# Police Use of Force and Protest Dynamics: A Time Series Analysis Using Crowd Counting Consortium Data

Blinded version for review

**Abstract**

This paper analyzes the Crowd Counting Consortium (CCC) data set, which documents protest events and corresponding police responses in the United States. We find that police use of force—specifically rubber bullets and chemical agents like tear gas and pepper spray—is associated with larger and more violent protests in subsequent days. We build time series models for the number of protests, the number of protesters, injuries to both protesters and police, and property damage. Across all outcomes, our findings show that increased use of force by police is statistically associated with higher levels of protest activity and harm in subsequent days. Furthermore, the impact of police actions on the number of protesters exhibits a "long memory" phenomenon, suggesting persistent structural effects. These results provide a statistical foundation for future work on the dynamics of collective action and offer new insights to police deciding between the escalated force approach and the negotiated management approach to protests.

## 1 Introduction

The United States has a long and complicated relationship with protest movements. While the right to assemble and express dissent is enshrined in the First Amendment, and protests often serve the important function of safely diffusing social tension [1], the state's response to protest has often involved the use of coercive policing tactics. From civil rights marches to contemporary demonstrations against racial injustice and police violence, protest events in the United States are often shaped not only by the grievances of the protesters but also by the actions taken by law enforcement. Understanding the consequences of those policing strategies is central to both public policy and the broader study of social movements.

Recent studies have highlighted the potential for police use of force to escalate rather than suppress protest activity. For example, analysis of the Armed Conflict Location & Event Data Project data set (ACLED) in the context of the 2020–2021 U.S. protests found that the use of kinetic impact projectiles (KIPs) by police was followed by increases in the number of protests and in protest-related deaths [1, 23]. Similar conclusions emerged from studies of data from the Center for Social and Labor Research in Ukraine [24], where forceful state responses during the 2013 Euromaidan protests in Ukraine were associated with sustained protest escalation and broader political unrest [2].

In this paper, we build on that work by analyzing a new source of empirical data: the Crowd Counting Consortium (CCC) data set [22], covering protests in the United States from 2017 through 2024 and including event-level information on protest activity and police responses, beyond what is present in the ACLED data set. We carry out time series analyses to replicate the findings from [1, 2] regarding the association between the use of rubber bullets, crowd injuries, police negativity, and subsequent increases in protest activity. We then expand the scope of inquiry in several important ways.

First, we examine the effects of chemical agents (e.g., pepper spray and tear gas) on protest dynamics. Tear gas has been widely used by law enforcement in recent protest contexts, yet its effects on protest escalation remain understudied in the applied statistics literature. Second, we broaden the set of outcomes beyond the number of protests to include the number of protesters, data that was largely absent in the data sets studied in [1, 2]. Third, we study whether police use of force (e.g., KIPs, chemical agents, mass arrests, and beatings) are associated with an increase in the violence of protests as measured by the number of reported injuries in subsequent days. We study this both for reported injuries to protesters and reported injuries to police officers. In every case, we find that greater use of force by police is associated with increases in these outcomes over the days that follow. Fourth, we study the relationship between police use of force and property damage. Here we find that property damage both precedes and follows police use of force, each potentially driving the other. The use of KIPs is associated with increases in property damage (both in the same day and in the next day), while the use of chemical agents is associated with less property damage the next day, but more in subsequent days.

Methodologically, our study employs time series analysis [25]. As in [1, 2], we fit seasonal autoregressive integrated moving average (SARIMA) models. We found that each time series exhibited long-memory behavior, and we additionally employed models using fractional differencing (ARFIMA). Such models were originally developed in geophysical settings such as the study of glacial varves, and are designed to capture persistent dependence over long time horizons [25]. Our results suggest that the use of force by police can have enduring effects on protest dynamics, contributing to a cycle of escalation that is not easily diffused.

We now provide an overview of the paper. In Section 2, we review related work. In Section 3 we discuss the CCC data and how we extracted our time series from it. To our knowledge, only one previous academic article has used the CCC data source [30], so we conduct an extensive exploratory data analysis in Section 3.1 and record precisely what this data set contains. In Section 4 we describe our methodology. In Section 5 we provide our results. We include a changepoint analysis (Section 5.1), univariate time series models for each of our time series to determine how they depend on their past (Section 5.2), and time series regression models to study the relationships between them (Section 5.3).

This paper provides an empirically grounded statistical perspective on the dynamics of protest and policing, to establish a basis for future interdisciplinary research on collective action, crowd behavior, and state response. In future work, we hope to employ mathematical models (e.g., extending the Hawkes process and agent-based models of [1, 2]) to capture the dynamics and simulate potential futures.

## 2 Related Works

Our work builds on two prior papers that analyze protest dynamics using time series methods and spatiotemporal modeling. Building on the methodology of [3], in [1], the authors used SARIMA models to analyze the effect of police repression techniques—particularly the use of rubber bullets—on the number of protests and on the number of deaths during protest events. They also fit agent-based models of social tension to simulate different policing strategies and subsequent changes in protest dynamics. They found that forceful police tactics were associated with increases in protest activity in subsequent days, a finding that provides evidence against the escalated force strategy of policing. In [2], the authors focused on the 2013–2014 Euromaidan protests in Ukraine, building from [5], and again using SARIMA models to relate negative response events (such as police beatings of protesters or mass arrests) to the number of protests. They also applied Hawkes process models to study the self-excitation of protests, the spatiotemporal spread of protest activity, and the role of political rather than spatial proximity in diffusion dynamics. That work again found that negative response events by police were associated with more protests in the subsequent days. The authors connected this research to protests in the Republic of Georgia in [8].

These studies reflect growing academic interest in the statistical modeling of protest and repression. Within sociology, scholars have long analyzed how social tension and conflict unfold in collective action settings. For example, [26, 27] emphasize the buildup of latent emotional pressure prior to protest, while [28] argues that protest can serve as a socially regulated release of such tension.

A parallel literature in criminology focuses on how law enforcement strategies shape protest outcomes. The two main strategies of policing are the escalated force approach and the negotiated management approach [6]. The escalated force approach involves using progressively more aggressive tactics like pepper spray, rubber bullets, tear gas, beatings, and mass arrests [7]. The negotiated management approach emphasizes communication and cooperation between law enforcement and protesters, with designated spaces for safe dialogue with protest organizers, built-in de-escalation periods, and agreement between protest organizers and police regarding the protest route, acceptable noise levels, and how to handle arrests [8]. In the history of the United States, periods where police followed the escalated force model (including the Civil Rights Era, the 1968 Democratic National Convention (DNC), the Rodney King riots, the Ferguson protests, and the George Floyd protests in many cities) have tended to be associated with larger and more violent protests subsequent to police escalation, while periods featuring the negotiated management approach (the 1970s and 1980s, the 2012 DNC, Newark NJ during the George Floyd protests) have tended to see protests cool as social tension is successfully diffused, and no escalation of violence [9]. The same is true for protests internationally [10]. For this reason, researchers have called for a shift to the negotiated management approach [11], and the U.S. Department of Justice echoed this finding in the post-Ferguson era [12].

Quantitative studies of protest-police interaction have also advanced. For example, [13] use time series methods to assess the short-term impact of police repression, and [14] studied the militarization of police under the 1033 program, finding increased use of deadly force. Studies of psychological outcomes, including [15, 16], document the mental health toll on both protesters and community members, especially when violence or repression is involved.

Finally, mathematical modeling approaches have gained traction in the study of unrest. For example, [4, 17, 18, 19] use diffusion-based and contagion-based models to explain riot dynamics, while other scholars have adapted frameworks from epidemiology, game theory, and nonlinear systems to understand protest escalation and participation (e.g., [20, 21, 29]). Our current work differs in focusing squarely on statistical models of protest dynamics using the new CCC data set. By incorporating both SARIMA and ARFIMA models, we show that police use of force not only has short-term associations with protest size and violence, but also exhibits long-memory effects, suggesting that recent force deployment shapes protest dynamics over extended periods. This finding adds a new dimension to the literature on protest-police interactions and strengthens the case for modeling these systems with tools developed for persistent processes in other domains.

# 3 Data acquisition and cleaning

We analyzed data from the Crowd Counting Consortium (CCC), an initiative of the Nonviolent Action Lab at the Ash Center for Democratic Governance and Innovation at Harvard University [30]. The CCC data set is curated and maintained by a team of scholars and research fellows at Harvard, and it compiles detailed records of protest events across the United States [31]. The cleaned and compiled data set was publicly available through the Nonviolent Action Lab's GitHub repository [22], and spans from January 1, 2017 until December 8, 2024. We downloaded the data on February 9, 2025 and it consisted of two comma-separated value files: one covering 2017–2020 and one for 2021–2024. All data wrangling and transformation were conducted in R.

Each of the 210,186 rows in the data set corresponds to a protest-related event, with 70 columns providing detailed information about the location, date, actors involved, claims made, and outcomes such as crowd size, police and participant injuries, property damage, and arrests. Missing data is marked `NA`. Columns of interest for our analysis included:

- **date:** the date on which the event happened.
- **size_mean:** the number of participants/protesters of each event. This is the average of `size_low` and `size_high`. We have 127,371 rows with NAs in the `size_mean` column, more than half of all rows. The data appears to be missing at random.
- **arrests:** the number of people arrested in an event. This variable is of type character. Most rows contain a number (expressed as a string) but some contain text (e.g., "dozen" or "numerous"). We conservatively converted vague descriptors to approximate numeric values (e.g., "multiple" → 2, "at least a few dozens" → 24). There is also `arrests_any`, which is 1 if any arrests occurred and 0 otherwise.
- **injuries_crowd:** the number of protesters injured. This variable was originally a character type. After conversion to numeric, we followed the same conservative procedure. There are 166,509 NAs (about 79%).
- **injuries_police:** the number of police injured at this event. There are 166,574 NAs (about 79%).

- **property_damage_any:** a binary variable (0 = no evidence of property damage, 1 = some property damage). The variable `property_damage` is a string indicating whether damage occurred and, in data from 2021 onward, sometimes describes the type of damage. It has 165,736 NAs (about 79%).
- **notes:** qualitative information about the event, including whether police used pepper spray, tear gas, rubber bullets, etc. There are 54,655 NAs (about 25%). Missingness appears random.
- **chemical_agents:** a binary indicator for police use of tear gas, pepper spray, or other irritants. Derived via natural language processing of `police_measures` by the CCC team. It has 59,662 NAs (about 26%). To be conservative, NA values were set to 0 (no evidence of use).

Unlike the data analyzed in [1, 2], the CCC data includes protest size estimates, detailed coding of police tactics, and new variables such as property damage, chemical agent usage, and participant injuries. These enhancements allowed us to explore richer models of protest escalation, policing response, and outcomes. Because of missing data issues, whenever we say "number of injuries" we mean "reported number of injuries." Text fields such as `notes` were further processed using regular expressions to extract consistent numerical indicators for specific police responses such as the use of KIPs.

## 3.1 Exploratory Data Analysis

Of the 70 variables in the CCC data set, 30 are fields holding sources (e.g., links to news articles) verifying the information in each row. In addition to the variables listed above, the data set has:

- **lat, lon:** geographic coordinates, which we plan to use in future work on spatiotemporal modeling [32].
- **locality, state, location_detail, resolved_locality, resolved_county, fips_code, resolved_state:** location descriptors based on geocoding.
- **type:** protest type (march, strike, counter-protest, etc.).
- **valence:** partisan lean of the protest theme (0 = neither, 1 = anti-Trump, 2 = pro-Trump for 2017–2020). Missing for only 62 events.
- **macroevent:** links counter-protests to the protest(s) they opposed.

We plan to use the latter two variables in future agent-based models of protests and counterprotests.

Variables excluded due to missingness include: `police_measures` (199,405 NAs; not recorded before 2021), `participant_measures` (173,317 NAs), `participant_deaths` (19 total, with known missingness), and `police_deaths` (1 total, also incomplete). Arrests data had 165,026 NAs (78%), consistent with [2].

Several variables not used in our analysis may interest future researchers, e.g., text descriptions of property damage, protest claims and issues, or which organizations were involved.

Summary statistics show:

- Median protest size: 55 people; mean: 621; maximum: 1.5 million; SD: 13,309.
- Mean arrests: 1.09 per protest; maximum: 1265.
- Mean crowd injuries: 0.04; maximum: 121.
- Mean police injuries: 0.02; maximum: 140.

On average, more protesters are injured than police. Distributions of protest size and protest count were highly skewed, motivating log-transformations.

## 3.2 Creation of the time series

We created a new column, `kip`, equal to TRUE if kinetic impact projectiles (rubber bullets, foam rounds, bean bag rounds, pepper balls, or munitions) were used. The strings to match precisely followed the analysis of [1], so that we can replicate its findings in Section 5. We searched for these strings in `police_measures` (available only 2021–2024) and `notes`.

We next created `police_negativity`, which is TRUE for an event if any of the above strings are found, as well as strings for tactical gear, riot gear, barricade, baton, bullet, flash bangs, grenade, mace, pepper spray, or tear gas. This aligns with the "negative response" variable in [2], but in the Ukraine data negative responses also included things like a protester being fired, harassed after a protest ended, or any confrontation with police. We don't have access to that kind of post-protest data in the CCC and the CCC data cannot distinguish suppressive from lawful arrests. For 2017–2020, we extracted this information from `notes`. For 2021–2024, extracting it from either `notes` and `police_measures` gave identical results.

As in [1, 2], we wrangled the data into a new form where each row is a day, with the following columns, each of which we view as a time series:

- `num_protests:` protests per day.
- `num_protesters:` protesters per day.
- `kips:` protests involving KIPs.
- `chemical:` protests where chemical agents were used.
- `negativity:` protests featuring police negativity.
- `property:` protests with property damage.
- `injuries_crowd:` total injuries to protesters.
- `injuries_police:` total injuries to police.

This transformed data frame had 2,899 observations (days). On average, property was damaged once every two days. Police negativity was present in about 0.8 protests per day, chemical agents were used in about 0.17 protests per day, and kips were used in about 0.04 protests per day.

The variable `chemical_agents` is TRUE for 497 protests (<1%), with NA values treated as 0. Since this variable had 26% missing data, most of the time when it's FALSE, it means we have evidence that no chemical agents were used. Similarly, `property_damage_any` is TRUE for 1,429 protests (<1%), and because `property_damage` is missing for 79% of the events, a value of 0 for `property_damage_any` (and hence for property) does not indicate strong evidence that no property damage occurred. Most often, it indicates missing data. Hence, both chemical and property are biased downwards.

Because of heavy skew, we used log-transformations for protest size and protest count to obtain better-behaved residuals (see Section 5).

Finally, because the original CCC data set and data dictionary were recently removed from their public repository, we host a copy of the cleaned data (as of February 9, 2025) and our code at `https://github.com/xxxxx/ccc-analyses-2025`.

# 4 Statistical Methods

We analyzed the time series variables listed above using both changepoint detection and classical time series models, including Seasonal Autoregressive Integrated Moving Average (SARIMA) models and Seasonal Autoregressive Fractionally Integrated Moving Average (SARFIMA) models. All data wrangling and transformation were conducted in R, version 4.4.1. We begin by describing changepoint detection, then we describe univariate then bivariate time series models.

A changepoint in a time series is a moment in time when the series changes in some substantial way. Most time series models require that the series has the same behavior over the entire window of time in question. Before fitting these models, we used changepoint detection, which considers whether the mean of the time series (e.g., average number of protesters per day) is constant over time, or whether it jumps after a certain date. Essentially, this algorithm first computes the mean over the entire time series, as well as the sum of squared errors (SSE) against that mean, then considers all possible times where the time series could be split into two pieces and for each splitting, finds the means of both pieces. The algorithm then computes the SSE against the two means, and determines which splitting time (that is: changepoint) provides the lowest SSE, and whether that is statistically significantly better than the SSE from the single mean model, which has no changepoint. The algorithm can also attempt to find more changepoints, resulting in models with many levels. In our case, we did find a changepoint, so we fit our time series models only on the latter portion of the data, representing the 1728 days after May 25, 2020, when the George Floyd protests began.

We now describe the univariate time series models that we used. First, the SARIMA model helps us understand and predict time series data by accounting for three main features:

- **Autoregression (AR):** How the current value relates to its past values. For example, the number of protesters today might be related to the number of protesters yesterday or several days ago.

- **Integration (I):** Transforming the time series to make it stationary, which means its statistical properties like mean and variance do not change over time. This is why we cut the time series at the changepoint and retained only the 1728 most recent days.
- **Moving Average (MA):** How the current value is influenced by past errors or shocks – unexpected spikes or drops in the time series.

The "seasonal" part means the model can also consider repeating patterns over fixed intervals, like weekly cycles (for example, similar protest levels every 7 days). The differencing transformation means shifting from the original time series $x_t$ to the differenced time series

$$X_t = x_t - x_{t-1} = (1 - B)x_t$$

where $Bx_t = x_{t-1}$ is the backshift operator. Higher-order differencing transforms to $(1 - B)^d x_t$. First-order seasonal differencing transforms to $x_t - x_{t-7}$. SARIMA models are described using notation like SARIMA$(p, d, q) \times (P, D, Q)[S]$, where:

- $p$ = number of autoregressive terms,
- $d$ = degree of differencing (to make the data stationary),
- $q$ = number of moving average terms,
- $P, D, Q$ = the same terms but for the seasonal part,
- $S$ = the length of the seasonal cycle (e.g., 7 for a weekly pattern).

For example, our best univariate model for property damage was a SARIMA$(0, 1, 5) \times (1, 0, 1)[7]$ model. This means, if $x_t$ denotes the number of protests on day $t$, then we first transform to representing first-order differencing, then fit the model:

$$X_t = \Phi_1 X_{t-7} + \psi_1 e_{t-1} + \psi_2 e_{t-2} + \psi_3 e_{t-3} + \psi_4 e_{t-4} + \psi_5 e_{t-5} + \Psi_1 e_{t-7} + e_t.$$

This model posits that $X_t$ is affected linearly by $X_{t-7}$–the value of the time series on the same day last week–and how far that day differed from what was expected (this is $e_{t-7}$, the error term on day $t - 7$), plus how much each of the past five days differed from what was expected. The term $e_t$ is the error term on day $t$, which will be used in modeling $X_s$ for $s > t$.

We fit these models by first determining the seasonality, which was always $S = 7$ for us, representing the weekly pattern, since protests tend to occur in larger numbers on weekends than on weekdays. We next determined the degree and type (regular vs seasonal) of differencing required by plotting the time series. We tested that our differenced time series were stationary using plots and the Augmented Dickey-Fuller (ADF) test. To determine $p, q, P$, and $Q$, we used autocorrelation functions (ACFs) and partial autocorrelation functions (PACF) as in [25]. The ACF has $h$ on the x-axis and has the correlation between $x_t$ and $x_{t-h}$ on the y-axis. After differencing, the time series was stationary, so the ACF (measuring the dependence of $x_t$ on its past $x_{t-h}$) decays exponentially, and the ACF and PACF can show which $h$ have statistically significant correlations.

We compared the models determined in this way with those produced by the Hyndman-Khandakar algorithm, which seeks the model that minimizes the Akaike Information Criterion (AIC), a measure rooted in information theory that helps evaluate and compare statistical models. AIC balances model fit and complexity by penalizing models with more parameters, favoring those that explain the data well with fewer assumptions. Lower AIC values indicate a model that is more likely to predict future data accurately without overfitting.

After identifying the specific SARIMA model to fit, we determined the coefficients using maximum likelihood estimation, and then checked that the residuals are random and independent (i.e., no further unmodelled temporal autocorrelation) using the ACF and the Ljung-Box test. In Section 5 we report summary tables showing which variables were statistically significant and their coefficients. For more summary information, such as the AIC, we direct readers to the supplementary file at `https://github.com/xxxxx/ccc-analyses-2025`. In many cases, the residuals are not normally distributed, but the Central Limit Theorem nevertheless guarantees that the coefficients of the model are approximately normally distributed, thanks to the large sample sizes [25].

For our time series regression models, we first use cross-correlation analysis to determine whether our explanatory variable(s) $x_t$, e.g., use of chemical agents, leads or lags behind our response variable $y_t$, e.g., number of protests. If $x_t$ lags behind $y_t$ then we reverse the order of our explanatory and response variables, because we do not want models that try to predict the past using the future. For example, police use of chemical agents tends to lag behind property damage, meaning property damage usually happens first. In some cases, we use lagplots to visually represent the correlations between $x_{t-h}$ and $y_t$ for various values of the lag $h$.

After cross-correlation analysis has identified which lags of the explanatory variable–e.g., $x_t, x_{t-1}$, and $x_{t-2}$–are relevant for predicting $y_t$, we fit a regression model for $y$ based on those lags of $x$, then use the SARIMA procedure discussed above on the residuals of that model. The result is a model for $y$ in terms of $x$, the history of $x$, and the history of $y$, whose error terms are random and independent. Such a model allows us to determine which coefficients are statistically significant. For example, one model we fit for the effect of chemical agents on crowd injuries is:

$$y_t = \beta_0 x_t + \beta_1 x_{t-1} + \beta_2 x_{t-2} + \beta_3 x_{t-3} + \beta_4 x_{t-4} + \psi_1 e_{t-1} + e_t.$$

The first part is the time series regression (or `xreg`) part, and the last two terms are the SARIMA part. The residuals $e_t$ were random and independent, so we were able to determine that $a_0 = 2.1$ and $a_4 = 1.2$ were statistically significant and positive. This means one additional use of chemical agents today (increasing $x_t$ by 1) is associated with 2.1 more injuries today, and 1.2 more injuries four days from now, due to protests becoming more violent. If multiple time series regression models fit (and leave good residuals), we bias in favor of simpler models, dropping non-significant variables.

We use SARFIMA models when the ACF decays too slowly for a stationary time series (where it would decay exponentially) but faster than a time series requiring full differencing

[25]. This involves fractional differencing, a transformation defined via Taylor series, which allows the influence of past events on the present to decay more slowly over time compared with regular differencing. The first step is to determine the degree of differencing (e.g., $d = 0.4$) and whether it is statistically significant. Then, after applying the fractional differencing transformation

$$X_t = (1 - B)^d x_t,$$

one fits a SARIMA model as usual. The fact that $x_t$ and $x_{t-h}$ remain correlated–even for large $h$–explains the "long memory" name. Formally, the term "long memory" means that the autocorrelation function $f(h) = \text{corr}(x_t, x_{t-h})$ is not absolutely summable. Although the term is usually associated with SARFIMA models, note that it also applies if full differencing is required. For example, Figure 3 shows that the time series `log(num_protesters)` has a long memory regardless of whether we model it using SARFIMA or SARIMA. In some cases, both a SARIMA model and a SARFIMA model were able to obtain random and independent residuals. In those cases, we use cross-validation to determine which model performed best. Even if the SARFIMA performed best, we still also report the SARIMA results as they are easier to understand, in line with the principle of parsimony.

We assessed each model by first checking that the residuals were random and independent (i.e., the ACF showed no statistically significant correlations between $e_t$ and $e_{t-h}$), then checking whether all coefficients in the model were statistically significant and considering simpler models by dropping non-significant variables (e.g., shifting from 4 lags to 3), and when we had multiple models that fit well we compared them using AIC and adjusted $R^2$. We report the best models in Section 5, and our R code (available at `https://github.com/xxxxx/ccc-analyses-2025`) shows all models we fit.

In summary, our modeling procedure involved:

1. **Checking Stationarity:** We first tested whether the time series were stationary or needed differencing to stabilize their mean and variance.

2. **Selecting Model Parameters:** We used statistical tools such as plots of autocorrelation and partial autocorrelation to identify appropriate lag terms for AR and MA components.

3. **Fitting Models:** We fit candidate SARIMA and SARFIMA models to the data.

4. **Diagnostic Checks:** We examined the residuals (the differences between observed and predicted values) to ensure they behaved like random noise, indicating a good fit.

5. **Model Comparison:** To decide which model best described each variable, we used the Akaike Information Criterion (AIC), which balances model fit with complexity. The model with the lowest AIC is considered the best compromise between accuracy and simplicity. We report both the AIC and the adjusted $R^2$.

By fitting models that produced random and independent residuals and achieving optimal Akaike Information Criterion scores, we successfully captured the underlying dynamics of protest events. These models enabled us to identify statistically significant relationships between explanatory variables and outcomes, allowing us to rigorously quantify the effects of specific police actions—such as the use of tear gas or rubber bullets—on protest activity and associated violence.

# 5 Results

In this section we state our main results. We begin with the results of our changepoint detection, then the results of our univariate and bivariate time series models, to show the effects of police actions on protest dynamics.

## 5.1 Changepoint detection

For each of the eight time series–`log(num_protests)`, `log(num_protesters)`, `chemical`, `kips`, `negativity`, `property`, `injuries_crowd`, `injuries_police`–we first assessed stationarity. As part of this, we considered whether there was evidence of a changepoint.

We found that a model with two means—one for dates before May 25, 2020 and one for dates after—fit substantially better than having one mean for all dates. Figure 1 shows the jump in protest activity (for the log of the number of protesters), aligning with the death of George Floyd. That is, the time series was categorically different after this date, meaning a single univariate time series model would not fit, because the assumption of stationarity is violated. For example, before May 25, 2020, there were several days with zero protesters (where the time series drops to the bottom of the frame in Figure 1), but after May 25, 2020 there were none.

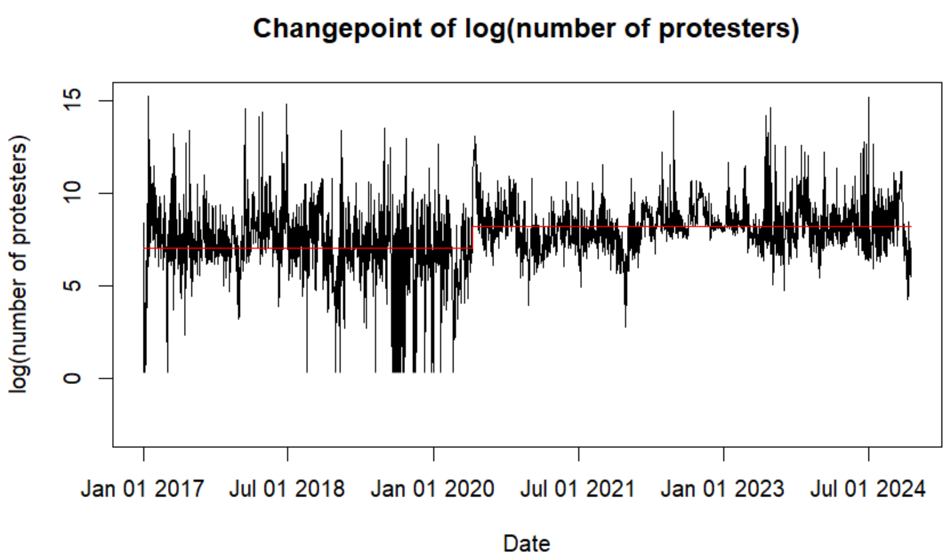


Figure 1: Time plot of the log-transformed number of protesters, showing the changepoint at May 25, 2020.

This finding, that the intensity of protests jumped in spring of 2020, aligns with the analysis of [1]. In addition to checking if the mean changed, we also checked whether the variance was constant over time. We did not detect any statistically significant changepoint for the variance.

There were no significant changepoints, in any of the time series, after May of 2020. That is, once the protests escalated in 2020, they never returned to their previous levels.

Consequently, we sliced all the time series to begin from May 26, 2020, and proceeded to fit time series models, explained in the next sections. Alternatively, we could have added a binary variable telling whether each date was before or after the changepoint, allowing different models for each half of the time series, but because our focus is on more recent protests, we preferred to slice the data and discard the earlier portion when protests were categorically different.

## 5.2 Univariate time series models

For each of the eight time series, we fit univariate time series models, to learn how the variable depended on its own history. Two of the variables—num_protests and num_protesters—required log transforms because they were spread across several orders of magnitude, with large outliers, as Figure 2 demonstrates. Differencing would not reduce the presence of these outliers, since $x_t - x_{t-1}$ would still be an outlier if $x_t$ was.

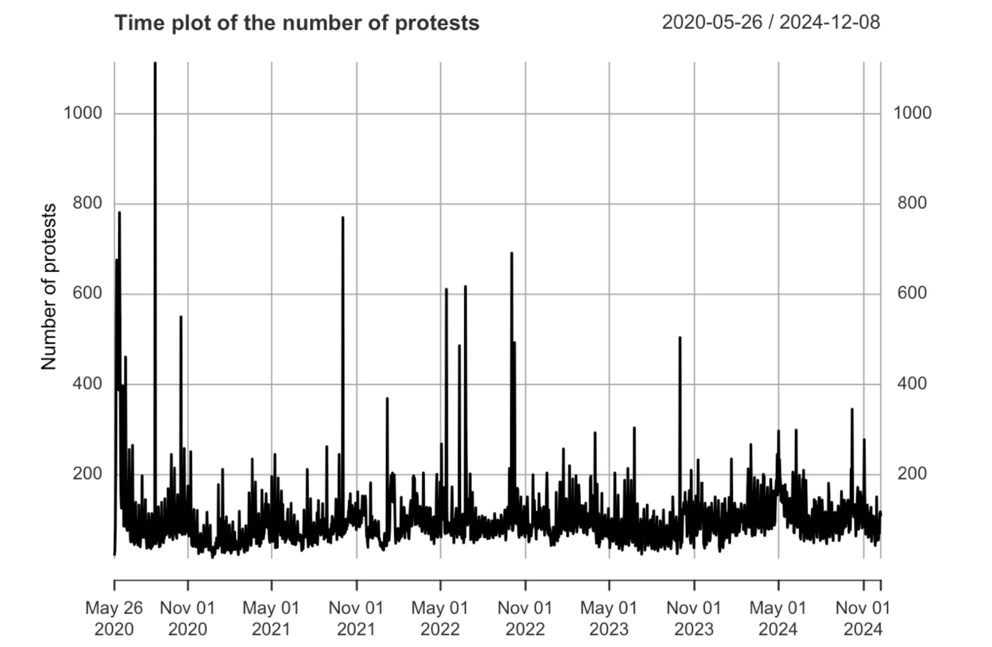


Figure 2: A time plot of the number of protests, before a log transform was applied.

As part of the SARIMA model fitting procedure described in Section 4, we inspected the ACF plots of all eight time series. The ACF plot for `log(num_protesters)` suggested that a SARFIMA model might fit well (see Figure 3). As discussed below, we were able to model `log(num_protesters)` using both a SARFIMA model and a simpler SARIMA model (with seasonal differencing), achieving random and independent residuals in both cases.

### 5.2.1 Summary of univariate models

As our main focus is on the time series regression models—like the influence of chemical agents on protest violence—described in Section 5.3, we only briefly state the results from

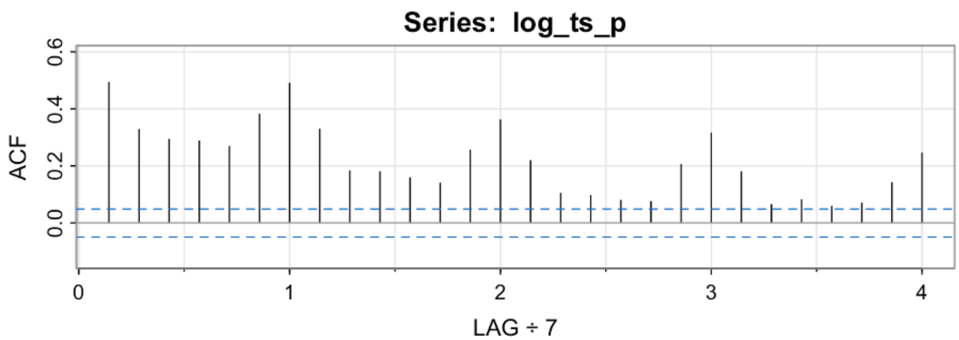


Figure 3: ACF plot for the log of the number of protesters, showing a long-memory phenomenon.

the univariate models in Table 1. Each of these models was the simplest model that achieved random and independent residuals. When we had multiple models that achieved random and independent residuals, we compared them using AIC or cross-validation when AIC was not possible (e.g., when different transformations had been applied). For example, we had both SARIMA and SARFIMA models for both `log(num_protests)` and `log(num_protesters)`. Cross-validation suggested that the SARIMA model was best for the former and SARFIMA for the latter (but, it only beat the SARIMA model by 1%). In Table 1, we use the notation of Section 4, and report relevant statistically significant coefficients, using `ar1` (resp. `sar1`) to denote the coefficient on $x_{t-1}$ (resp. $x_{t-7}$).

Table 1: Summary of univariate models.

| Variable | Best-Fitting Model(s) | Notes |
|---|---|---|
| Kips | ARIMA(3,1,4) | ar3 = 0.59 |
| Chemical | ARIMA(1,1,4) | ar1 = 0.35 |
| Negativity | ARIMA(3,1,3) | ar1 = 0.92 |
| Property | SARIMA(0,1,5)(1,0,1)[7] | sar1 = 0.99 |
| Injuries (crowd) | ARIMA(0,1,1) | ma1 = −0.97 |
| Injuries (police) | ARIMA(0,1,4) | ma1 = −0.61 |
| log(num_protests) | SARIMA(1,0,1)(1,1,1)[7] or SARFIMA(2,0,2)(2,0,1)[7] with $d = 0.173, d_7 = 0.321$ | ar1 = 0.90 |
| log(num_protesters) | SARIMA(1,0,2)(1,1,2)[7] or SARFIMA(2,0,2)(2,0,1)[7] with $d = 0.334$, $d_7 = 0.322$ | ar1 = 0.92, sar1 = 0.82 |

Each model that had an autoregressive term had a positive and statistically significant coefficient, suggesting that they may be self-exciting as in [1, 2], i.e., a high value today tends to predict high values tomorrow. For example, police negativity today is strongly correlated with police negativity tomorrow. Often, the moving average coefficients were negative, meaning an unusually high value today is counterbalanced by lower values in the subsequent days, perhaps in response to public outrage. A supplementary file showing all the models, their coefficients, and summary statistics, is hosted at `https://github.com/xxxxx/ccc-analyses-2025`.

The long-memory phenomenon in `log(num_protesters)` from Figure 3 means, in practice, that means large protest events, like the 2017 Women's March—where over 500,000 people marched in Washington, D.C., as well as hundreds of thousands in other cities—are remembered for a long time and can influence protests in the distant future.

### 5.2.2 Replicating related research

Our analysis allows us to attempt to replicate findings from [1, 2]. Specifically, in [1], based on the ACLED data, the authors found that for `num_protests`, a SARIMA(3,1,2)×(2,0,0)[7] fit well. We fit this model to our `num_protests` and found that it did not fit well, i.e., it left a pattern in the residuals, including the outliers from Figure **??** and a pattern in the ACF of the residuals. This makes sense because the CCC data contains more protests than the ACLED data, and our analysis covers a wider period; the analysis in [1] runs from January 1, 2020 until March 12, 2021. Although the CCC variable `num_protests` requires a log transform, we did replicate the finding of [1] that this time series has a positive and statistically significant `ar1` term, suggesting that the dynamical system model of [1] might again be appropriate.

In [1], the authors found that an ARIMA(1,0,1) fit their KIP time series, with an ar1 = 0.6. The CCC data is richer than the ACLED data regarding usage of KIPs, so it makes sense that a more complicated model is needed for the CCC data (and, indeed, the model of [1] does not fit well), but it is interesting that the autoregressive coefficients are similar. Our model shows that the use of KIPs today can affect KIP use for the next three days. This is related to a feedback loop phenomenon that we will investigate more below with our time series regression models. We also remark that [1] includes a model for a time series of deaths, but the CCC data set is missing too much data for us to attempt a similar analysis.

Similarly, in [2], based on data from Ukraine, the authors found that for `num_protests`, a SARIMA(2,1,3)×(1,0,2)[7] model fit well, but this model did not fit our `num_protests` well. We are not surprised that a model focused on Ukraine in 2013 would not fit a data set of protests in 2020–2024 in the United States, especially since Ukraine has a smaller population and far fewer outlier days with an unusually large number of protests. Interestingly, the strength of the dependence on the previous day was similar: ar1 = 1.45 in [2] was similar to the ar1 = 1.49 we found when we fit the model of [2] to our `num_protests`.

The fact that our `log(num_protests)` requires seasonal differencing, whereas the time series of [1, 2] did not, suggests that protests in recent years in the United States have an autocorrelation structure with a longer time horizon. In other words, the average protest today is remembered longer and can affect protests in the distant future, more than the average protest in 2020 or in Ukraine in 2013.

## 5.3 Time series regression models

In this section, we consider the effect of three explanatory variables—use of KIPs, use of chemical agents, and police negativity—on four response variables: `log(num_protests)`, log(num_protesters), injuries_police, and injuries_crowd. We first determine the lead/lag relationship between the variables using cross-correlation analysis. Then we fit the time series regression models as described in Section 4. In Section 5.3.3 we replicate the results from [1, 2].

### 5.3.1 Cross-correlation analysis

For each pair (response variable, explanatory variable), we carried out cross-correlation analysis to determine the lead/lag relationship. The first step is to pre-whiten the response variable $y_t$, meaning to fit a univariate model that leaves random and independent residuals (as in Section 5.2), then save the residuals $r_t$, which are now white noise and do not depend on their own history. The machinery of SARIMA provides a transformation that turns $y_t$ into $r_t$. We next apply that same transformation to $x_t$ to produce a filtered version $f_t$ of $x_t$ that can be compared to $r_t$. Finally, we compute the correlation between $f_t$ and $r_{t+h}$ for every lag $h$. Figures 4 and 5 provide examples of the resulting cross-correlation function (CCF).

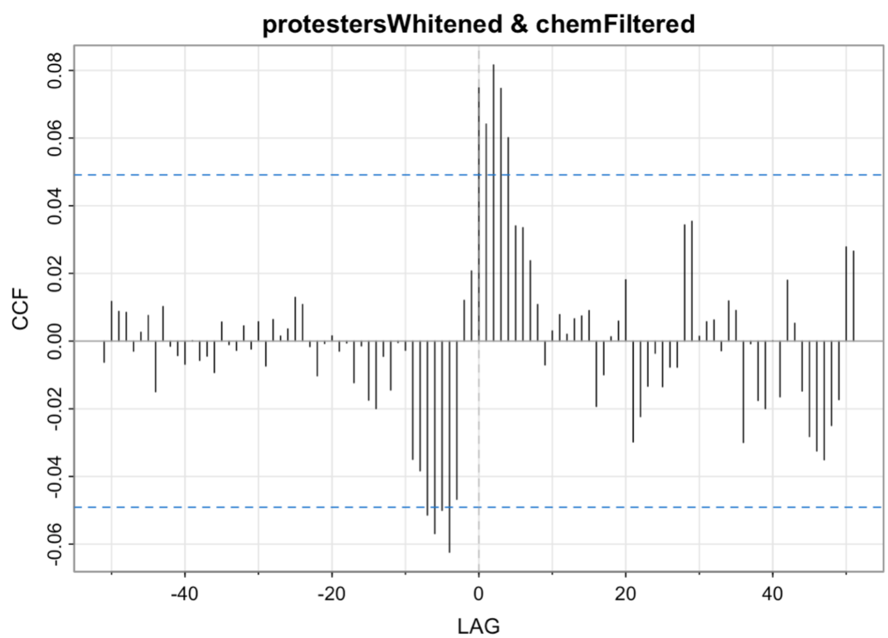


Figure 4: Use of chemical agents leads the number of protesters.

Figure 4 shows positive correlations at lags $h = 0, 1, 2, 3$, and 4. So, an increase in $f_t$—the uses of chemical agents today—is associated with an increase in $r_{t+1}$ and hence with the number of protesters tomorrow. We learn that chemical agent use is associated with larger protests in the subsequent four days, and we say that chemical agent use leads the number of protesters. The left side of Figure 4 describes the lag relationship. Taking $h = -4$ shows a negative correlation between $f_t$ and $r_{t-4}$. So, a large number of protesters on day $t - 4$ is associated with less chemical agent usage on day $t$. The blue dotted lines show that most other values of $h$ do not yield a statistically significant cross-correlation.

Sometimes, the lead/lag relationship does not tell a clear story. For example, Figure 5 has positive cross-correlations on both the left and the right side of lag $h = 0$.

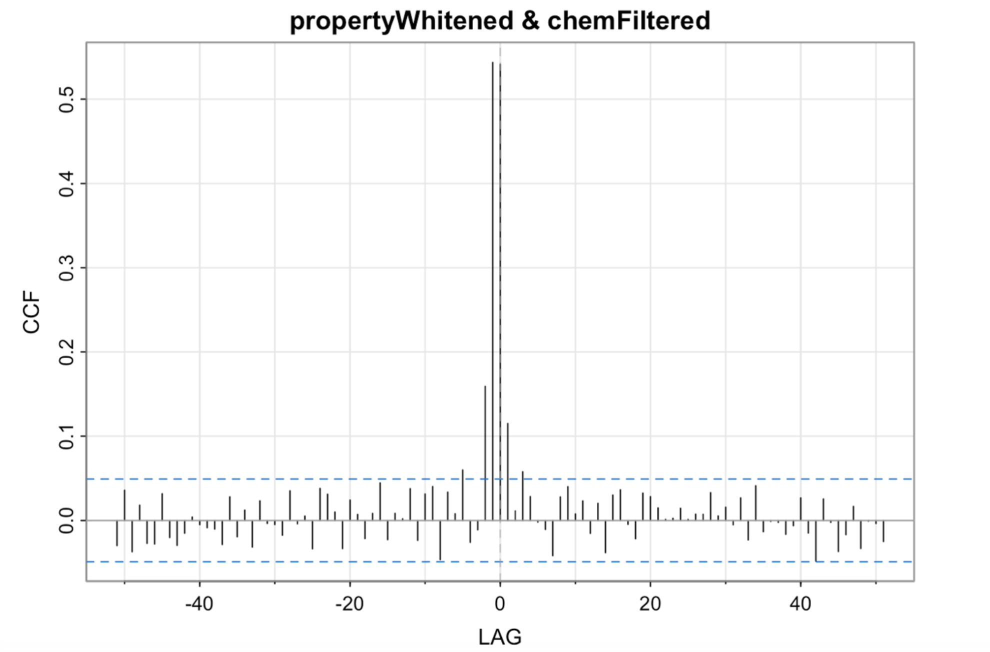


Figure 5: The use of chemical agents both leads and lags property damage.

In Figure 5, the large correlation at lag $h = 0$ suggests that days with a lot of property damage also have a large use of chemical agents. The large correlation at lag $-1$ suggests that when property damage occurs, more chemical agents tend to be used the next day. However, there is also a statistically significant cross-correlation at lag $h = 1$ (albeit smaller), suggesting that chemical agent use today is associated with more property damage tomorrow. We refer to this as a feedback effect, where each variable drives the other higher.

We summarize the lead/lag relationships in Table 2. Because police negativity was essentially the sum of KIP use and chemical agent use, we omit the cross-correlation analysis for this explanatory variable and the four response variables.

The last four CCFs of Table 2 have positive cross-correlations in both directions. For the `injuries_police` rows, we refer to this as a cycle of violence, because police use of force is associated with subsequent injuries to police, which is itself associated with subsequent police use of force. The `injuries_crowd` rows show a similar pattern, but with far more associations on the right side of the graph, meaning police use of force is associated with subsequent injuries to the crowd. However, the cross-correlations on the left side of the graph suggest that sometimes crowd injuries occur first, and police use of force occurs later, perhaps to attempt to quell violence in ongoing protests. Regardless, this police use of force fails to quell violence, as it is associated with violence at subsequent protests.

We note that cross-correlation analysis does not imply causality. For example, a significant cross-correlation at lag 0 could mean that $x_t$ causes $y_t$ or that $y_t$ causes $x_t$ or neither (e.g., both could be caused by some exogenous event). Our full analyses and collection of all CCF plots we made are hosted at `https://github.com/xxxxx/ccc-analyses-2025`.

Table 2: Summary of cross-correlation analysis.

| Response variable | Explanatory | Significant lags | Interpretation |
|---|---|---|---|
| log(num_protesters) | KIP use | 0 | Same-day association |
| log(num_protesters) | Chemical | 0, 1, 2, 3, 4; −4 | Chemical leads (up to 4 days); less after large protests |
| log(num_protests) | KIP use | 0 | Same-day association |
| log(num_protests) | Chemical | 0, 1; −3, −4, −5; seasonal ±7 | Chemical leads; fewer chemical uses after many protests; feedback effect with seasonal cycles |
| Property damage | KIP use | 0 | Same-day association |
| Property damage | Chemical | $-2, -1, 0, 1$ | Feedback loop: chemical ↔ damage |
| Injuries (police) | KIP use | $-3, 6$ | Cycle of violence |
| Injuries (police) | Chemical | $-2, -1, 0, 1$ | Cycle of violence |
| Injuries (crowd) | KIP use | $-2, 0, 2, 4, 5, 6$ | Feedback effect plus KIPs → more crowd injuries over following days |
| Injuries (crowd) | Chemical | −3 to 8 | Feedback effect plus chemical → more crowd injuries up to 8 days later |

### 5.3.2 Results of time series regression models

For each response variable $y_t$, we fit three time series regression models as described in Section 4 one based on KIP use $k_t$, one based on chemical agent use $c_t$, and one based on both together, which is equivalent to police negativity $n_t$. For each explanatory variable $x_t$, this involves first regressing $y_t$ on $x_t$ and any relevant lags $x_{t-h}$, then modeling the residuals using SARIMA or SARFIMA as in Section 5.2.

We only present models that achieved random and independent error terms of this two-step process. In situations where multiple models did so, we used AIC to decide between them, and most often present the simpler model. Since we were interested in the effect of police use of force on subsequent protests, we did not include terms like $x_{t+h}$ in our model, which would be predicting protests today based on future police use of force.

Note that the correlation between $k_t$ and $c_t$ was only 0.04 at lag 0, and essentially zero at all other lags, so including both of these variables in a model does not introduce multicollinearity. We summarize our models in Table 3, stating the statistically significant coefficient values and how to interpret each model (the symbol "→" should be read "is associated with" not "causes"). In many cases, the residuals had a positive and statistically significant ar1 term, and the explanatory variable had a statistically significant and positive coefficient. In such cases, increases in police use of force are associated with larger values of $y_t$ and hence also of $y_{t+1}$, via the ar1 term.

In a few cases, both a SARIMA and a SARFIMA model fit the residuals, but in all cases the SARIMA model did a better job, and was also simpler, so we only present the SARIMA models below. The SARIMA models still required differencing, and sometimes seasonal differencing. The ACFs make it clear that the effect of police use of force on protest size, destructiveness, and violence, has a long memory.

| Response var | Predictor | Coefficient(s) | Interpretation |
|---|---|---|---|
| log(num_protests) | $k_t$ | $k_t : 0.0598$ | KIP use → more protests today and in subsequent days (positive ar1). |
| | $c_t$ | $c_t$: 0.0129<br>$c_{t-1}$: 0.0151 | Chemical agent use → more protests today and tomorrow, and subsequent days (positive ar terms). |
| | both | $k_t : .057$, $c_t : .012$,<br>$c_{t-1} : 0.0159$ | Both → more protests today and tomorrow, and subsequent days. |
| log(protesters) | $k_t$ | $k_t : 0.2552$ | KIPs → more protesters today and in subsequent days. |
| | $c_t$ | $c_t : 0.0440$ | Chemical agents → more protesters today and in subsequent days. |
| | both | $k_t : .218$, $c_t : .037$,<br>$c_{t-4} : 0.0381$ | Both → more protesters. Chemical effects persist at least 4 days. |
| property damage | $k_t$ | $k_t : 0.2927$ | KIPs → more property damage today and in subsequent days. |
| | $c_t$ | $c_t : 1.4199$,<br>$c_{t-1} : -0.308$ | Chemical agents → more damage today but a bit less tomorrow. |
| | both | $k_{t-1} : .376$,<br>$c_t : 1.4724$,<br>$c_{t-3} : 0.1$ | Chemical agents → more property damage today and in 3 days; KIPs → more property damage tomorrow. |
| injuries_police | $k_t$ | $k_{t-5} : 1.4709$ | KIPs → more police injuries several days later. |
| | $c_t$ | $c_t : 2.1031$ | Chemical agents → more police injuries. |
| | both | $k_{t-5} : 1.3805$,<br>$c_t : 2.0838$ | Same findings as individual models. |
| injuries_crowd | $k_t$ | $k_t : 1.2946$, $k_{t-4} : 1.5785$,<br>$k_{t-5} : 1.069$, $k_{t-6} : 0.99$ | KIPs → crowd injuries for the next week, as protests escalate. |
| | $c_t$ | $c_t : 2.4641$, $c_{t-1} : 0.7347$,<br>$c_{t-3} : -0.756$,<br>$c_{t-4} : 1.2415$ | Chemical agents → crowd injuries today and in later days; short-term suppression before escalation. |
| | both | $k_t : 0.5255$, $k_{t-4} : 0.8892$,<br>$k_{t-5} : 0.989$, $k_{t-6} : 0.851$,<br>$c_t : 1.9709$, $c_{t-2} : -0.667$,<br>$c_{t-4} : 1.1838$ | Same findings as individual models; KIPs → crowd injuries for next week, chemical → more injuries today, less in two days, and more in four days. |

Table 3: Results of time series regression models.

Full specifications of each model may be found in the supplemental file hosted at `https://github.com/xxxxx/ccc-analyses-2025`. The coefficients in Table 3 give estimates of the magnitude of the effect of each variable. For example, the first model for police injuries shows that one additional protest at which police use KIPs is associated with 1.47 more police injuries five days later. For multivariate regression models, one usually interprets each coefficient as the expected change in the response variable for a unit change in the given explanatory variable, after allowing for simultaneous change in the other explanatory variables. Since $k_t$ and $c_t$ have such a small correlation, knowing that $c_t$ increased does not tell us much information about whether $k_t$ changed, so there is little harm in interpreting coefficients as for regression models, e.g., the last model shows that one additional event with chemical agents is associated with an average of 1.1838 additional injuries to protesters four days later. Lastly, when the response variable has a log transform, we must apply the inverse function and interpret the coefficient multiplicatively, e.g., in the first model the KIP coefficient of 0.0598 transforms to $e^{0.0598} \approx 1.06$ and we say that one additional KIP event is associated with a 6% increase in the number of protests.

### 5.3.3 Replicating related work

We close this section by discussing the models from [1] and [2]. In [1], the best model for the number of protests as a function of KIP usage, based on the ACLED data, was a regression on $k_t$ with SARIMA(2,1,2)×(2,0,0)[7] residuals. Our model from Table 3 replicates the finding that $k_{t-1}$ is not statistically significant but that the ar1 term is statistically significant and large, representing a pathway for an increase in $k_t$ to affect the number of protests in subsequent days.

We also fit the model from [1] to our (untransformed) num_protests. The model did not achieve independent residuals, but this is common in time series analysis since the CCC data time series is very different from the ACLED time series. Nevertheless, both models had a large coefficient for $k_t$—it was 10 for [1] and 7 for us—and a statistically significant positive ar1 term.

In [2], the best model for the number of protests as a function of negative response events (their version of our `police_negativity`), based on the CSLR data in Ukraine, was a regression on $n_t$, $n_{t-1}$, and $n_{t-2}$, with SARIMA(1,0,2)×(1,0,2)[7] residuals. We found that this model fit well, leaving independent residuals. Furthermore, we replicate the finding from [2] that all three explanatory variables have positive and statistically significant coefficients—2.3247, 0.7656, and 0.4810 in Ukraine and 5.7393, 1.1120, and 1.9256 in the USA—and nearly identical ar1 values of approximately 0.9. So, police negativity in the USA is associated with more protests in the subsequent two days, with even larger coefficients than in Ukraine, though this aligns with the higher frequency of protests in the CCC data.

In [2], the best model for the number of protests as a function of crowd injuries was a regression on $i_t$ and $i_{t-1}$ with SARIMA(2,1,2)×(2,0,0)[7] residuals. We fit this model to the CCC data and found that it did leave autocorrelation in the residuals. However, we replicate the finding that both explanatory variables had statistically significant and positive coefficients.

In [2], the best model for the number of protests as a function of crowd injuries and negative response events was a regression on $i_t$ and $n_t$, with ARIMA(1,1,1)×(0,0,2)[7] errors. A competing model had ARIMA(2,1,1) errors. We fit both models and found that both left autocorrelation in the residuals. We replicated the finding that negative response rates are associated with large increases in protests (each additional negative response event was associated with 8 additional protests) and that the ar1 term is large and positive. However, unlike [2], we find that in the presence of $n_t$, the variable $i_t$ no longer has a strong effect, probably because our model in Table 3 shows that $n_t$ has strong predictive power for $i_t$. Likely the difference between our results and [2] is due to the CCC data set having far more complete data regarding both injuries and police negativity, compared with the Ukraine data.

Additional models are presented in [1] and [2] that we could not attempt to replicate. For example, in [1] the authors also modelled deaths as a function of KIP usage, but we did not because of missing data in the CCC data set. Similarly, in [2] the authors modelled number of protests as a function of number of arrests, but we did not because CCC lacks data on arrests. Lastly, [2] contains many models where one explanatory variable is the number of protests associated with the Euromaidan movement, but there is no analogous movement driving protests in the CCC data.

# 6 Conclusion

Our analysis of the CCC data shows that protest dynamics in the United States since May 2020 can be modeled effectively using modern time series techniques, including SARIMA and SARFIMA models. We confirm earlier findings that protest-related variables are often self-exciting: large protests and police use of force today are statistically associated with increases in both tomorrow. We extend prior work by investigating the role of police use of chemical agents (like tear gas and pepper spray), and determining the effect of police use of force on the number of protesters, injuries at protests (both to protesters and police), and property damage. Additionally, we document long-memory effects, particularly for the number of protesters, suggesting that unusually large events may influence protest dynamics well into the future.

Our cross-correlation analysis finds the lead/lag relationships between police use of force and protest frequency, size, and violence, while time-series regression models show positive correlations between police use KIPs and chemical agents, and subsequent increases in the number of protests, protest size, injuries, and property damage. We discover feedback loops and cycles of violence in which police use of force and injuries to police are linked across multiple days. These findings reinforce concerns raised in prior literature about escalated force strategies, which do not appear to reduce unrest and may instead be linked to its persistence.

More broadly, our results highlight the importance of treating protest-police interactions as dynamic systems with memory, rather than as isolated events. Statistical models that capture feedback and persistence provide useful tools for understanding these systems, and our replication of earlier findings in new data strengthens the evidence base for considering alternatives to escalated force. Future work could expand on this approach by incorporating missing variables (such as the presence of counter-protesters or protester demographics) or by extending analyses to international contexts.

# References


[1] Rodriguez, N., & White, D. (2023). An analysis of protesting activity and trauma through mathematical and statistical models. *Crime Science*, 12:17, 1–20.

[2] Bahid, Y., Kutsenko, O., Rodriguez, N., & White, D. (2024). The statistical and dynamic modeling of the first part of the 2013–2014 Euromaidan protests in Ukraine: The Revolution of Dignity and preceding times. *PLoS ONE*, 19(5): e0301639.

[3] Ma, L., Tran, L., & White, D. (2021). A statistical analysis of drug seizures and opioid overdose deaths in Ohio from 2014 to 2018. *Journal of Student Research*, 10(1).

[4] Berestycki, H., Nadal, J.-P., & Rodríguez, N. (2015). A model of riot dynamics: Shocks, diffusion, and thresholds. *Networks and Heterogeneous Media*, 10(3), 1–34.

[5] Nicoara, O., & White, D. (2015). A contextual model of the secessionist rebellion in eastern Ukraine. *2015 Annual Meetings of the Public Choice Society*, available as arXiv:1606.02748.

[6] della Porta, D., & Reiter, H. (1998). *Policing Protest: The Control of Mass Demonstrations in Western Democracies*. Minneapolis: University of Minnesota Press.

[7] White, D., Rodriguez, N., & Topaz, C. (2024). Cooperation can end violent protests. *New York Daily News*, May 16, 2024.

[8] White, D., & Topaz, C. (2024). Foster negotiation between police and protesters for a peaceful DNC. *Chicago Tribune*, August 16, 2024.

[9] White, D., & Topaz, C. (2024). Will Milwaukee police fuel violence at Republican National Convention protests? *Milwaukee Journal Sentinel*, July 11, 2024.

[10] White, D., Rodriguez, N., Topaz, C., & Khiguradze, T. (2024). Police must refrain from escalating violence. *JAM News*, May 27, 2024.

[11] Maguire, E. R. (2015). New directions in protest policing. *Saint Louis University Public Law Review*, 35(1).

[12] U.S. Department of Justice, Office of Community Oriented Policing Services, & Institute for Intergovernmental Research. (2015). *After-action assessment of the police response to the August 2014 demonstrations in Ferguson, Missouri*.

[13] Earl, J., & Soule, S. A. (2010). The impacts of repression: The effect of police presence and action on subsequent protest rates. In *Research in Social Movements, Conflicts and Change* (Vol. 30, pp. 75–113). Emerald Group Publishing Limited.

[14] Delehanty, C., Mewhirter, J., Welch, R., & Wilks, J. (2017). Militarization and police violence: The case of the 1033 program. *Research & Politics*, 4(2), 2053168017712885.

[15] Ni, M. Y., Kim, Y., McDowell, I., Wong, S., Qiu, H., Wong, I., Galea, S., & Leung, G. (2020). Mental health during and after protests, riots and revolutions: A systematic review. *Australian and New Zealand Journal of Psychiatry*, 54(3), 232–243.

[16] First, J. M., Danforth, L., Frisby, C. M., Warner, B. R., Ferguson Jr., M. W., & Houston, J. B. (2020). Post-traumatic stress related to the killing of Michael Brown and resulting civil unrest in Ferguson, Missouri: Roles of protest engagement, media use, race, and resilience. *Journal of the Society for Social Work and Research*, 11(3), 369–391.

[17] Berestycki, H., & Nadal, J.-P. (2010). Self-organised critical hot spots of criminal activity. *European Journal of Applied Mathematics*, 21(4–5), 371–399.

[18] Yang, C., & Rodrıguez, N. (2020). A numerical perspective on traveling wave solutions in a system for rioting activity. *Applied Mathematics and Computation*, 364.

[19] Davies, T. P., Fry, H. M., Wilson, A. G., & Bishop, S. R. (2013). A mathematical model of the London riots and their policing. *Scientific Reports*, 3, 1303.

[20] Lang, J., & De Sterck, H. (2014). The Arab Spring: A simple compartmental model for the dynamics of a revolution. *Mathematical Social Sciences*, 69, 12–21.

[21] Berenji, B., Chou, T., & D'Orsogna, M. (2014). Recidivism and rehabilitation of criminal offenders: A carrot and stick evolutionary game. *PLoS ONE*, 9(1), e85531.

[22] Nonviolent Action Lab. Crowd Counting Consortium. https://github.com/nonviolent-action-lab/crowd-counting-consortium

[23] Raleigh, C., Linke, A., Hegre, H., & Karlsen, J. (2010). Introducing ACLED: An armed conflict location and event dataset: Special data feature. *Journal of Peace Research*, 45(7), 651–660.

[24] The Center for Social and Labor Research. Ukrainian Protest and Coercion Data Project; 2023. https://www.cslr.org.ua/en/ukrainian-protest-and-coercion-data-project/

[25] Shumway, R. H., & Stoffer, D. S. *Time Series Analysis and Its Applications* (4th ed.). Springer.

[26] Smelser, N. (1994). *Sociologiya (Sociology)*. Moscow: Phoenix.

[27] Artemov, G., Aleinikov, A., Daur, A., Pinkevich, A., & Abalian, A. (2017). Social tension: the possibility of conflict diagnosis (on the example of St. Petersburg). *Economics & Sociology*, 10, 192–208.

[28] Cosner, L. (1956). *The Function of Social Conflict*. New York: The Free Press.

[29] Drury, J., & Reicher, S. (2000). Collective action and psychological change: The emergence of new social identities. *British Journal of Social Psychology*, 39(4), 579–604.

[30] Chenoweth, E., Hammam, S., Pressman, J., & Ulfelder, J. (2024). Protests in the United States on Palestine and Israel, 2023–2024. *Social Movement Studies*, 1–14. https://doi.org/10.1080/14742837.2024.2415674

[31] Ulfelder, J. (2025). Crowd Counting Consortium U.S. Protest Event Data. *Harvard Dataverse*, V1, UNF:6:B9cfwGHNn5ss5fu7EHFhbQ. https://doi.org/10.7910/DVN/6OPP7H

[32] Bermingham, N., Willey, N., & White, D. (2025). Tracking the spatiotemporal spread of the Ohio overdose epidemic with topological data analysis, arXiv:2509.22705.