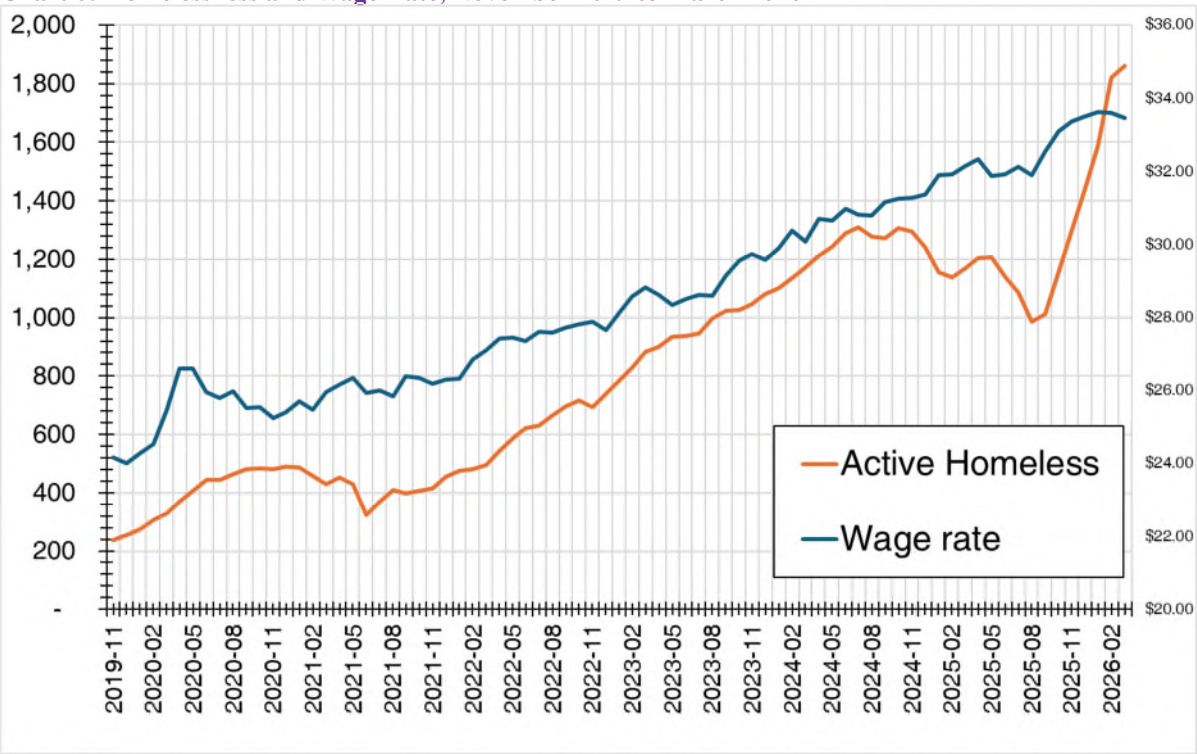

Source: AHANS and Statistics Canada

What Falls Away Under Closer Examination

A range of additional variables were tested, including unemployment, population growth, housing starts, rent inflation, and new home prices. Many of these are commonly cited as key drivers of homelessness.

Individually, they appear relevant. But once persistence and housing availability are accounted for, their independent effects largely disappear. This does not mean they are unimportant. Rather, it suggests that their influence is indirect. They operate through broader economic and housing market dynamics rather than exerting a separate, identifiable effect on homelessness.

For example, population growth increases demand, but its impact is ultimately reflected in vacancy. Similarly, housing starts affect supply, but their effect depends on whether they meaningfully alter market balance.

The Rent Cap

Rent regulation is often positioned as a central driver of housing outcomes, including homelessness. Yet there is little agreement on how, or even whether, it improves those outcomes. The data provides a more nuanced view.

In the baseline model, **the rent cap appears to be associated with a modest reduction in homelessness. At first glance, this suggests a measurable policy effect. However, this result does not hold under closer examination.**

When a general time trend is introduced, the estimated effect becomes statistically insignificant. This points to a key limitation in the data. The rent cap was implemented early in the sample period and remains in place for most of the timeframe, effectively dividing the data into a short “pre-period” and a much longer “post-period.” As a result, the rent cap variable is closely aligned with time itself.

In the baseline specification, it captures not only any potential effect of the policy, but also broader structural changes that occurred after 2020, including shifts in housing demand, economic conditions, and political changes. Once these broader time dynamics are explicitly accounted for, the independent contribution of the rent cap disappears. This suggests that the initial result reflects differences between periods rather than a clearly identified causal effect of rent regulation.

Model Performance

Our main model closely tracks observed homelessness over time and reproduces past levels with a high degree of accuracy. As shown in Chart 7, the fitted values move almost in tandem with the observed series, capturing not only the overall upward trend, but also the smaller fluctuations along the way. This includes the early rise through 2020, the dip in mid-2021, and the steady increase that follows through 2022 and 2023.

Importantly, the model also captures key turning points. The plateau observed in late 2024 and early 2025, followed by a decline into mid-2025, is reflected in the predicted values. Likewise, the sharp increase toward the end of the sample period is mirrored closely. While there are small deviations, particularly during more volatile periods, the model consistently follows the direction and magnitude of change. This indicates that it is not simply fitting a long-term trend, but is responsive to shifts in underlying conditions.

The model captures the core dynamics of homelessness in Halifax. It reflects both the gradual build-up driven by underlying momentum and the shorter-term movements associated with changes in housing availability and economic conditions. As a result, it provides a reliable representation of how homelessness evolves over time and a credible basis for interpreting the relationships identified in the analysis.

Chart 7: Actual Homelessness and Predicted Homelessness, November 2019 to March 2026

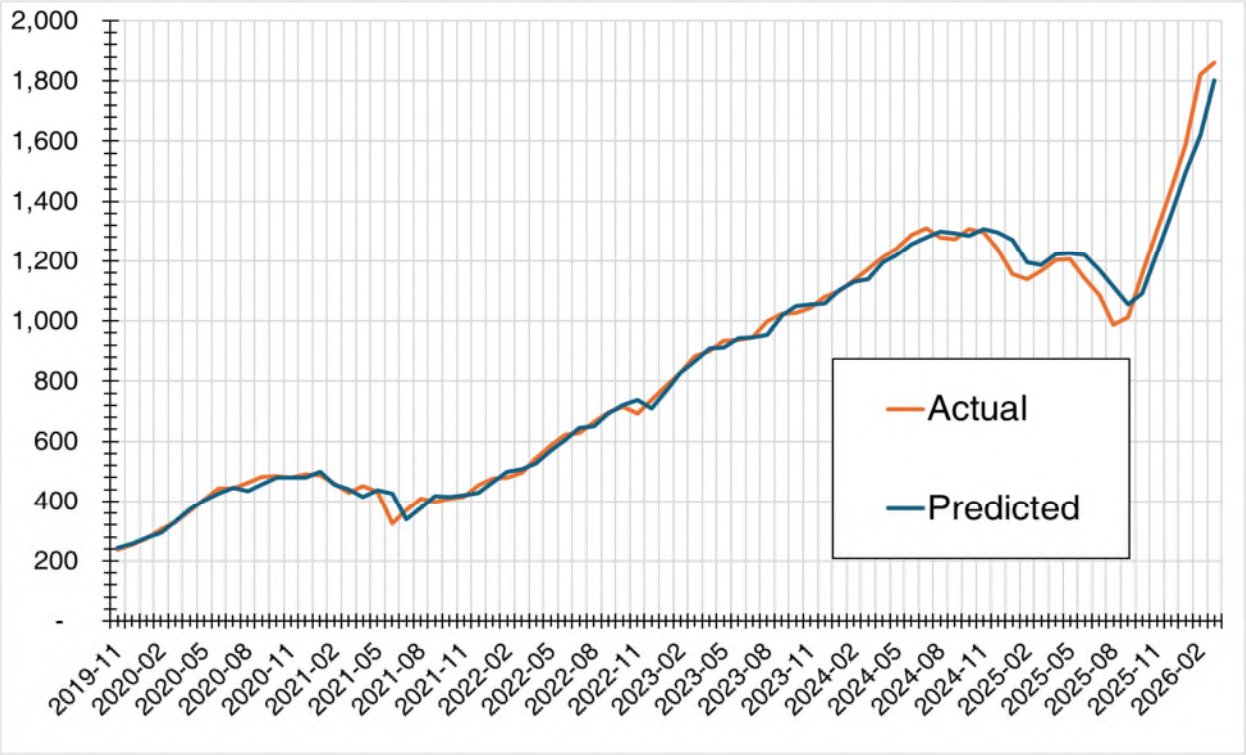

Source: AHANS & TDP EIU

Implications

Homelessness in Halifax is not primarily driven by short-term shocks or isolated policy interventions. It is the outcome of a persistent system operating within a constrained housing environment.

Persistence ensures that increases accumulate over time. Housing availability determines the system's capacity to absorb pressure. Economic conditions shape the level of that pressure. Income supports play a supporting role by reducing entry into the system, but they do not fundamentally alter its structure. These findings have several practical implications.

- First, timing is crucial. Because of persistence, delays in response allow the system to build. Once elevated, homelessness declines only gradually. *In plain language: it is tough to eliminate homelessness once it has taken hold.*

- Second, housing availability is key. Even modest increases in vacancy have large effects on homelessness, suggesting that supply conditions are a critical lever. *But housing must be available at prices people can afford, especially rental stock since it is an expense, rather than an investment for the renter.*

- Third, preventative measures are effective but limited. Income supports reduce inflows, but do not quickly reduce the existing stock of homelessness. *Or to put it another way, income support helps reduce the number entering homelessness but has limited impact on reducing the numbers of those already unhoused.*

- Fourth, economic growth must be matched by housing supply. If growth outpaces supply, increased demand translates into increased pressure. *It is comforting, but illusionary, to assume that homelessness will decline as the economy improves because the latter encourages in-migration which in turn increases demand.*

- Finally, policy effects are difficult to isolate in short time series. Apparent impacts may reflect broader structural changes rather than direct causal relationships. *We were not able to measure the impact of the rent cap for example, because the time period over which it had not been in force was too short.*

- *Every city is different. We can extract causes of homelessness from this study of HRM but they cannot be extrapolated in the same way for other communities.*

The most important insight from the analysis is not tied to any single variable. Homelessness in Halifax is not simply rising or falling in response to current conditions. It is being shaped by a system that carries forward its own history, operating within the constraints of a tight housing market. Until those constraints change in a meaningful way, the system is unlikely to adjust quickly. More than any individual factor, this explains why homelessness has proven so difficult to reverse.

The results also suggest that policy needs to distinguish between measures that reduce inflows into homelessness and measures that increase exits from homelessness. Income supports may help with the former, while housing availability is more directly related to the latter.

For a detailed discussion of methodology, model development, and diagnostics, see the Appendix below:

Appendix: Empirical Approach, Model Development, and Diagnostics

Overview

This analysis examines the drivers of homelessness in Halifax Regional Municipality using monthly data from October 2019 to March 2026, from the Affordable Housing Association of Nova Scotia (AHANS). The dependent variable is the number of active homeless individuals in HRM, while all explanatory variables are measured at the Nova Scotia level. This reflects both data availability and the broader reality that housing markets, labour conditions, and policy environments operate at a regional scale.

The objective was to develop a model that is not only statistically robust, but also economically interpretable; a model that isolates meaningful structural relationships rather than spurious correlations driven by shared trends.

Variables: Definitions, Form, and Rationale

The final specification (model) incorporates a focused set of variables selected based on both statistical performance and economic interpretability. Each variable is included to capture a distinct dimension of the system such as persistence, housing market conditions, income stability, and broader economic pressure.

The dependent variable is the number of active homeless individuals in Halifax from AHANS, expressed in natural logarithmic form. Modelling homelessness in logs allows coefficients to be interpreted as approximate percentage changes and helps stabilize variance over time. Given the upward trend and widening dispersion in the raw series, this transformation improves both the statistical properties of the model and the clarity of interpretation.

A lagged dependent variable is included, also in logarithmic form, to capture persistence in homelessness. This is a critical component of the model. Homelessness behaves as a stock rather than a flow, meaning that current levels are heavily influenced by past levels. Including the lagged term accounts for this inertia and prevents omitted variable bias arising from temporal dependence. Empirically, it is the strongest predictor in the model, reinforcing the idea that homelessness adjusts slowly over time.

Housing market conditions are captured through the vacancy rate, included with a one-period lag. The vacancy rate serves as the primary measure of housing market tightness. Lower vacancy implies fewer available units and a reduced ability for the system to absorb shocks, increasing the likelihood of displacement. The lag structure reflects the fact that changes in housing availability do not translate into homelessness immediately, but rather with a short delay. This variable emerges as the dominant structural driver in the model. Data on vacancy was sourced from Canada Mortgage and Housing Corporation (CMHC). Since vacancy rate data from CMHC is available only at an annual basis, we used linear interpolation to fill the gap[1]. Data for the months of 2026 were sourced from Rentals.ca.

Income stability is proxied using the number of sickness benefit recipients, lagged by 4 months and scaled per 1,000 individuals. This variable is intended to capture the strength of the social safety net, particularly for individuals facing health-related employment constraints. The lag reflects the time required for income supports to influence housing outcomes. Scaling the variable improves interpretability, allowing coefficients to be read in terms of changes per 1,000 recipients. The results suggest that income supports play a preventative role, reducing the flow of individuals into homelessness over time. Data on benefits were sourced from Statistics Canada, Table 14-10-0009-01: Employment insurance beneficiaries by type of income benefits, monthly, unadjusted for seasonality.

Broader economic conditions are represented by the wage rate. While wages are often interpreted as a measure of affordability, in this context they function more accurately as a proxy for overall economic activity. Rising wages tend to coincide with stronger labour markets, increased in-migration, and heightened demand for housing. This variable captures demand-side pressure rather than direct improvements in affordability. Its inclusion helps account for macroeconomic conditions that influence the housing market indirectly. Data was sourced from Statistics Canada, Table 14-10-0063-01: Employee wages by industry, monthly, unadjusted for seasonality.

Finally, a binary indicator is included to capture the period during which provincial rent control measures were in effect. This variable is intended to test whether policy intervention is associated with measurable changes in homelessness, independent of broader market conditions. However, given the limited number of observations prior to

[1] Because vacancy is central to the model, the use of interpolated annual CMHC vacancy data is an important limitation. The results should therefore be interpreted as capturing broader movements in housing availability rather than precise month-to-month changes in vacancy.

its implementation, and its overlap with broader time trends, this variable is interpreted cautiously. Subsequent robustness checks suggest that it captures period effects rather than a clearly identified causal relationship.

In summary, these variables provide a structured representation of the system. They allow the model to capture persistence, housing market constraints, income stabilization mechanisms, and macroeconomic pressures, while maintaining a level of parsimony necessary for clear interpretation.

Additional Variables Considered

In addition to the variables retained in the final specification, a broader set of indicators was initially considered, including population, unemployment rate, housing starts, new home prices, and rent inflation.

While these variables are commonly associated with housing market dynamics, they were not retained in the final model. In most cases, they did not exhibit statistical significance once persistence and housing market tightness were accounted for, and their inclusion did not improve model fit or interpretability.

This suggests that their influence on homelessness operates indirectly, primarily through their impact on housing availability and broader economic conditions, rather than as independent drivers.

Initial Specification: A Static Model

We began with a standard static specification, regressing homelessness on contemporaneous housing, economic, and policy variables.

At first glance, the results appeared strong. The model explained over 90% of the variation in homelessness, and most variables were statistically significant. Vacancy rates, income supports, wages, and the rent cap all appeared to have meaningful effects.

Source	SS	df	MS	Number of obs	=	77
				F(4, 72)	=	249.89
Model	19.3667093	4	4.84167733	Prob > F	=	0.0000
Residual	1.39499972	72	.019374996	R-squared	=	0.9328
				Adj R-squared	=	0.9291
Total	20.761709	76	.273180382	Root MSE	=	.13919

ln_homeless	Coef.	Std. Err.	t	P>|t|	[95% Conf.	Interval]
L_vacancy	-19.01182	3.81837	-4.98	0.000	-26.6236	-11.40003
L_benefits	-.062686	.0193322	-3.24	0.002	-.101224	-.0241481
Wagerateperhour	.2025742	.0092477	21.91	0.000	.1841393	.2210091
Provincialrentcapdummy	.1223282	.0602867	2.03	0.046	.0021487	.2425076
_cons	1.247233	.1829936	6.82	0.000	.8824416	1.612024

However, a closer inspection revealed clear issues. Coefficient magnitudes were implausibly large. For example, the vacancy rate was associated with a very large reduction in homelessness, and several variables that are conceptually distinct were simultaneously highly significant. This is a common symptom of time-series models that omit persistence, i.e., variables trending together over time can appear related even when no direct causal relationship exists.

In short, while the static model fit the data well, it did so for the wrong reasons. It was not suitable for interpretation.

Introducing Dynamics: Capturing Persistence

To address the issue of persistence, we added a lagged version of the dependent variable as an independent variable. The rationale is that one of the main drivers of homelessness this month is its level last month. This adjustment fundamentally changed the results.

Source	SS	df	MS	Number of obs	=	77
				F(5, 71)	=	1400.30
Model	20.5532849	5	4.11065699	Prob > F	=	0.0000
Residual	.208424115	71	.002935551	R-squared	=	0.9900
				Adj R-squared	=	0.9893
Total	20.761709	76	.273180382	Root MSE	=	.05418

ln_homeless	Coef.	Std. Err.	t	P>|t|	[95% Conf.	Interval]
L_ln_homeless	.8465972	.0421089	20.10	0.000	.7626344	.9305599
L_vacancy	-5.04447	1.640634	-3.07	0.003	-8.315802	-1.773138
L_benefits	-.0189625	.0078329	-2.42	0.018	-.0345809	-.0033441
Wagerateperhour	.0381261	.0089365	4.27	0.000	.0203072	.055945
Provincialrentcapdummy	-.0556412	.0250804	-2.22	0.030	-.1056502	-.0056322
_cons	.151153	.0896986	1.69	0.096	-.0277009	.330007

The coefficient on lagged homelessness was approximately 0.85, indicating a very high degree of persistence. Once homelessness increases, it tends to remain elevated over time. This aligns with the underlying structure of the system: homelessness is a stock that adjusts gradually, not a flow that responds immediately to short-term changes.

At the same time, the inclusion of this dynamic term corrected the earlier distortions. Coefficient magnitudes became more realistic, and variables that previously appeared important lost statistical significance. Model fit improved substantially, with the R^2 rising to approximately 0.99. More importantly, the model became economically coherent.

Diagnosing the Model

Despite these improvements, diagnostic testing indicated that the model still suffered from serial correlation. The Breusch-Godfrey test identified statistically significant autocorrelation at higher lags, confirming that residuals were not fully independent over time. This is not unusual in monthly time-series data, particularly in the presence of persistence.

Tests for heteroskedasticity, including Breusch-Pagan and White's test, did not indicate any major issues. Variance appeared broadly stable.

We also assessed multicollinearity. Variance Inflation Factors (VIF) were elevated for the wage variable and the lagged dependent variable. This reflects the fact that both variables trend over time and is not unexpected. Importantly, coefficients remained stable across specifications, suggesting that multicollinearity was not distorting the results. Importantly, the presence of multicollinearity does not affect the unbiasedness of the estimates, but may inflate standard errors.

Correcting Inference: Newey–West Estimation

Given the presence of autocorrelation, we re-estimated the model using Newey–West (heteroskedasticity and autocorrelation consistent) standard errors. This step does not change the coefficients themselves, but ensures that standard errors, and therefore statistical inference, are reliable.

Regression with Newey-West standard errors — Number of obs = 77
maximum lag: 3 — F(5, 71) = 4917.03
Prob > F = 0.0000

ln_homeless	Coef.	Newey-West Std. Err.	t	P>|t|	[95% Conf.	Interval]
L_ln_homeless	.8465972	.0339886	24.91	0.000	.7788258	.9143685
L_vacancy	-5.04447	1.603483	-3.15	0.002	-8.241724	-1.847217
L_benefits	-.0189625	.0059049	-3.21	0.002	-.0307364	-.0071885
Wagerateperhour	.0381261	.0092491	4.12	0.000	.0196839	.0565682
Provincialrentcapdummy	-.0556412	.0158561	-3.51	0.001	-.0872575	-.024025
_cons	.151153	.0775782	1.95	0.055	-.0035335	.3058396

The results were highly consistent with the dynamic OLS model. All core variables remained statistically significant, and coefficient magnitudes were virtually unchanged. This provided confidence that the relationships identified were not a result of incorrect inference. From this point forward, the Newey–West specification was treated as the baseline model.

Testing a Broader Specification

To ensure that relevant drivers were not being omitted, we estimated a fully specified model including a wider set of variables: population, unemployment, housing starts, new home prices, and rent inflation. The results were instructive.

Regression with Newey-West standard errors — Number of obs = 77
maximum lag: 3 — F(10, 66) = 1968.68
Prob > F = 0.0000

ln_homeless	Coef.	Newey-West Std. Err.	t	P>|t|	[95% Conf.	Interval]
L_ln_homeless	.9074081	.0443954	20.44	0.000	.8187699	.9960464
L_vacancy	-3.649638	1.93445	-1.89	0.064	-7.511893	.2126171
sicknessbenefitrecipients_lagg	-.0000173	6.37e-06	-2.71	0.009	-.00003	-4.54e-06
Wagerateperhour	.0662468	.0187497	3.53	0.001	.0288119	.1036817
Unemploymentrate	-.7579345	.6778648	-1.12	0.268	-2.111335	.5954663
Populationthousands	-.003493	.0018868	-1.85	0.069	-.0072602	.0002742
rent_inflation	.4860595	.339221	1.43	0.157	-.191217	1.163336
HousingStarts	-.0000131	.0000197	-0.66	0.509	-.0000524	.0000262
NewHomePrices	-1.11e-08	8.36e-08	-0.13	0.895	-1.78e-07	1.56e-07
Provincialrentcapdummy	-.0613666	.0244627	-2.51	0.015	-.1102079	-.0125252
_cons	1.993825	1.012174	1.97	0.053	-.0270455	4.014696

While the core variables remained broadly consistent, most additional variables were statistically insignificant. Their inclusion did not improve model fit and, in some cases, reduced coefficient stability.

This suggests that many commonly cited drivers of homelessness operate indirectly, primarily through their effect on housing market conditions, rather than as independent determinants.

The Preferred Specification

The final model includes five variables: lagged homelessness, vacancy rate, income supports, wages, and the rent cap indicator. The results are both statistically strong and economically intuitive.

Homelessness exhibits a high degree of persistence, with approximately 85% carrying forward from one month to the next. Vacancy rates are strongly associated with lower levels of homelessness, indicating that housing availability is a primary driver. Income supports reduce homelessness with a lag, consistent with a preventative role. Wages are positively associated with homelessness, reflecting broader economic pressures and increased housing demand.

The rent cap variable is statistically significant in this specification. However, further testing suggests that this result should be interpreted with caution.

Regression with Newey-West standard errors — Number of obs = 77
maximum lag: 3 — F(5, 71) = 4917.03
Prob > F = 0.0000

ln_homeless	Coef.	Newey-West Std. Err.	t	P>|t|	[95% Conf.	Interval]
L_ln_homeless	.8465972	.0339886	24.91	0.000	.7788258	.9143685
L_vacancy	-5.04447	1.603483	-3.15	0.002	-8.241724	-1.847217
L_benefits	-.0189625	.0059049	-3.21	0.002	-.0307364	-.0071885
Wagerateperhour	.0381261	.0092491	4.12	0.000	.0196839	.0565682
Provincialrentcapdummy	-.0556412	.0158561	-3.51	0.001	-.0872575	-.024025
_cons	.151153	.0775782	1.95	0.055	-.0035335	.3058396

Robustness Checks

To test the stability of the results, we conducted several robustness checks.

First, we removed the rent cap variable entirely. The remaining coefficients were largely unchanged, and model performance was effectively identical, indicating that the rent cap is not a core structural driver of homelessness in this specification. The variable is nonetheless retained in the final model to test for policy effects and to demonstrate that the core relationships are robust to its inclusion.

Second, we introduced a linear time trend to capture broader movements in homelessness over the sample period. When this trend is included, the rent cap variable becomes statistically insignificant. This reflects an important identification issue in the data. The rent cap was implemented early in the sample and remains in place for the vast majority of observations, meaning the variable effectively divides the dataset into a short “pre-period” and a much longer “post-period.” As a result, the rent cap dummy is highly correlated with time itself. In the baseline specification, this allows it to capture not only any potential policy effect, but also other structural changes occurring after 2020, including shifts in housing demand, economic conditions, and measurement practices. Once a general time trend is introduced, these broader pre- and post-period differences are explicitly accounted for, and the independent contribution of the rent cap variable disappears. This suggests that the original coefficient reflects differences between periods rather than a clearly identified causal effect of the policy.

Finally, we replaced wages with the unemployment rate as an alternative macroeconomic control. In this specification, the model becomes unstable: the coefficient on lagged homelessness increases to approximately one, and other variables lose statistical significance. This is indicative of the persistence term approaching unity, meaning that current homelessness is almost entirely explained by its past value. When this occurs, the model effectively attributes nearly all variation in homelessness to its own past trajectory, leaving little explanatory power for other variables. This is a sign of misspecification rather than a true change in underlying relationships. Including wages helps capture broader economic conditions and demand-side pressures in a way that stabilizes the model and allows other structural variables to retain explanatory power. By comparison, unemployment appears to absorb or distort these dynamics, leading to an over-reliance on persistence and a loss of interpretability.

Predictive Performance

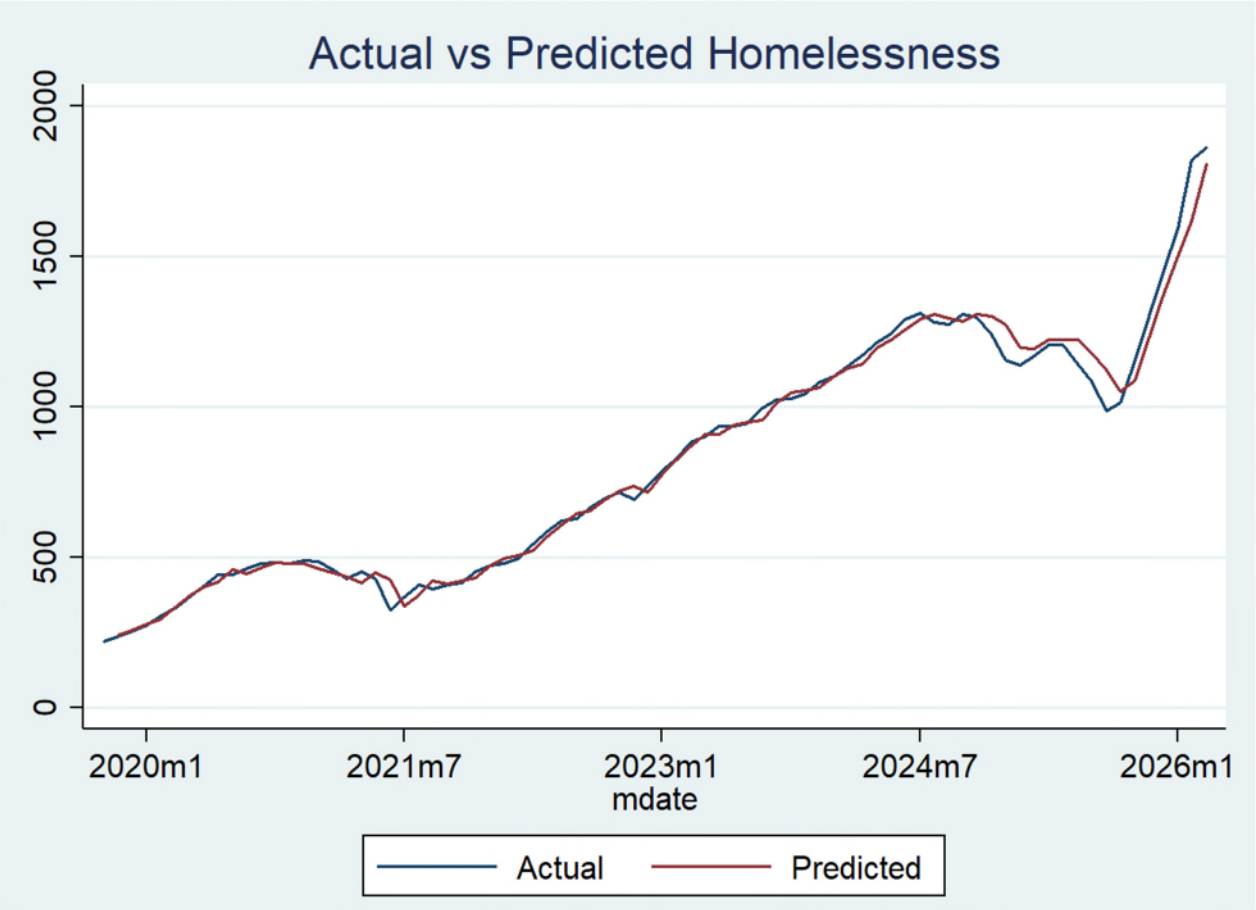

The model's predictive performance provides an additional check on its validity. Predicted values, constructed using observed lagged homelessness rather than recursive forecasts, closely track the actual data over time. The model reproduces both the level and trajectory of homelessness, with only minor deviations.

This confirms that the high explanatory power of the model reflects genuine structure in the data, rather than overfitting or spurious correlation.

Key Takeaways

Several consistent conclusions emerge from the analysis.

- Homelessness is highly persistent and adjusts slowly over time.
- Housing market conditions, as captured by the vacancy rate, are the dominant driver.
- Income supports play a meaningful but lagged role in preventing entry into homelessness.
- Broader economic conditions influence outcomes through demand-side pressure.
- Policy effects, while present in some specifications, are not robust once broader time dynamics are taken into account.

Limitations

Explanatory variables are measured at the provincial level, while homelessness is specific to HRM. Some variables, particularly wages, capture multiple underlying forces of the economy. Housing supply dynamics may operate over longer time horizons than those captured in the model. Data on vacancy rates were complied from multiple sources, as well as through interpolation.

Conclusion

The analysis progresses from a simple static specification to a dynamic, parsimonious model supported by extensive diagnostics and robustness testing.

The final model corrects for persistence and autocorrelation, isolates structurally meaningful variables, and provides a reliable empirical foundation for understanding homelessness in Halifax.

Importantly, it shifts the focus away from short-term fluctuations and toward the underlying conditions that sustain homelessness over time.