OXFORD JOURNAL
OF STUDENT
SCHOLARSHIP

Predicting Property Crime in Chicago: A Time-Series Analysis of Macroeconomic, Social, and Seasonal Factors

Richard Chen
rich.long.chen@gmail.com

ABSTRACT

In 2025 alone, roughly 108,000 reported instances of property crime in Chicago have caused over 865 million dollars in estimated societal costs [1]. This paper aims to answer three questions central to criminology and economics studies in Chicago: (1) What factors have been associated with property crime historically? (2) How accurately can models forecast property crime on city and police district levels? and (3) In which police districts are future crimes concentrated? Our study uses Random Forest, Ridge Regression, and Autoregressive Distributed Lag OLS models to analyze how well socioeconomic factors, weather trends, and public unrest predict property crime. Namely, we discovered that property crime had significant positive associations with temperature, and negative relationships with unemployment, protest interest, and precipitation. Our SARIMA forecast anticipates persistent crime trends through 2029. Damped ETS suggests sizable movements of districts in their relative rankings of crime risk, although wide prediction intervals make projections uncertain. Using predictive analysis results in conjunction with forecasts identifying future high-risk districts, we propose recommendations for Chicago to mitigate victimizations and damages of property crime.

INTRODUCTION

According to aggregated FBI and crime data reporting, property-related offenses account for approximately 80% of all crimes [2]. In 2022, there were 13.4 million property victimizations of U.S. households, and similar numbers were reported in the following years [2]. The millions of estimated property crime accounts that occur annually emphasize its substantial role in shaping urban safety and community well-being.

With a city population of 2.74 million, Chicago is the most populous city in Illinois, and it is the third-largest city in the United States. Demographically, Chicago is one of the most diverse cities in the United States. Its neighborhoods range from dense commercial districts to residential communities. As a major metropolitan hub with shifting demographics and evolving economic conditions, Chicago provides a complex environment for studying patterns of criminal activity. However, while prior scholarship [3]

analyzes the spatial concentration of crime, this paper takes a time-series approach and focuses on property crime specifically.

High tourism amplifies the importance of addressing property damage because the activity increases exposure to theft, vandalism, and other property-related offenses in commercial and public spaces. Further, in large metropolitan areas such as Chicago, acts of vandalism occur frequently, requiring substantial public resources to restore affected spaces to their original condition [4]. Those most at risk include residents and vehicle owners, especially in high-foot-traffic areas which face higher risks of theft, burglary, and vandalism [5].

Risk of property crime in Chicago is driven by a combination of macroeconomic, environmental, and situational factors. Over the past twenty-five years, Chicago has experienced significant fluctuations in criminal damage due to major structural disruptions such as the 2008 financial crisis, Covid-19 pandemic, and large-scale protest activity. Chicago fell into a severe, sustained recession during the 2008 financial crisis characterized by rising unemployment, declining household income, and increased financial instability [6]. More than a decade later, the Covid-19 pandemic caused severe social and economic disruption [7]. Business closures, abrupt unemployment spikes, and shifts in routine activity expressed changes in crime trends. Large-scale protests such as the George Floyd protest in 2020 produced short-term surges in social unrest and affected policing practices [8]. In particular, Chicago experienced a series of civil disturbances due to protests over Floyd's murder that turned into riots, causing the government to send in the Illinois National Guard, and economic damage caused by the disturbances exceeding $66 million [9].

Overall, the overlapping economic, environmental, and social activity make Chicago an analytically rich setting for examining criminal damage trends over time. We will explore the patterns of property crime in Chicago over time, explained by unemployment and poverty, macroeconomic events, post-pandemic behavioral changes, and structural conditions. We will also discuss the potential risk mitigation concepts for homeowners and vehicle owners at risk, strengthening safety measures in high-exposure areas and analyzing the benefits of policing strategies and economic initiatives.

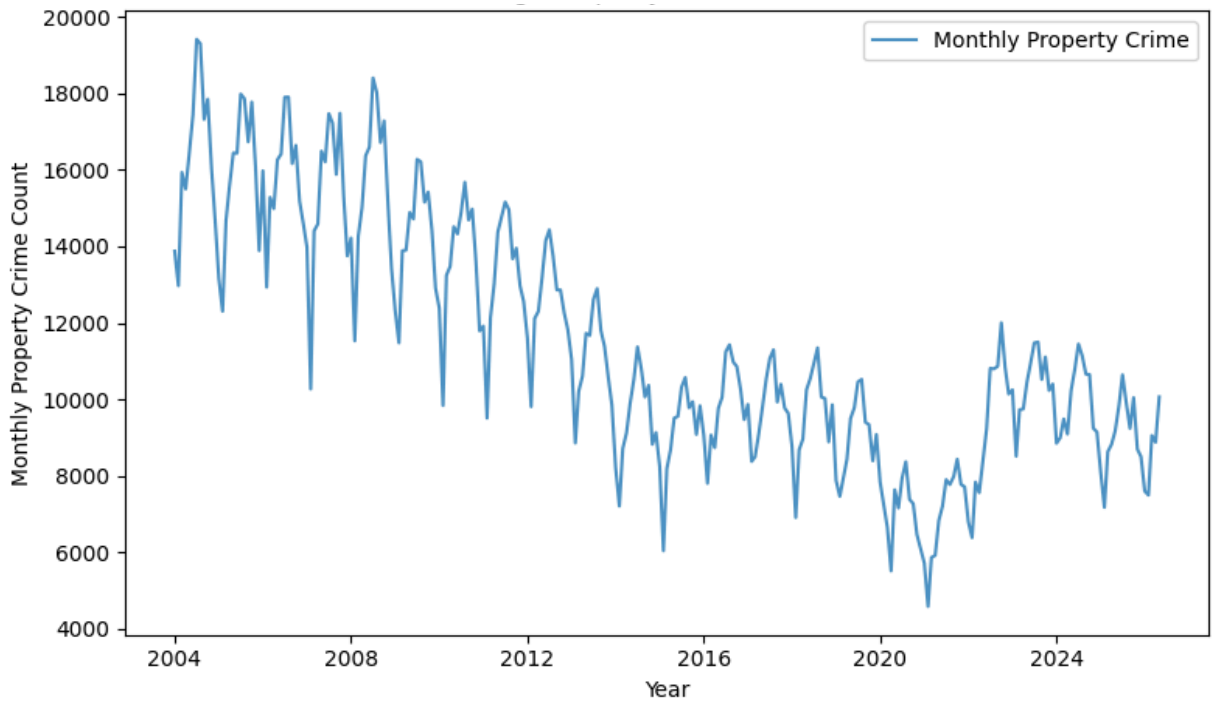
METHODS AND DATA

We have compiled the sources listed below to answer three main research questions: (1) What factors have been associated with property crime historically? (2) How accurately can models forecast property crime, and (3) Which police districts are future crimes concentrated in? This section explains our process for selecting and processing data sources along with graphs for contextualization. For every data source, we used observations from January 2004 to May 2026.

Dataset 1: Property Crime [10]
From the Chicago Data Portal, we retrieved time-series data for property crime counts in Chicago provided by the Chicago Police Department, the official law enforcement agency of the municipality. We

used the FBI's Uniform Crime Reporting National Incident-Based Reporting System (NIBRS) [11] to define property crime, under which the system puts burglary, larceny-theft, motor vehicle theft, arson, and vandalism. Although our data had no classification for vandalism, criminal damage closely corresponds to it, as outlined in Illinois statute 720 ILCS 5/21-1 [12]. Thus, we filtered our data for the following: burglary, theft, motor vehicle theft, arson, and criminal damage. Each instance of crime has multiple categorizations; namely, we used its type and location data for models and analytics. The trend is shown in Figure 1.

Figure 1: Chicago Monthly Property Crime



Dataset 2: Unemployment Rate [13]

For our unemployment rate variable, we used monthly data from the Federal Reserve Bank of St. Louis, which is a federal organization widely trusted for its data. Although the dataset is of the Chicago-Naperville-Elgin Metropolitan Statistical Area region, the data is still relatively accurate given that Chicago's large population means it represents a large portion of the region's unemployment rate. The unemployment rate for October 2025 was missing, but we filled it with the average of September and November values.

Dataset 3: Poverty Rate [14]

The Federal Reserve Bank of St. Louis provided our data for the number of impoverished people in Cook County. Cook County data is a reasonable proxy since the majority of the county is the city of Chicago. We divided the number of people in poverty by the population of Cook County to find the poverty rate,

which avoids correlation with the population predictor. For the purposes of keeping all variables monthly, we assumed the same yearly poverty rate for every month.

Dataset 4: Real Income [15]
Using the Federal Reserve Bank of St. Louis' data for median household income and consumer price index (CPI), we were able to deflate income by CPI to calculate real income, which is income adjusted for inflation. The adjustment gave a better measure of real household purchasing power important to modeling crime, ensuring reliability in our methodology. Since the data is still yearly, we used the yearly income for every month.

Dataset 5: Interest in Protests [16]
As a proxy for the interest in protests, we used Google Trends' data for Chicago's interest in the term "protest" over time. Google searches are a credible proxy because it has been used in professional research as an effective predictor for protests—as demonstrated in a study by Timoneda and Wibbels to be even more accurate than structural predictors [17]. We use this data to determine how protest activity, which could divert police attention or create social unrest, predicts property crime.

Dataset 6: Population [18]
We pulled our population data again from the Federal Reserve Bank of St. Louis. The data is of Cook County, a reliable metric of Chicago because the city of Chicago makes up a large share of the county's total population. We adjusted from annual to monthly by applying the data across all 12 months, preserving our monthly time-series data for modeling.

Dataset 8: Weather [19]
We got both monthly temperature and precipitation data from Climate at a Glance: City Time Series site by the National Oceanic and Atmospheric Administration (NOAA). NOAA provides a federal repository for US climate data popularly used in many papers for its quality weather records. Our study uses temperature and precipitation as a measure of outdoor activity, as we believe that it could predict crime. The November 2012 measure of temperature was missing, so we assumed the same temperature as the November average of the prior year.

Datasets for Crime Rates per Police District:
Section 3.3.2, which models and maps crime rates per 100,000 residents in each police district, requires combining crime data (source 1) with three new datasets: sources 9, 10, and 11.

Dataset 9: Police-district Boundaries [20]
Police district multi-polygon coordinates were provided by Chicago's Data Portal and were used to create choropleth maps with geographic boundaries.

Dataset 10: Census Tract Population [21]

Since districts have different populations, fair evaluations of risk necessitated crime rate—crimes per 100,000 people. Because population data was not available per district, we used census tract population data, obtained from the U.S. Census Bureau. Since census tracts were the smallest regions with population data, this choice gave us the most accurate overlay onto police districts. The same population was used for 2025 and 2029, so forecasted crime rates represent changes in crime, not population.

Dataset 11: Census Block To Police District Mapping [22]
Census units do not perfectly fill police districts, so we used area-weighted crosswalks provided by Zenodo's datasets to overlay the two. We then allocated tract populations based on the proportion of Census-Block area falling in the given district.

Variable Selection

We selected our variables on the basis of two criteria: policy relevance and past trends visualized in data characterization.

We prioritized many macroeconomic factors in our selection because they are theoretically connected to crime and are the most actionable. Importantly, economic conditions are strongly influenced by fiscal and monetary policy, which is itself a large field of research. Thus, to give insightful recommendations, we focused on the most common, structured, and research-supported form of intervention, policymaking.

We also chose weather because we observed from the data characterization that crime oscillates seasonally. Monthly temperature and precipitation are included because the amount of outdoor activity logically would be a predictor of crime rates. This relationship has been explored in other studies, such as Cohn [23], which has found significant positive relationships. Because of observed trends in crime and the precedence of other papers, we believe that seasonal weather predictors would improve the accuracy of our model.

A proxy for police presence was omitted in our model because of reverse causality: it is uncertain whether police presence causes a change in crime rates, or that police presence responds to changes in crime. Thus, it would be difficult to interpret results and propose recommendations.

Although Covid-19 caused greater unrest and changes in routine activity that may impact crime, the variable was omitted because there were no reliable proxies. Originally, our models included Covid cases data maintained by the Chicago Department of Public Health [24]. However, the number of case counts reported depends on testing and giving treatment, which became much less prevalent after 2022. Also, the dataset was discontinued in May 2024, and to avoid treating the absence of observations as zero case counts, we removed it from our models. In section 3.2.4, I added dummy variables for pandemic periods as a robustness check, but the simpler model performed better.

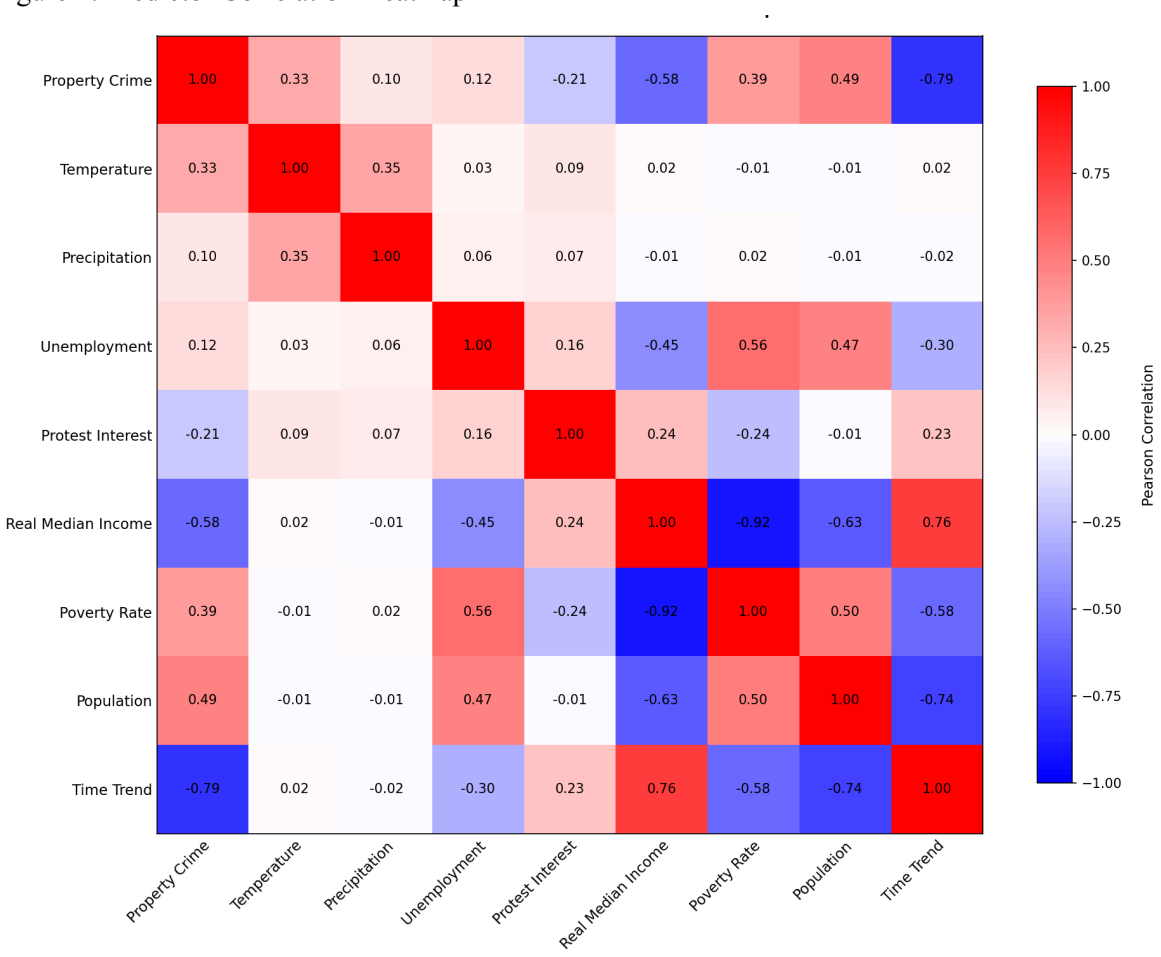
Exploratory Analysis

Before we begin modeling, we must ensure the statistical soundness of our data. Many macroeconomic variables tend to move together, so we used a heatmap to measure each variable's correlation with another. A heatmap, as visualized in Figure 2, allows us to identify predictors that conflate with each other and remove them to maintain interpretability. Our heatmap uses a score commonly found in statistical studies, the Pearson correlation coefficient, which is calculated using the formula:

$$r = \frac{\sum_{i=1}^{n}(x_i-\bar{x})(y_i-\bar{y})}{\sqrt{\sum_{i=1}^{n}(x_i-\bar{x})^2}\sqrt{\sum_{i=1}^{n}(y_i-\bar{y})^2}}$$

Where, r is the correlation coefficient, n is the number of observations, and X and Y are the variables.

Figure 2: Predictor Correlation Heatmap



We set the Pearson Correlation threshold to $|r| \geq 0.80$ since an absolute value as high as 0.8 indicates collinearity likely exists [25]. 0.8 is a general threshold where an absolute value above it likely suggests collinearity, whereas a value below it is relatively safe. We decided to remove the annual variables—real median income, poverty rate, and population—primarily since they do not temporally match with the

monthly ones. Our heatmap results echoed the concerns about median income and poverty rate: while annual variables appear generally problematic, all monthly predictors comfortably fall below the threshold.

PREDICTIVE AND FORECASTING METHODOLOGY

Assumptions

Assumption 1: Data from Cook County and the Chicago-Naperville-Elgin Metropolitan Area reflects the city of Chicago. Many of the variables we use only have available data for the broader region of Cook County of the metropolitan area. We assume these measures reflect the conditions in Chicago because the city makes up a large portion of the region's population and economic output. While suburban areas do create some unaccounted variation or noise, broader trends of the region are similar to those in the city of Chicago.

Assumption 2: The relationship between predictors and the target variable can be accurately represented using linear models, such as Ridge Regression and ARDL. Though the actual relationship may best be modeled using nonlinear methods, such as our Random Forest, the use of linear models is justifiable in our project given that they are widely used in the professional research fields of crime, economics, and policy.

Assumption 3: The main predictors of property crime are unemployment [26], poverty, income [27], protests, temperature [23], and precipitation. We assume that these variables capture the majority of the variation in crime because we included different types of drivers, including economic, social, political, and seasonal. Other criminology studies also focus on these variables, including the reference ones attached above.

Assumption 4: There are no major external shocks in the forecasted period.
In order for forecasting to be accurate, the structure of our model would need to continue and be relevant in the future. Thus, we assume that there are no unprecedented shocks, such as a pandemic, changes to policing, legal reform, mass migration, etc. However, major shocks have been historically rare, and given the relatively stable economic conditions and police policy in Chicago, unlikely.

Predictor Modeling Design and Results

Train-test split:

- For the purpose of consistency, all models used Jan 2004-May 2023 for the training and June 2023-May 2026 for the testing. We prevented the training model from using any values in the 36-month test window.

Calendar time trend:

- For each model, we also tested whether a calendar time trend, a variable whose value increases by 1 every elapsing month, could improve model accuracy. If crime trends are systemically and gradually upward or downward, the calendar time trend would capture this pattern.

Naive Baselines:

- All models were compared to naive baselines as a performance benchmark to test whether our complex models improved upon the simpler model. For this predictor analysis section, the objective was to evaluate performance sequentially as new information became available, so we used rolling origin models as the comparison model. I designed three types of baselines to represent the sources of accuracy: seasonal naive to measure improvement over yearly seasonality, previous-month naive for improvements over short-term persistence, and lag-only models to isolate the effect of predictors. Results are depicted in Table 1.

Table 1: Predictive Modeling Performances

Model	R^2	RMSE	MAE	RMSE improvement vs Seasonal Naive	RMSE improvement vs Lags-Only
Seasonal Naive t-12	0.486	820.8	687.6	-	-
Prev-month Naive t-1	0.507	804.0	666.9	2.0%	-
Random Forest - Lags Only	0.691	636.8	550.5	22.4%	-
Random Forest - Full	0.805	505.4	429.9	38.4%	20.6%
Ridge - Lags only	0.673	654.5	519.1	20.3%	-
Ridge - Full	0.644	683.1	548.0	16.8%	−4.4%
AR Only, p = 15	0.649	678.3	505.4	17.4%	-
ARDL - Full, p = 15, q = 1	0.834	466.8	378.5	43.1%	31.2%

We use three standard scores in statistical studies to determine the accuracy of all models.

1. The R^2 score measures how much of the variance in the dependent variable is explained by the model. In context, it is the percentage of variation in monthly property crime on the test set.
2. The RMSE (root mean squared error) measures the square root of the average of squared prediction errors. In context, a higher RMSE compared to the MAE tells us there are some months with larger errors, and on a numerical scale. However, RMSE weights large errors significantly higher than small ones (sensitive to outliers).
3. The MAE (mean absolute error) indicates the average absolute difference between predicted and actual results, reflecting the typical error size. In context, it is the number of crimes the prediction is off by per month.

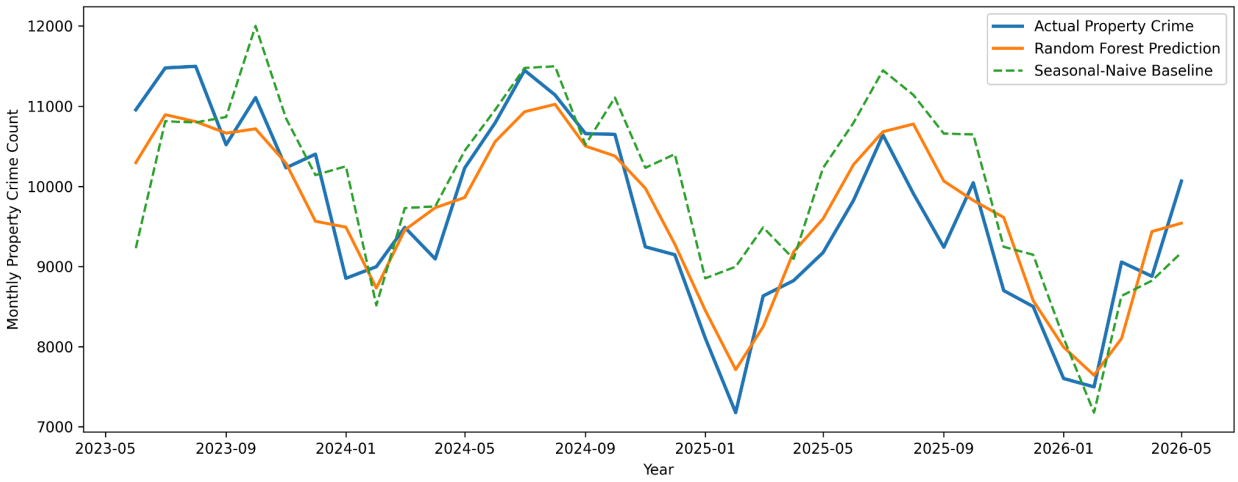
Random Forest

To accurately model and predict crime, we implemented a Random Forest Regressor Prediction Model. Unlike traditional statistical models, which may struggle with complex, nonlinear relationships, the Random Forest Regressor effectively captures variable interactions and complex patterns. Instead of relying on one decision tree, random forests aggregate the output of multiple trees to reduce error and variance. Each decision tree learns patterns by repeatedly splitting the data based on feature thresholds to minimize prediction error, but individual trees tend to overfit. Thus, random forests train each tree on a different bootstrap sample of the data and restrict each split to a random subset of features, reducing correlation between trees and preventing dominant variables from overwhelming the model.

The hyperparameters were selected using expanding-window time-series cross-validation. Combinations were evaluated across multiple validation periods, with each model trained only on observations preceding the corresponding validation window. The specification with the strongest out-of-sample was adopted for the model, with the number of trees being 1,200, maximum depth being 12, minimum split being 6, and minimum leaf being 2.

After testing multiple variations, the model with the best validation accuracy used the following: all the monthly predictors, a 1-month crime lag, and a 12-month crime lag. Notably, the addition of variables such as sine/cosine seasonal terms and a calendar-time trend did not improve model validation performance, so they were omitted. As shown in Figure 3, our model closely follows the actual trend while improving upon the seasonal naive baseline.

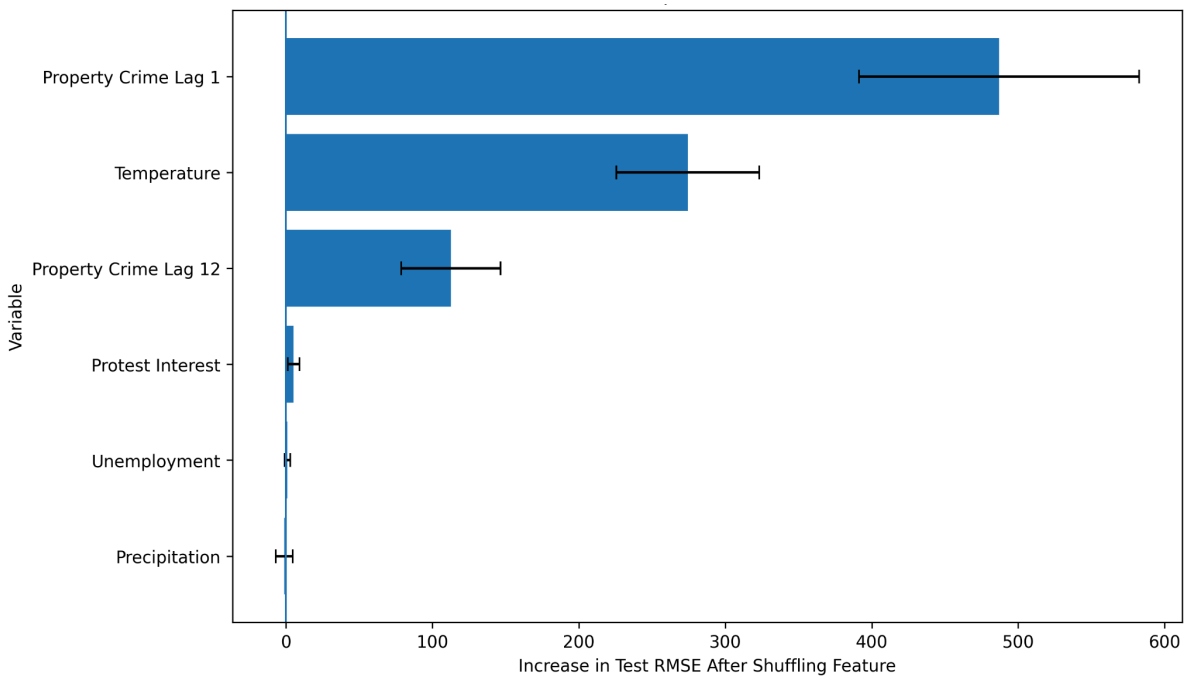
Figure 3: Random Forest (Actual, Predicted, and Baseline)



The Random Forest is overall a great predictor of crime. On average, predictions are off by about 430 crimes per month. With test-period monthly counts often in the 7,000–11,500 range, that's roughly 4–6% of typical levels, which is reasonable.

Then, we calculated how much RMSE increased after shuffling the predictors, a common method of measuring feature importance in non-linear models such as Random Forest. 1-month crime lag, 12-month crime lag, and temperature contribute almost all the predictive power, as portrayed in Figure 4. Random Forest seems to be primarily capturing crime persistence and seasonality, which emphasizes predictive accuracy, though not having the best interpretability.

Figure 4: Random Forest Permutation Importance



Ridge Regression

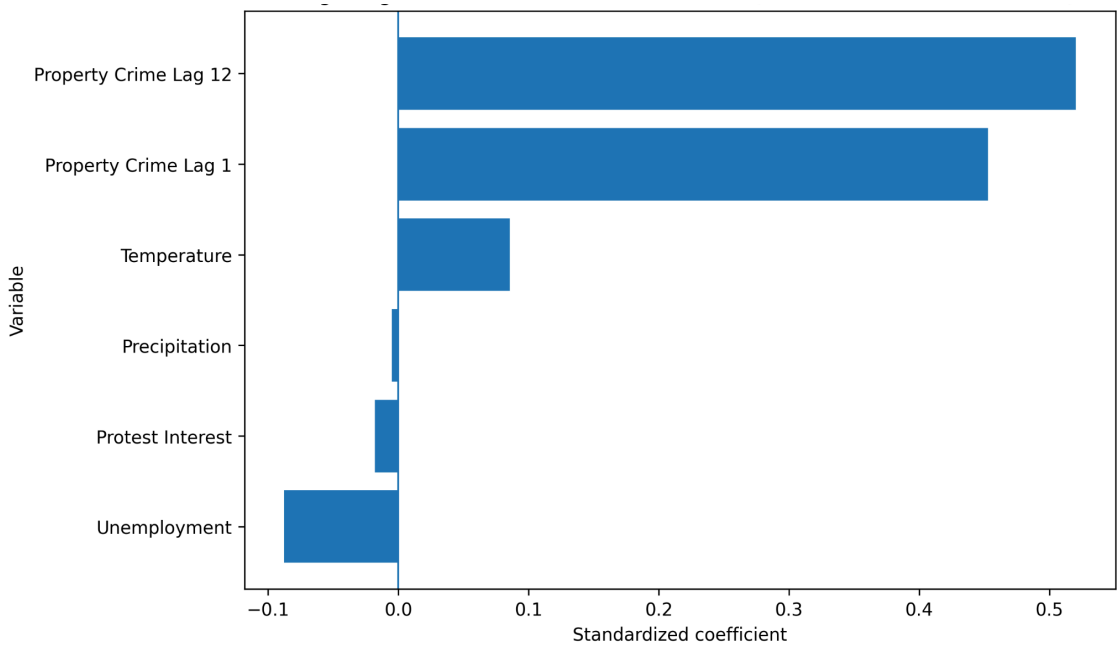
The ridge regression model is a variation of the Ordinary Least Squares (OLS) method of linear regression. Because we are conducting multivariable analysis, ridge regression specifically addresses multicollinearity in the variables by introducing an "L2" regularization parameter defined as $\lambda \sum \beta_j^2$ that constrains the optimization process and helps for multicollinearity. The optimal value of λ is selected using time-series cross validation to best balance bias and variability. Similar to the random forest model, we used the same three methods of evaluating accuracy for the ridge regression.

Comparing Ridge Regression with and without calendar time trend, we found that the one without performed better across expanding validation windows. The final, best performing model included the same variables as the Random Forest.

Overall, the accuracy metric results are not as strong as the other two models. We believe this is because of nonlinear trends and autoregressive patterns in crime, which the ridge regression model does not account for as deeply as other models.

We used the ridge regression model to evaluate the importance of each feature, concluding that the lags are the biggest predictor, followed by unemployment and temperature. Figure 5, as shown below, is calculated using standardized coefficients, which allow us to compare the magnitude of each coefficient by homogenizing their units to standard deviations.

Figure 5: Ridge Regression Standardized Coefficients



Autoregressive Distributed Lag OLS

Finally, we used the Autoregressive Distributed Lag model (ARDL). The method uses Ordinary Least Squares (OLS) to fit coefficients, but it has an autoregressive structure that uses past values of the target variable (crime) to predict itself.
ARDL general form:

$$y_t = \alpha + \delta t + \sum_{i=1}^{j} a_i y_{t-i} + \sum_{r=1}^{P} \sum_{l=0}^{k} b_{r,l} x_{r,t-l} + \sum_{m=2}^{12} \gamma_m D_{m,t} + \epsilon_t$$

Where j and k represent the number of months lagged.
To choose ARDL lags, we used a Bayesian Information Criterion (BIC), which tests combinations of target variable lags and predictor lags for the best training-period fit, but punishes the addition of extra parameters. The BIC searched orders j = 1–24 and k = 0–6. We screened candidates for residual autocorrelation, and then continued the specification with the lowest BIC. For both the models including and not including calendar time trend, the BIC chose {j=15,k=1}. Omitting the calendar time trend improved the RMSE across five expanding validation windows from 737 to 717, so we proceeded with the version without the trend. Monthly fixed effects were implemented to further account for seasonal patterns. Performance of our final ARDL model is visualized in Figure 6.

Figure 6: ARDL (Actual, Predicted, and Baseline)

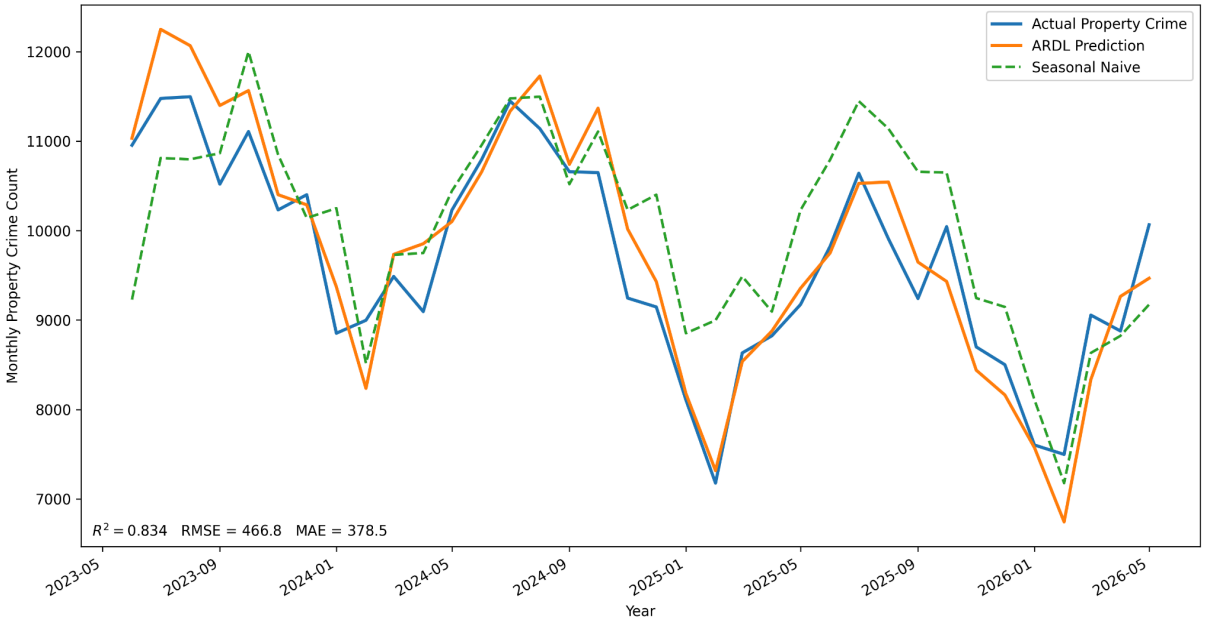
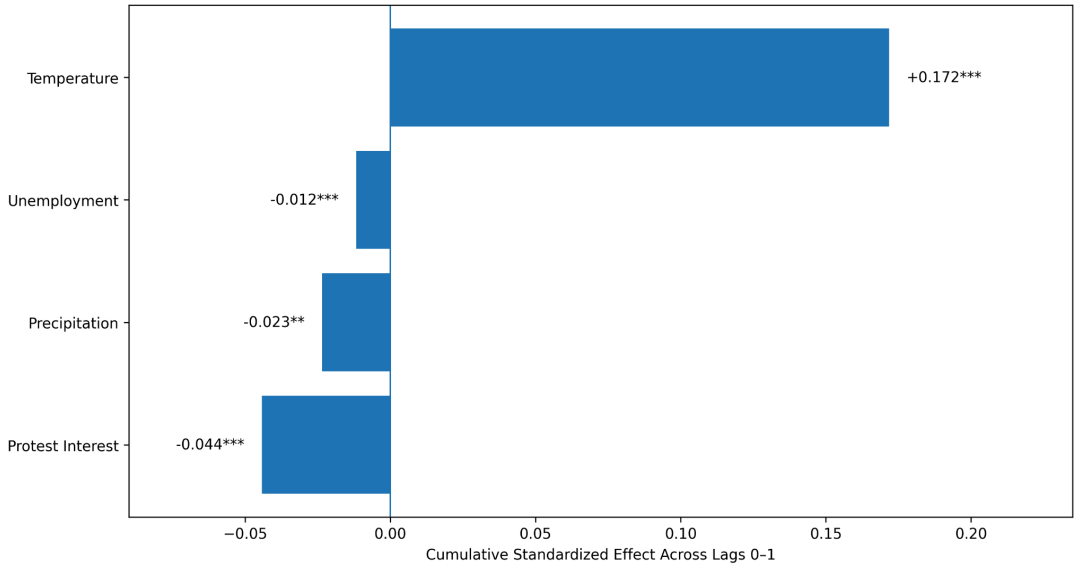


Figure 7 reports cumulative short-run effects, calculated by adding each predictor's coefficients from the current month through one month later. The significance stars are derived from joint Wald tests, which test whether the two lagged coefficients for each predictor simultaneously equal 0.

Figure 7: Short-Run Standardized Effects



Significance notation:
*** p < 0.001
** p < 0.01
* p < 0.05

ARDL Diagnostics

We used diagnostics to investigate four areas of nuance in the ARDL.

1. Covid dummy variables. As a robustness check, we added a binary variable for initial shocks from March-May of 2020 and another beginning in June 2020. The resulting model performed worse, with R^2 = .796, RMSE = 516.65, MAE = 417.80, and neither dummy being statistically significant.
2. Stationarity and bounds test. To verify that coefficients are relationships rather than coinciding trends, we used Augmented Dickey-Fuller (ADF) and KPSS tests on all variables. The results pointed to a mix of stationary and nonstationary variables, and since there was no definite evidence that any variable was integrated beyond the first order, ARDL is a generally suitable model. Property crime produced mixed results after first differencing, likely because of its strong annual seasonality; a twelve-month seasonal-difference ADF test rejected the unit-root null (p = .029). Full results can be found in Appendix 1. To further validate the existence of long-term relationships, we used a bounds test for ARDL. The test produced an F-statistic (Pesaran–Shin–Smith Case III specification) of 5.72, above the 5% I(1) upper critical value of 3.79; thus, we can reject the hypothesis and conclude that long-run relationships exist.
3. Short-run and long-run effects. While Figure 7 only accumulates short-run effects, long-term autoregressive patterns are not represented. The fifteen crime-lag coefficients sum to 0.976, meaning crime is persistent over time. Long-term effects can be calculated by dividing each predictor effect by $1 - 0.976$, which yielded highly sensitive results: +7.279 for temperature, -0.995 for precipitation, -0.498 for unemployment, and -1.876 for protest interest.
4. Residual diagnostics. To see whether errors are autocorrelated, we used a 12-lag Breusch-Godfrey test, which found no serial correlation (LM = 5.144, p = .953). A 24-lag test reinforced the

conclusion (LM=17.852, p=.810). Nonetheless, we maintained the HAC/Newey–West standard errors for robustness.

Model Discussion

Out of our three predictive models, the ARDL model achieved the highest accuracy on all metrics, so we used its output as the primary interpreter of our results.

The rest of the section interprets the standardized effect graphs for the predictors: temperature, precipitation, protest interest, and unemployment. To understand what standardized effects are in real-world circumstances, we put each variable within its historic range of variation to see how much predicted crime would change. For context, average monthly crime hovers around 11,500 occurrences during the training period.

Temperature:
The positive relationship between temperature and crime is strong and consistent at a raw cumulative effect of +30.23 crimes per 1°F. Since the monthly fixed effects already carry the seasonal pattern, the temperature coefficient is best for interpreting unusually cold or hot months, rather than changes between regular seasons. In context, a month that is 10°F warmer than usual is associated with 302 additional crimes when applying cumulative 0-1 month effects.

Precipitation:
Precipitation has a relatively small cumulative effect (−35.57 crimes per additional inch). The 10th percentile of monthly precipitation in Chicago is around 1.23 inches, while that number climbs to 6.24 inches for the 90th percentile. Our model indicates that between these months, crimes per month would drop by 178 across cumulative coefficients.

Unemployment:
Between unemployment and crime, we found a negative effect (−14.90 crimes per percentage-point increase in unemployment). According to the model, crime per month would drop by 91 instances across the two coefficients if Chicago's unemployment moves from 4.1% to 10.2%, the 10th and 90th percentiles respectively. The idea that crime decreases when the economy is doing worse is counterintuitive, and past studies have yielded findings both in support and against our results. We provide rationales for this relationship in the following subsection and derive policy implications in Recommendations.

Protest Interest:
Protest search interest has a negative cumulative effect on crime (−18.62 crimes per index point). If protest interest were to rise from the median value of 3 to the 90th percentile of 7, our model would predict 74 fewer crimes per month in contemporaneous and 1-month lagged effects. Likely, the model is capturing changes in business activity and policing, not so much the isolated effect of protests.

Unemployment Robustness Checks and Discussion

As previously discussed, unemployment (unintuitively) has a negative association with property crime. Our result is theoretically surprising because it contradicts traditional economic theories, which state that worsening labor opportunities make crime more attractive [28]. However, the relationship between unemployment and crime has been a contentious topic in academia. While a 1987 metastudy by Chiricos [29] found more positive relationships, a metastudy by Box [30] from the same year identified 19 studies that pointed to a negative effect.

The contradiction in results for empirical studies can be explained by differentiating property crime from other types of crimes. Most property offenses are committed with the intent of theft and the motivation of financial gain, unlike violence, which may be emotionally motivated. Economic downturns could deter crime because there is less opportunity—both the lack of good targets and the increase of guardianship, as identified by Cantor and Land [31]. The term effective guardianship refers to the presence of persons or mechanisms that can prevent, deter, or interrupt criminal activity [32]. When unemployment declines, a greater proportion of individuals are engaged in work outside their homes, reducing residential and vehicle occupancy during day hours [33] and increasing opportunities for property crimes.

Since the heatmap correlation between unemployment and crime is positive ($r = 0.12$, Figure 2), shuffling unemployment barely influences test errors in the Random Forest, and ARDL's negative association was small (-0.012), we conducted two tests that further validated our findings.

1. We tested our model again but used seasonally adjusted unemployment rates from FRED [34]. Since our original unemployment data reflected seasonal changes, some associations may have been falsely attributed. Our model concluded a cumulative standardized effect of -0.0121 with a p-value $< .001$, being essentially identical to the original -0.0118, which validates the model design.
2. We also re-ran the model while excluding Covid-19 periods to remove coincidental shocks in routine-activity. The first version omitted March-May of 2020, when lockdown controls were implemented, and the second version broadly omitted 2020 and 2021. Both yielded highly significant results, being -0.0180 and -0.0160 respectively. The negative association persists.

This paper is also positioned well to explore Cantor and Land's hypothesis that the increase of guardianship deters crime. If true, we would expect a contemporaneous negative opportunity effect as more people stay home and a lagged positive motivation effect as financial strain accumulates. We found results that broadly align: current month unemployment had an ARDL coefficient of -183.91 ($p < .001$), while the 1-month unemployment lag had a coefficient of $+169.01$ ($p < .001$). Our model generally affirms Cantor and Land's hypothesis, which we use later to formulate recommendations.

Forecasting Results and Discussion

To keep the forecasting models' evaluation and methods consistent with the predictor analysis models, we

used the same train-test split and compared each model with and without the calendar time trend.

For the seasonal naive baseline, we took a slightly different approach. Since the purpose is to create forecasts, we used a seasonal model with a fixed-origin, as opposed to the rolling-origin. Under this new approach, all future predictions were generated using only information available at the forecast origin, because data beyond the origin point would be from the future and hence unknown to forecasting models.

SARIMA

The goal of this section is to forecast Chicago's crime from June 2026 to December 2029. We chose a Seasonal AutoRegressive Integrated Moving Average (SARIMA), which extends upon a simple ARIMA model by adding seasonality, since we confirmed crime's seasonal variation both graphically and through our previous models.

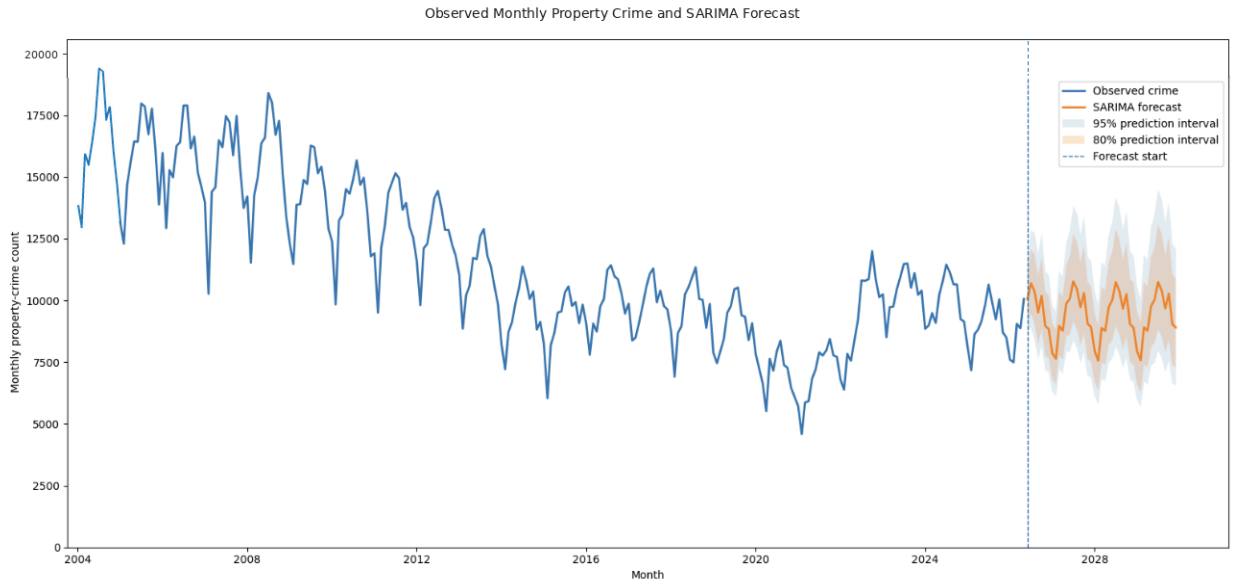
We tested variations of the SARIMA model—first with calendar time trends, and then with exogenous variables (SARIMAX), whose future values were forecasted using individual SARIMAs. Each model used a BIC algorithm to determine the best parameters and compare these versions fairly. SARIMA, by itself, had the highest out-of-sample accuracy. The best specifications were $\text{SARIMA}(0,1,3)(1,1,2)_{12}$, where notably, d=1 creates a first difference order that emphasizes recent trends, rather than long-term trends, which may include recessions and shocks of the past; q=3 creates short term persistence, P=1 adds one seasonal autoregressive term, D=1 removes yearly seasonal effects, Q=2 accounts for forecast errors at 12-month and 24-month lags, and m=12 adds monthly seasonality. This SARIMA model is fairly accurate and vastly outperforms the seasonal naive baseline, as depicted in Table 2.

Table 2: SARIMA Performance

Model	R^2	RMSE	MAE	RMSE improvement vs Seasonal Naive
Seasonal Naive	-0.077	1188.1	1008.1	-
SARIMA	0.570	751.0	673.7	36.8%

As illustrated in Figure 8, SARIMA forecasts a slight increase in property crime, from an observed annual total of 108,775 instances in 2025 to 112,438 forecasted instances in 2027. Beyond 2027, the trend levels out, with the model forecasting 111,911 occurrences of property crimes in 2028 and 112,067 in 2029. However, since prediction intervals are wide, as exhibited in Appendix Table 2, the forecasted increase is not definite. 2028 and 2029 figures are likely steady due to model design: double differencing with no drift term.

Figure 8: SARIMA Forecast Through December 2029



Police District Forecasting: Damped ETS

For our Geospatial analysis, we divided the monthly property crime data into police districts to create more informative forecasts of high risk and low risk districts. Chicago has 22 police districts numbered 1-25; districts 13, 21, and 23 had been removed by the city. We projected property crime for those districts until December 2029, a timeframe we hope would balance model accuracy and utility for recommendations. We approached this police district forecasting section with two goals in mind: (1) explore which model most accurately forecasts district crime and (2) create forecasts that could inform police allocation.

Using nested fixed-origin validation produced within the January 2004–May 2023 period , we selected the final forecasting model across wide candidates, including SARIMA, ETS, alternative transformations, ensemble methods, and hierarchical approaches. Damped Exponential Smoothing (Damped ETS) had the greatest NRMSE performance in the validation windows, and thus we selected it as the final model.

We evaluated additionally whether crime trends had cross-district spillovers. Spatial relationships were represented using row-standardized Queen-contiguity weights, which defined areas sharing a boundary or vertex as neighbors. Moran's I indicated a small positive spatial autocorrelation ($I = 0.206$, $p = 0.001$). However, adding multiple spatial specifications did not improve the accuracy of our model, so we still used the non-spatial Damped ETS model to create forecasts. Its accuracy metrics are listed in Table 3.

Table 3: Damped ETS Performance

Evaluation	Metric	Result
Monthly temporal accuracy	Median R^2	0.026
	Median RMSE	69.22
	Median MAE	56.23
	Median NRMSE	14.9%
	Districts outperforming seasonal naive by RMSE	22 of 22
Annual cross-sectional accuracy	Mean R^2	0.877
	Mean NRMSE	11.4%
	Mean Spearman ρ	0.966
	Seasonal Naive NRMSE	11.8%
	Seasonal Naive Spearman ρ	0.960

The model had a higher monthly temporal accuracy than the seasonal naive models across all districts, indicating the Damped ETS consistently adds predictive accuracy. However, the average error is 56 crimes, and with monthly district averages being 523 crimes, the model is not strong enough to confidently determine exact crime trends. Thus, we pivoted to annual cross-sectional accuracy, which captures how well our model *ranks* the 22 districts by crime, to inform police allocation. Understanding which districts will have higher crime than others can identify crime hotspots or under resourced areas.

Accuracy declined moderately the further into the forecast, with R^2 being .937 in year 1 to .815 in year 3. A Spearman correlation of 1 means a perfect ranking, so our model's 0.966 suggests modest improvement over the seasonal naive, which implies that the high accuracy is partly due to the persistence of crime rankings. For further robustness, we checked whether the district forecasts sum to a citywide forecast. The districts summed to 116,018 in 2027, 115,058 in 2028, and 114,158 in 2029, which is only +3.18%, +2.81%, and +1.87% of the citywide forecast for respective years. For some districts, the number of people exposed to property crime may not reflect population data, such as District 1, a commerce center with many daytime workers, shoppers, and tourists, but only few residents. As a result, we included crime counts alongside rates for a holistic picture. Below, Table 4 lists annual crime rate and rankings change for the top 7 ranking police districts.

Table 4: Police District Annual Crime Rate and Ranking

Police District	2025 Count	2025 Rate	2025 Rank	2029 Count	2029 Rate, 95% PI	2029 Rank	Rank Change
1	7,012	7,635	1	7,043	7,669 (2,083-9,510)	1	0
12	8,312	6,037	6	10,378	7,538 (6,353-10,212)	2	+4
3	5,177	6,584	3	5,276	6,710 (6,000-9,013)	3	0
11	3,705	5,329	8	4,580	6,588 (5,047-7,358)	4	+4
2	6,599	6,279	5	6,724	6,398 (5,024-9,085)	5	0
6	5,823	7,005	2	5,305	6,383 (5,582-8,288)	6	-4
7	3,525	6,426	4	3,385	6,171 (4,561-7,391)	7	-3

* Crime rates are per 100,000 residents.

We present findings as uncertain projections rather than definite allocation recommendations due to wide 95% prediction intervals.

Using data for each district's geographic boundaries, we created a choropleth map that shades districts based on their annual crime rates, as depicted in Figure 9. Again, note that this forecast should best be used to rank risk across districts, rather than exact increases or decreases in crime.

Figure 9: Choropleth Map of Annual Crime Rates

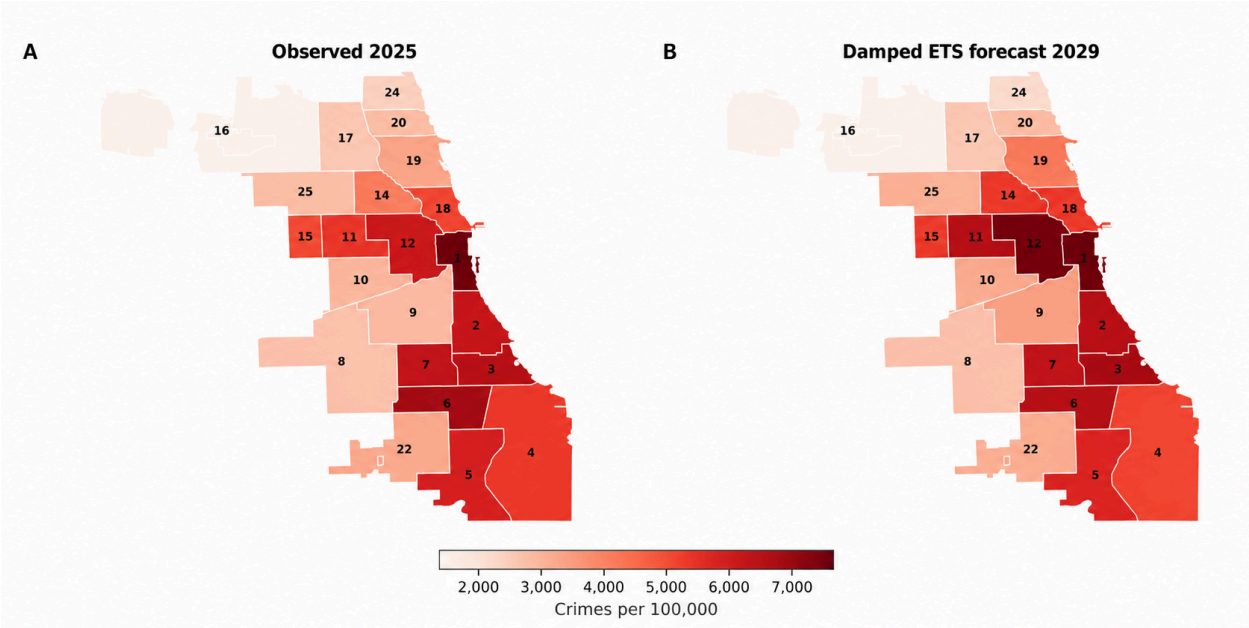


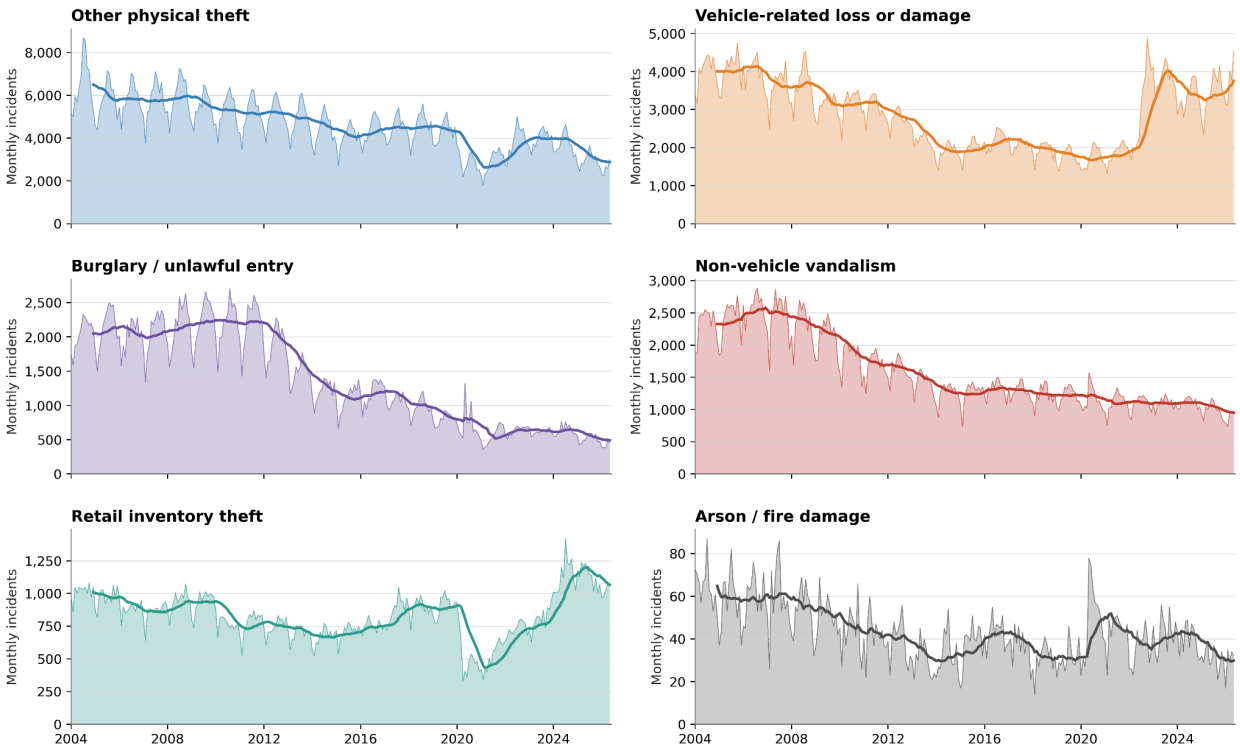
Table 4 and Figure 9 indicate that crime is concentrated in central and southern Chicago, with districts 11 and 12 seeing the most significant jumps up in the predicted rankings. These geospatial changes imply that crime within districts varies substantially over time.

RISK ANALYSIS

Risk Characterization

A major component of risk analysis and proposing mitigation strategies is understanding which location categories have the highest risk. We split property crime into six major types—vehicle-related, burglary/unlawful entry, vandalism, retail theft, arson, and other—and plotted them over time in Figure 10.

Figure 10: Chicago Property Crime by Likely Loss Exposure



While overall crime sees a downtrend over time, vehicle damage and retail theft remain prolific over time. These trends likely reflect improvements to security systems in private buildings, while public places and vehicles remain vulnerable.

Expected Value of Loss

In a 2010 study, McCollister et al. analyzed the societal cost per each type of criminal offense, which closely matches our definitions [1]. These costs are the most holistic, since they account for downstream societal effects: criminal justice costs, victim productivity losses, and monetized intangible costs such as pain, suffering, and reduced quality of life. Since values were in 2008-dollars, we multiplied costs by 1.495 to find inflation-adjusted costs in 2025, calculated by comparing 2008 and 2025 national CPI-U [35, 36]. We then calculated the total cost of Chicago's reported property crime in 2025—shown in Table 5.

Table 5: Estimated Societal Cost 2025

Crime Category	Burglary	Theft/larceny	Motor Vehicle Theft	Arson	Vandalism/ Criminal Damage	Total
2025 Observed	9,741	55,171	17,250	374	26,239	108,775

Incidents						
McCollister Cost (2008 Dollars)	$6,462	$3,532	$10,772	$21,103	$4,860	-
Estimated 2025 Societal Cost	$94,123,789	$291,380,481	$277,852,526	$11,801,704	$190,683,210	$865,841,710

Using this approach to contextualize historically, 2004 is the highest costing year in our window, with approximately $1.52 billion in 2025 dollars.

RECOMMENDATIONS

Given how property crime is predicted to remain persistent with annual losses reaching almost $1 billion, immediate and effective action can mitigate its harms. Section 5.1 covers a solution derived from our predictive analysis models, while section 5.2 uses our forecasts to identify target districts for these solutions.

Surveillance Technology

As explored in Section 3.2.5, unemployment's negative relationship with property crime supports Cantor and Land's hypothesis [31] that guardianship deters crime. Further validation in Section 3.2.6 found a contemporaneous negative effect that precedes a lagged positive effect—aligning with an expected drop in initial crime as more people stay home, which provides guardianship, and a rebound as financial strain accumulates, which incentivizes crime. These results broadly support guardianship via surveillance as a mitigation strategy, at the very least framing it as a solution worth exploring.

Surveillance functions to deter or hinder criminals by increasing their perception of the risk of getting caught, increasing the actual risk of getting caught, and denying criminals the opportunity to gain information or access about their target property. We suggest the installation of CCTV cameras in areas at risk to property crime because it maintains stable guardianship levels despite availability of natural surveillance (neighbors in residential areas, informal observation) or any change in economic conditions. CCTV cameras, also known as closed-circuit television, is a video surveillance that can autonomously monitor alarm systems, automated license plate readers, and access-control systems [37]. The government of Chicago, namely the Chicago Police Department, will be involved in the installation of the surveillance cameras, which requires funding allocation and regulatory oversight. The victims of property crime, residents, businesses, and visitors, will be positively affected by this recommendation because it actively minimizes criminal damage, increases perceived safety, and improves economic stability. An Urban Institute study [38] implemented CCTV surveillance systems in Humboldt Park and West Garfield Park, two community areas in Chicago. While Humboldt Park's crime rate decreased by 20% after the introduction of surveillance, no significant change was measured in West Garfield Park. Nevertheless, for

every $1 spent on surveillance, $4.30 in crime costs was saved, with Humboldt Park's benefits offsetting West Garfield Park's costs. These results support surveillance systems as an effective solution, and highlights how impacts differ across areas.

Furthermore, risk characterization in Section 4.1 depicts that vehicle related damage has persisted across the years. Surveillance may be especially effective at reducing this category of crime. Piza et al. [39] conducted a meta analysis on 40 years of CCTV research and found the largest and clearest benefit in reducing vehicle crime and general crime in car parks. The authors state a continued need for CCTV to be targeted on property crime and vehicle crime specifically, which has remained at large in Chicago.

Resource Allocation

Adapting proven solutions and reallocating resources into emerging crime hotspots is a cost-effective solution to property crime in Chicago. More strategic placement of intervention is an easy way to increase the effectiveness of any intervention, including the surveillance systems discussed earlier. Our forecast of district level crime rates loosely suggests that the districts moving up in rankings are as follows: 12 (+4), 11 (+4), 15(+2), 14 (+2), 9 (+2), and 25 (+2). Importantly, these estimates should highlight districts for closer monitoring rather than directly informing allocation because of uncertain prediction intervals. The rest of this section considers three additional strategies that are empirically successful, which can be significantly improved if preemptively placed in high-risk districts.

Organized Policing:

- A recommendation worth exploring is for the Chicago Police Department to hire more officers and also allocate units specifically to districts increasing in crime risk. Anticipating crime rather than reacting is key to effective intervention as it gives police officers time to gain expertise, adapt to neighborhood patterns, and specialize in their newly assigned region [40]. This recommendation is most directly supported by our analysis because the central contribution of our forecasting model—identifying emerging high-risk districts before future crime is observed—is an important prerequisite for effective organized policing.
- Our district-level forecasts can be used to identify districts with persistently high crime, which can undergo further study to locate hotspots for crime. Hotspot policing has been effective: Braga et al. [40] systematically reviewed 65 studies and 78 test hotspots and found that 62 of the 78 saw reductions in crime with little spillover to surrounding areas. Raising officer salaries is the most common method to expand the police force, as higher wages are more competitive with opportunity costs [41]. Chalfin and McCrary [42] use panel data of large US cities from 1960 to 2010 to support how the benefits outweigh the costs, as each $1 spent on the police in general reduces victimization costs by $1.60.
- Importantly, there are a few commonly cited problems with organized policing. When more officers are deployed, more crimes also get reported, which Lum and Isaac [43] and Ensign et al. [44] describe as a feedback loop where predictive policing justifies itself. Another concern is overpolicing and privacy infringements for many already vulnerable communities. These risks should be considered and studied before implementing organized policing.

Damage-limiting designs:

- Damage-limiting designs aim to reduce the severity of property crime by installing engineering or architectural modifications such as anti-shatter material on glass, anti-graffiti coatings, and crime prevention through environmental design (CPTED). Historically, security window film and polycarbonate glazing panels have been used in New York during the 2020 riots [45]. Although designs had limited mitigation effects, improvements are still worth consideration. Our forecasts can help increase the effectiveness of damage-limiting designs, since they take a long time to plan and should best be placed preemptively of crime surges.

Community intervention programs:

- Community intervention programs, health or employment programs that identify and provide individual support to motivated offender pools, are another common option to mitigate crime risks. Since these programs build upon pre-existing public health infrastructure, they are scalable to high-density districts. The most evidence-backed intervention program in Chicago is the One Summer Chicago Plus program in 2012 [46]. Evaluations of the program have found reductions in arrests for participants, but evidence is far stronger for reducing violence arrests rather than for criminal damage [47].

LIMITATIONS

1. Results are associational rather than causal. Future research can use causal models to better justify Cantor and Land's hypothesis and derive recommendations. Another research area unlocked by causal inference is studying how policing affects property crime, which we had to omit in this paper because we couldn't resolve reverse causality—policing can affect crime but so can crime affect policing.
2. Data limitations. The unemployment data comes from the Chicago metropolitan area rather than the city itself, which could result in small prediction errors. With more local data available and the use of a causal approach, researchers can study with more accuracy the predictors of property crime and the policy solutions we need for public safety. Additionally, the protest interest data can be improved. Because Google Trends scales values from 0 to 100 relative to the highest data point, most values are close to the bottom of the scale, which means variation is dominated by a few spikes and integer rounding introduces inaccuracies. Searches for "protest" in Chicago may be a result of national and international headlines, rather than local events. Future studies should consider datasets that isolate Chicago or an event-based methodology. Our paper also removed the Covid-19 variable since the dataset stopped coverage after 2024; our next option, pandemic period dummy variables, were not significant and harmed model performance. With a more comprehensive proxy for Covid-19, future research can explore the relationship between property crime and external shocks such as pandemics.
3. The calendar time trend was strongly correlated with several variables in the heatmap, having a Pearson Correlation of -.79 with property crime. Although correlation is expected given the time

trend's purpose is to capture long-term direction of variables, it can cloud interpretability. Predictor effects may become hard to separate, so coefficients are mere estimates.

4. The Damped ETS model, which forecasts district-level crime, has accuracy limitations. Although the model marginally outperformed the naive benchmark in terms of annual cross-sectional accuracy, the prediction intervals revealed much uncertainty in forecasting district rankings. Moreover, our district rates assumed populations constant through 2029, which doesn't account for population shifts across the city. Because of these concerns, the resource allocation proposals in Recommendations must be interpreted cautiously.
5. Finally, our forecasts suppose that historical relationships hold true in the future under Assumption 4. However, our model cannot anticipate economic, legislative, or public safety shocks. For example, Illinois' Pretrial Fairness Act [48], which ended cash bail, was enacted in September 2023. A national surge in Kia and Hyundai vehicle thefts [49] coincided with a huge 2022 spike in Chicago's property crime, showing how unaccounted shocks can influence crime. The uncertainty propagates as our model forecasts further into the future.

CONCLUSION

Ultimately, our paper successfully answers the research questions outlined in the Abstract and reiterated in Data Methodology. First, we found four variables able to significantly predict property crime: temperature had a positive association, while precipitation, unemployment, and protest interest had negative associations. Second, our models forecast monthly Chicago crime with moderate accuracy ($R^2 = 0.570$), monthly police-district crime with weak accuracy ($R^2 = 0.026$), and annual police-district crime rankings with strong accuracy (Spearman $\rho = 0.966$), but only modest improvement over the baseline. Third, our model forecasts the top four districts with the highest crime per capita in 2029 to be 1, 12, 3, and 11, with districts 12 and 11 predicted to climb four ranks. Our findings for the unemployment predictor parallels the theory that guardianship deters crime, although our method doesn't imply causation. We proposed recommendations that used forecasted police-district risk to better allocate resources and preemptively place interventions. Such improvements are immediate and actionable ways to bolster surveillance technology, policing, damage-limiting designs, and community intervention programs.

REFERENCES

[1] McCollister, Kathryn E., Michael T. French, and Hai Fang. "The Cost of Crime to Society: New Crime-Specific Estimates for Policy and Program Evaluation." *Drug and Alcohol Dependence*, vol. 108, no. 1–2, 2010, pp. 98–109. https://pmc.ncbi.nlm.nih.gov/articles/PMC2835847/

[2] Bureau of Justice Statistics. *Criminal Victimization, 2022*. U.S. Department of Justice, 2023. https://bjs.ojp.gov/press-release/criminal-victimization-2022

[3] Wang, X. "Exploring the Impact of Socioeconomic and Built Environment Factors on Theft Crime Trends in Chicago During the COVID-19 Pandemic." *Journal of Applied Economics and Policy Studies*, vol. 14, no. 1, 2024, pp. 21–30. https://doi.org/10.54254/2977-5701/2024.17977

[4] CBS Chicago. *Graffiti Removal Blasters Streets and Sanitation*. CBS News Chicago. https://www.cbsnews.com/chicago/news/graffiti-removal-blasters-streets-and-sanitation/

[5] ABC7 Chicago. *Chicago Crime: Car Theft and Vehicle Crime Trends*. ABC7 Chicago. https://abc7chicago.com/post/chicago-crime-car-theft-vehicle-police-department/14212073/

[6] Chicago Metropolitan Agency for Planning. *Regional Economy and Clusters*. 2017. https://cmap.illinois.gov/wp-content/uploads/Regional-Economy-and-Clusters-snapshot.pdf

[7] Bartik, Alexander W., et al. "The Impact of COVID-19 on Small Business Outcomes and Expectations." *Proceedings of the National Academy of Sciences*, vol. 117, no. 30, 2020, pp. 17656–17666. https://doi.org/10.1073/pnas.2006991117

[8] City of Chicago Office of Inspector General. *Report on Chicago's Response to George Floyd Protests and Unrest*. 18 Feb. 2021. https://igchicago.org/wp-content/uploads/2023/08/OIG-Report-on-Chicagos-Response-to-George-Floyd-Protests-and-Unrest.pdf

[9] Office of Governor JB Pritzker. "Gov. Pritzker Activates Illinois National Guard at Request of City of Chicago." *State of Illinois*, 31 May 2020. https://www.illinois.gov/news/release.html?releaseid=21617

[10] City of Chicago. *Crimes – 2001 to Present (About Data)*. City of Chicago Data Portal. https://data.cityofchicago.org/Public-Safety/Crimes-2001-to-Present/ijzp-q8t2/about_data

[11] Federal Bureau of Investigation. (n.d.). National Incident-Based Reporting System (NIBRS). https://www.fbi.gov/how-we-can-help-you/more-fbi-services-and-information/ucr/nibrs

[12] Illinois General Assembly. (n.d.). 720 ILCS 5/21-1: Criminal damage to property. Illinois Compiled Statutes. https://www.ilga.gov/Documents/legislation/ilcs/documents/072000050K21-1.htm

[13] Federal Reserve Bank of St. Louis. *Unemployment Rate in Chicago-Naperville-Elgin, IL-IN-WI (CHIC917URN)*. FRED. https://fred.stlouisfed.org/series/CHIC917URN

[14] U.S. Census Bureau. *Estimate of People of All Ages in Poverty in Cook County, IL (PEAAIL17031A647NCEN)*. FRED. https://fred.stlouisfed.org/series/PEAAIL17031A647NCEN

[15] U.S. Census Bureau. *Estimate of Median Household Income for Cook County, IL (MHIIL17031A052NCEN)*. FRED, Federal Reserve Bank of St. Louis. https://fred.stlouisfed.org/series/MHIIL17031A052NCEN

[16] Google. *Google Trends: "Protest" Search Interest, Chicago DMA.* Google Trends. https://trends.google.com/explore?q=protest&date=all&geo=602

[17] Timoneda, Joan C., and Erik Wibbels. "Spikes and Variance: Using Google Trends to Detect and Forecast Protests." *Political Analysis*, vol. 30, no. 1, 2022, pp. 1–18. https://www.cambridge.org/core/journals/political-analysis/article/abs/spikes-and-variance-using-google-trends-to-detect-and-forecast-protests/952C6EB342A7CC71A7D079F1A5624A5D

[18] Federal Reserve Bank of St. Louis. *Resident Population in Cook County, IL (ILCOOK1POP).* FRED. https://fred.stlouisfed.org/series/ILCOOK1POP

[19] National Centers for Environmental Information. *Climate at a Glance: City Time Series.* NOAA. https://www.ncei.noaa.gov/access/monitoring/climate-at-a-glance/city

[20] City of Chicago. *Boundaries – Police Districts (current).* Chicago Data Portal, 2024. https://data.cityofchicago.org/d/fthy-xz3r

[21] U.S. Census Bureau. *Sex by Age.* American Community Survey, 2024 ACS 5-Year Estimates Detailed Tables, Table B01001, 2024. https://data.census.gov/table/ACSDT5Y2024.B01001

[22] Graziul, Christopher, Drew Keller, and Margaret Beale Spencer. *CPD Beat to Census Block Crosswalk.* Zenodo, 2024. https://doi.org/10.5281/zenodo.18111458

[23] Cohn, Ellen G. "Weather and Crime." *The British Journal of Criminology*, vol. 30, no. 1, 1990, pp. 51–64. https://academic.oup.com/bjc/article-abstract/30/1/51/523870

[24] City of Chicago. *COVID-19 Daily Cases, Deaths, and Hospitalizations (About Data).* City of Chicago Data Portal. https://data.cityofchicago.org/Health-Human-Services/COVID-19-Daily-Cases-Deaths-and-Hospitalizations-H/naz8-j4nc/about_data

[25] Khan, Muhammad, et al. "Ridge Regression Analysis: A Review." *American Journal of Applied Mathematics and Statistics*, vol. 8, no. 2, 2020. https://pubs.sciepub.com/ajams/8/2/1/index.html

[26] Raphael, Steven, and Rudolf Winter-Ebmer. "Identifying the Effect of Unemployment on Crime." The Journal of Law and Economics, vol. 44, no. 1, 2001, pp. 259–283. https://doi.org/10.1086/320275

[27] Stucky, Thomas D., Seth B. Payton, and John R. Ottensmann. "Intra- and Inter-Neighborhood Income Inequality and Crime." Journal of Crime and Justice, vol. 39, no. 3, 2016, pp. 345–362. https://doi.org/10.1080/0735648X.2015.1004551

[28] Ehrlich, Isaac. "Participation in Illegitimate Activities: A Theoretical and Empirical Investigation." Journal of Political Economy, vol. 81, no. 3, 1973, pp. 521–565. https://doi.org/10.1086/260058

[29] Chiricos, Theodore G. "Rates of Crime and Unemployment: An Analysis of Aggregate Research Evidence." Social Problems, vol. 34, no. 2, 1987, pp. 187–212. https://doi.org/10.2307/800715

[30] Box, S. *Recession, Crime and Punishment*. Macmillan Press, 1987. https://doi.org/10.1007/978-1-349-18784-3

[31] Cantor, D., and Land, K. C. "Unemployment and Crime Rates in the Post-World War II United States: A Theoretical and Empirical Analysis." *American Sociological Review*, vol. 50, no. 3, 1985, pp. 317–332. https://doi.org/10.2307/2095542

[32] Hollis-Peel, Meghan E., et al. "Guardianship for Crime Prevention: A Critical Review of the Literature." *Crime, Law and Social Change*, vol. 56, no. 1, 2011, pp. 53–70. https://doi.org/10.1007/s10611-011-9309-2

[33] D'Alessio, S. J., Eitle, D., and Stolzenberg, L. "Unemployment, Guardianship, and Weekday Residential Burglary." *Justice Quarterly*, vol. 29, no. 6, 2012, pp. 919–932. https://doi.org/10.1080/07418825.2011.605073

[34] U.S. Bureau of Labor Statistics. Unemployment Rate in Chicago-Naperville-Elgin, IL-IN-WI (MSA) [CHIC917UR]. FRED, Federal Reserve Bank of St. Louis, https://fred.stlouisfed.org/series/CHIC917UR

[35] U.S. Bureau of Labor Statistics. *Consumer Price Index: December 2008.* U.S. Department of Labor, 16 Jan. 2009. https://www.bls.gov/news.release/archives/cpi_01162009.pdf

[36] U.S. Bureau of Labor Statistics. *Summary of Annual and Semi-Annual Indexes.* Mid-Atlantic Information Office. https://www.bls.gov/regions/mid-atlantic/data/consumerpriceindexannualandsemiannual_table.htm

[37] Space and Naval Warfare Systems Center Atlantic. *CCTV Technology Handbook*. U.S. Department of Homeland Security, July 2013. https://digitalcommons.csp.edu/cgi/viewcontent.cgi?article=1047&context=cup_commons_faculty

[38] La Vigne, Nancy G., et al. *Evaluating the Use of Public Surveillance Cameras for Crime Control and Prevention.* Urban Institute, 2011. https://www.urban.org/sites/default/files/publication/27546/412401-Evaluating-the-Use-of-Public-Surveillance-Cameras-for-Crime-Control-and-Prevention-A-Summary.PDF

[39] Piza, Eric L., Brandon C. Welsh, David P. Farrington, and Amanda L. Thomas. "CCTV Surveillance for Crime Prevention: A 40-Year Systematic Review with Meta-Analysis." Criminology & Public Policy, vol. 18, no. 1, 2019, pp. 135–159. https://doi.org/10.1111/1745-9133.12419

[40] Braga, Anthony A., and David L. Weisburd. "Does Hot Spots Policing Have Meaningful Impacts on Crime?" *Criminology & Public Policy*, 2023. https://pubmed.ncbi.nlm.nih.gov/37133274/

[41] National Institute of Justice. *Hiring and Keeping Police Officers.* U.S. Department of Justice. https://nij.ojp.gov/library/publications/hiring-and-keeping-police-officers

[42] Chalfin, Aaron, and Justin McCrary. *Are U.S. Cities Underpoliced? Theory and Evidence.* 2012. https://eml.berkeley.edu/~jmccrary/chalfin_mccrary2012.pdf

[43] Lum, Kristian, and William Isaac. "To Predict and Serve?" *Significance*, vol. 13, no. 5, 2016, pp. 14–19. https://doi.org/10.1111/j.1740-9713.2016.00960.x

[44] Ensign, D., Friedler, S. A., Neville, S., Scheidegger, C., & Venkatasubramanian, S. (2018). Runaway feedback loops in predictive policing. *Proceedings of the 1st Conference on Fairness, Accountability and Transparency, 81*, 160–171. https://proceedings.mlr.press/v81/ensign18a.html

[45] Thomas, Lauren. "U.S. Retailers Secure Stores as Worries About Election Unrest Mount." *CNBC*, 21 Oct. 2020. https://www.cnbc.com/2020/10/21/us-retailers-secure-stores-as-worries-about-election-unrest-mounts.html

[46] University of Chicago Crime Lab. *One Summer Chicago Plus.* University of Chicago Crime Lab. https://crimelab.uchicago.edu/projects/one-summer-chicago-plus/

[47] Heller, Sara B., et al. *Summer Jobs Reduce Violence Among Disadvantaged Youth.* Office of Juvenile Justice and Delinquency Prevention, U.S. Department of Justice, 2017. https://www.ojp.gov/pdffiles1/ojjdp/grants/251101.pdf

[48] State of Illinois Office of the Illinois Courts. "Illinois Supreme Court's Comprehensive Preparations for the End of Cash Bail." 27 Sept. 2023. https://www.illinoiscourts.gov/News/1279/Illinois-Supreme-Courts-comprehensive-preparations-for-the-end-of-cash-bail/news-detail/

[49] Illinois Secretary of State, Motor Vehicle Theft Prevention and Insurance Verification Council. *2024 Annual Report*. https://www.ilsos.gov/content/dam/mvtpiv/pdf/24annualreport.pdf

APPENDIX

Table A1: ADF and KPSS Stationarity Test Results

Variable	ADF at Level	KPSS at Level	After First Difference	Conclusion
Property Crime	.247	<.010	Mixed; KPSS supports stationarity	Mixed evidence; Likely I(1)

Temperature	.050	>.100	Stationary	At most I(1)
Precipitation	.020	>.100	Stationary	I(0)
Unemployment	.045	.071	Stationary	I(0)
Protest Interest	<.001	.46	Stationary	I(0)

Table A2: SARIMA Forecast Performance for Chicago-wide Crime

Year	Forecast Avg Monthly Crime	Forecast Annual Total	Average Lower 80%	Average Upper 80%	Average Lower 95%	Average Upper 95%
2027	9,370	112,438	8,659	10,261	8,259	10,712
2028	9,326	111,911	8,521	10,411	8,101	10,931
2029	9,339	112,067	8,448	10,502	7,934	11,134

*Crime figures are rounded to the nearest whole number

Table A3: Damped ETS Monthly Temporal Accuracy (All 22 districts)

District	R^2	ETS RMSE	ETS MAE	Seasonal-naive RMSE	RMSE improvement (%)
1	0.625	61.639	49.307	131.992	53.301
2	0.304	68.970	57.131	96.636	28.629
3	-0.519	76.344	65.824	102.412	25.454
4	0.053	70.801	58.933	106.697	33.642
5	-1.698	70.869	62.232	86.775	18.330
6	-0.864	93.487	82.586	119.487	21.760
7	-1.577	72.414	65.291	81.691	11.356
8	0.427	64.389	53.157	89.917	28.391
9	0.270	45.360	31.317	46.255	1.936

10	-0.601	63.474	55.231	70.213	9.598
11	-0.391	72.281	63.497	109.381	33.918
12	0.532	69.468	54.710	116.162	40.198
14	0.201	64.629	54.664	81.078	20.287
15	-1.634	48.126	41.279	94.907	49.292
16	-0.630	66.170	55.327	71.477	7.425
17	0.360	33.508	27.118	44.489	24.684
18	0.654	49.209	40.732	59.579	17.406
19	0.448	81.274	63.544	83.360	2.503
20	0.352	28.967	21.510	46.606	37.846
22	-3.246	82.962	77.792	91.748	9.577
24	-0.001	76.016	67.228	83.365	8.815
25	-1.029	72.867	64.234	86.158	15.427

Table A4: Damped ETS Annual Cross-Sectional Accuracy

Forecast Period	R^2	NRMSE	Seasonal Naive NRMSE	NRMSE improvement vs seasonal baseline	Spearman ρ
Year 1	0.937	7.650%	7.651%	0.014%	0.956
Year 2	0.879	11.479%	11.872%	3.314%	0.975
Year 3	0.815	15.069%	15.787%	4.545%	0.956

Table A5: Full Damped ETS Crime Rate and Ranking

Police District	2025 Count	2025 Rate	2025 Rank	2029 Count	2029 Rate, 95% PI	2029 Rank	Rank Change

1	7,012	7,635	1	7,043	7,669 (2,083–9,510)	1	0
12	8,312	6,037	6	10,378	7,538 (6,353–10,212)	2	+4
3	5,177	6,584	3	5,276	6,710 (6,000–9,013)	3	0
11	3,705	5,329	8	4,580	6,588 (5,047–7,358)	4	+4
2	6,599	6,279	5	6,724	6,398 (5,024–9,085)	5	0
6	5,823	7,005	2	5,305	6,383 (5,582–8,288)	6	−4
7	3,525	6,426	4	3,385	6,171 (4,561–7,391)	7	−3
5	3,825	5,993	7	3,599	5,638 (4,391–6,168)	8	−1
15	2,750	4,966	11	3,005	5,426 (4,518–6,854)	9	+2
14	4,841	4,095	12	6,248	5,284 (3,361–6,228)	10	+2
18	7,417	5,120	10	7,454	5,146 (2,370–6,676)	11	−1
4	6,079	5,319	9	5,654	4,947 (4,477–6,819)	12	−3
19	7,234	3,327	13	8,363	3,846 (3,143–4,863)	13	0
9	4,580	2,847	16	5,372	3,339 (2,712–3,938)	14	+2
10	3,053	2,967	15	3,202	3,111 (2,530–3,913)	15	0
25	5,276	2,771	18	5,874	3,086 (2,264–3,688)	16	+2
22	3,363	3,286	14	3,155	3,083 (1,853–4,043)	17	−3

20	2,610	2,843	17	2,620	2,854 (2,628–3,628)	18	−1
8	6,654	2,687	19	6,666	2,692 (1,962–3,254)	19	0
17	3,650	2,612	20	3,687	2,638 (1,749–3,327)	20	0
24	3,580	2,440	21	3,078	2,098 (1,883–3,079)	21	0
16	3,710	1,454	22	3,490	1,368 (1,051–1,944)	22	0