Political Science Research and Methods (2026), 1–18
doi:10.1017/psrm.2026.10119

ORIGINAL ARTICLE

Echoes of conflict: quantifying redundancy in fatality time series

Thomas Schincariol and Thomas Chadefaux

Political Science, Trinity College, Dublin, Ireland
Corresponding author: Thomas Schincariol; Email: schincat@tcd.ie

(Received 2 March 2025 UTC; revised 9 December 2025 UTC; accepted 12 March 2026 UTC)

Abstract
Do wars follow predictable patterns, or are they chaotic and inherently unpredictable? We explore this question by quantifying temporal redundancy—how often similar patterns recur—in conflict-related fatalities. Using entropy-based analysis—a measure of randomness—and Dynamic Time Warping, we compare conflict dynamics with other complex systems such as seismology and finance. Our findings reveal that the dynamics of fatalities within conflicts exhibit high levels of entropy: they are largely unpredictable. Conflicts generally display higher entropy than climate or earthquake data, which highlights the challenges of predictive modeling. Moreover, conflict patterns are not region- or decade-specific, which suggests a degree of universality. Our framework distinguishes stable from chaotic phases in conflict dynamics and clarifies where forecasting is feasible.

Keywords: conflict data; entropy; fatalities; temporal patterns; temporal redundancy

The adage "History repeats itself" has long intrigued scholars and laypeople alike. It suggests that historical events and patterns tend to recur over time and that human behavior, societal trends, and global events often follow recognizable sequences. As Mark Twain famously quipped, "History doesn't repeat itself, but it often rhymes." This idea remains contentious. Critics argue that viewing history through repetitive patterns oversimplifies complex realities, as each historical moment is unique and shaped by specific contexts. Proponents, on the other hand, contend that identifying patterns is crucial for developing explanatory frameworks and understanding underlying mechanisms in historical events. This debate is especially relevant to conflict studies, where recurring patterns of violence can help to identify underlying causes and improve predictive models.

This paper seeks to explore whether temporal patterns in history do indeed repeat. Specifically, we aim to determine the extent of repetition in conflict-related events and identify patterns across various historical periods and geographical boundaries. Conflicts throughout history have exhibited considerable diversity, presenting a unique opportunity to investigate the potential for consistent recurrence of violence and fatalities over time and space. By examining the temporal progression of conflict-related fatalities, we aim to uncover inherent cycles or motifs within these events. Identifying such patterns would suggest the existence of intrinsic temporal rhythms in conflicts, potentially shedding light on the fundamental mechanisms driving the dynamics of conflict.

Measuring the repetition of temporal patterns in conflict is valuable for several reasons. Firstly, it helps us to understand the stability and predictability of conflict dynamics. If certain conflict patterns are found to repeat over time, it suggests that there are underlying factors or conditions consistently

© The Author(s), 2026. Published by Cambridge University Press on behalf of EPS Academic Ltd. This is an Open Access article, distributed under the terms of the Creative Commons Attribution-NonCommercial licence (http://creativecommons.org/licenses/by-nc/4.0), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original article is properly cited. The written permission of Cambridge University Press or the rights holder(s) must be obtained prior to any commercial use.

https://doi.org/10.1017/psrm.2026.10119 Published online by Cambridge University Press

driving these events. This understanding of the temporal aspects that contribute to the recurrence of violence can inform theoretical models of conflict. Secondly, identifying redundant patterns enables the development of more effective intervention strategies. By recognizing when specific patterns are likely to reoccur, policymakers and practitioners can time their interventions to pre-empt and mitigate the impacts of conflict. Lastly, measuring redundancy in temporal patterns allows for a comparative analysis across different domains. By comparing the results obtained from conflict data with those in other fields such as seismology, finance, and epidemiology, we can better understand the uniqueness or commonality of conflict dynamics relative to other complex systems.

In this paper, redundancy refers to repeated temporal structure in the data. We distinguish (i) *pattern redundancy*, how often the same shape recurs, measured by the number of sequences grouped into a pattern; and (ii) *outcome redundancy*, how consistently those patterns are followed by similar futures, summarized by the (normalized) entropy of future clusters. High redundancy implies frequent shape recurrence and/ or low-entropy futures; low redundancy implies rare shapes and/or high-entropy futures. Redundancy captures empirical repetition in observed sequences. Formal definitions appear in Appendix E. We recognize that the question of identifying recurring patterns in history is ambitious and complex, so the scope of this paper is intentionally narrow; we focus specifically on conflict fatalities. By focusing solely on temporal sequences of conflict fatalities, we are limited in what we can measure and conclude. Our analysis thus does not account for the broader social, political, or economic factors that often influence conflicts, nor does it address the potential recurrence of non-fatality-related events. Despite this narrow scope, we aim to offer a methodological blueprint for investigating similar questions across different domains. Additionally, our inquiry contributes to the broader discussion on the predictability of historical events. We aim to assess the frequency and extent of temporal sequences' recurrence, questioning whether these sequences recur more frequently than would be expected by chance.

Our methodological approach centers on the concept of entropy, a measure of randomness or uncertainty in a dataset. By calculating the entropy of temporal sequences in conflict data, we can estimate the level of redundancy within these sequences. Low entropy would indicate higher predictability and suggest an underlying data-generating process favoring certain patterns. Conversely, high entropy would point to a more random distribution of patterns and low redundancy. We apply this entropy-based analysis to various aspects of conflicts, including conflict type, region, and temporal resolution. Additionally, we compare our results with findings from other fields such as seismology, finance, and epidemiology to contextualize the predictability of conflict events within the broader predictive landscape. Finally, we also focus on whether certain patterns tend to be consistently followed by certain outcomes. By calculating the entropy of these outcomes, we can quantify the information content available in these patterns.

We engage three longstanding debates in the study of conflict. First, our work speaks to research on the temporal dynamics of political violence, including the escalation and de-escalation of conflict over time. Scholars have long debated whether conflict trajectories are strategic and path-dependent or largely stochastic Fearon (2004); Cederman *et al.* (2010). By measuring the recurrence of specific fatality patterns, we provide empirical evidence that escalation often follows discernible structures, supporting the idea that conflict unfolds in temporally organized ways.

Second, we contribute to the literature on conflict forecasting and early warning. Existing models typically focus on structural risk factors such as political institutions, economic inequality, or demographic pressures Goldstone *et al.* (2010); Ward *et al.* (2010). In contrast, our approach captures endogenous temporal dynamics, identifying patterns that precede different types of conflict trajectories. This allows us to assess not only whether conflict occurs, but how its intensity and duration might evolve Chadefaux (2014).

Third, we engage the debate over whether conflict dynamics are universal or context-specific. Some scholars emphasize the local, idiosyncratic nature of political violence, while others identify recurring "conflict cycles" or macro-patterns that transcend regional boundaries Kriesberg (1998); Harish and

https://doi.org/10.1017/psrm.2026.10119 Published online by Cambridge University Press

Little (2017). Our findings show that certain temporal structures repeat across countries and decades, suggesting that while triggers of violence may differ, the unfolding of conflict often follows shared rhythms.

In all three areas, our contribution is to provide a quantitative method for evaluating temporal regularities in conflict data. Rather than assuming cyclical or patterned behavior, we use entropy-based analysis and Dynamic Time Warping (DTW) to test whether, and under what conditions, such structure exists. In doing so, we offer a methodological bridge between statistical forecasting and qualitative insights about historical repetition in conflict.

Our analysis reveals that conflict patterns are dominated by periods of peace, but when focusing on active conflict periods, these patterns exhibit a spectrum of predictability, ranging from highly unpredictable sequences (similar to white noise) to more stable patterns resembling earthquake data. Predictability varies by conflict intensity, with low-intensity conflicts showing more regular, predictable patterns and high-intensity conflicts exhibiting greater unpredictability. Despite differences in magnitude, conflict patterns remain consistent across regions and decades, with no significant region-specific or time-specific trends. Temporal resolution also affects predictability, with weekly data showing slightly higher entropy than monthly data. Overall, conflict patterns show higher entropy than other fields like inflation or temperature, though less than random data, highlighting both the potential and limits of forecasting conflict dynamics using historical data. Furthermore, our findings suggest a degree of universality in conflict dynamics, as no region or decade-specific patterns were detected. This suggests that conflict patterns may be driven by global processes rather than localized phenomena.

The paper is structured as follows: We begin with a review of relevant literature, followed by a detailed description of our methodology. We then present our findings, categorized by various conflict attributes. The discussion section interprets these results in the context of historical repetition and conflict predictability. We conclude by exploring the implications of our findings for conflict prevention and management strategies, as well as suggesting avenues for future research.

1. Patterns of conflict

The assertion that "those who cannot remember the past are condemned to repeat it" Santayana (1905) suggests that events and patterns in human behavior, societal trends, and global phenomena tend to repeat over time. Thucydides, in his analysis of the Peloponnesian War, observed recurring patterns in state and individual behavior Thucydides (1974). Similarly, Machiavelli argued that political states cycle through periods of order and disorder, driven by inherent human nature and societal dynamics Machiavelli (1532), while Marx famously remarked that "history repeats itself, first as tragedy, second as farce" Marx (1852).

The debate over historical repetition is central to historical and political analysis, with significant implications for understanding historical processes. Proponents of cyclical history argue that recognizing patterns helps uncover causal relationships Mahoney and Rueschemeyer (2003); Goodin and Tilly (2006), while critics argue that this oversimplifies unique events shaped by specific contexts Collier and Mazzuca (2006). This tension reflects a fundamental challenge in historical analysis: the need to balance the particularities of each historical moment with the search for generalizable insights.

For conflict specifically, there are clear reasons to expect temporal redundancy, including historical patterns, structural issues, and predictable trajectories. Much of the literature on conflict focuses on identifying recurring conditions that may lead to violence. From economic instability and inequality Collier (2009) to political transitions and weak governance Fearon and Laitin (2003), research often centers on these repeating factors as drivers of conflict. The concept of "conflict cycles" has been widely discussed in political science and related disciplines. It emphasizes the repetitive patterns of conflict escalation and de-escalation. For example, Kriesberg (1998) describes how conflicts typically

https://doi.org/10.1017/psrm.2026.10119 Published online by Cambridge University Press

transition through phases of intensification and resolution, shaped by internal dynamics and external interventions. Similarly, Acemoglu and Wolitzky (2014) offer an economic perspective, modeling how cyclical patterns of distrust and conflict arise due to incomplete information and coordination failures. In political contexts specifically, Harish and Little (2017) identify recurring cycles of electoral violence, with predictable phases of heightened conflict followed by periods of relative calm. This literature underscores the importance of understanding temporal dynamics in conflicts. In particular, it highlights predictable phases of escalation, peak violence, and de-escalation.

Our study complements this literature by empirically quantifying the degree of redundancy or repetition within conflict-related fatalities using entropy-based measures and DTW. By doing so, we directly assess the extent to which theoretically suggested conflict cycles manifest consistently across different conflicts and temporal contexts, providing a systematic method for evaluating the predictability implied, among others, by cyclical theories.

In methodological terms, our approach builds on established tools for time-series shape matching and information-theoretic assessment of predictability. DTW provides a robust measure of shape similarity under temporal misalignment and has a long record in time-series retrieval and indexing Keogh and Ratanamahatana (2005); we adopt a cross-case, shape-based matching strategy similar to recent work on conflict and migration dynamics Schincariol *et al.* (2025); Schincariol and Chadefaux (2025). To quantify outcome uncertainty, we draw on entropy-based measures widely used to capture regularity and complexity in time series Kumar *et al.* (1986); Pincus (1991). Our contribution complements the conflict forecasting literature, which has emphasized structural risk and algorithmic prediction Hegre *et al.* (2013); Ward *et al.* (2013); Chadefaux (2017), by focusing on the recurrence of temporal shapes and the predictability of their subsequent trajectories. Finally, our cross-domain comparisons are informed by benchmark findings from complex systems: earthquake occurrence is notoriously hard to forecast Geller *et al.* (1997), whereas climate variables often exhibit strong seasonal predictability Papacharalampous *et al.* (2018), with further application-specific evidence in seismology Banna *et al.* (2020).

However, there are also reasons why conflict patterns may not repeat, including strategic behavior, data limitations, and stochastic events. Conflicts involve complex strategic, dynamic, and stochastic human actions that are not entirely dissimilar to stock markets. Any discernible pattern in behavior risks being exploited by actors, in such a way that the repeating pattern disappears Tversky and Kahneman (1974). Even with a comprehensive understanding of conflict causes, understanding the timing, intensity, and evolution of conflicts can be an uphill task. Conflict dynamics are often nonlinear, subject to minute changes in conditions. Furthermore, data is often incomplete or unreliable.

Additionally, the fickle nature of alliances and international relations can drastically alter the conflict landscape. The secretive and delicate nature of diplomatic negotiations, peace treaties, or shifting geopolitical interests can transform a potential conflict into a situation of relative stability or vice versa Holsti (1996). These complexities are challenging to encapsulate within simple temporal patterns. Furthermore, individual decision-making plays an important role in conflicts as well. Influential figures, political leaders, or even single actors can trigger, escalate, or defuse conflicts. Influenced by personal biases, beliefs, or misinformation, these decisions are highly stochastic Tetlock (2005). "Black swan" events—unforeseeable, highly impactful incidents such as sudden technological advancements, natural disasters, or pandemics—can significantly alter conflict dynamics Taleb (2007).

The extent to which data "repeats" varies widely across scientific fields, driven by differences in the nature of phenomena, data availability, and the complexity of models. In seismology, predicting earthquakes is notoriously difficult Geller *et al.* (1997). Similarly, in finance, the efficient market hypothesis suggests that market prices reflect all available information, so that consistently achieving returns above the market average is almost impossible. In contrast, climate models,

https://doi.org/10.1017/psrm.2026.10119 Published online by Cambridge University Press

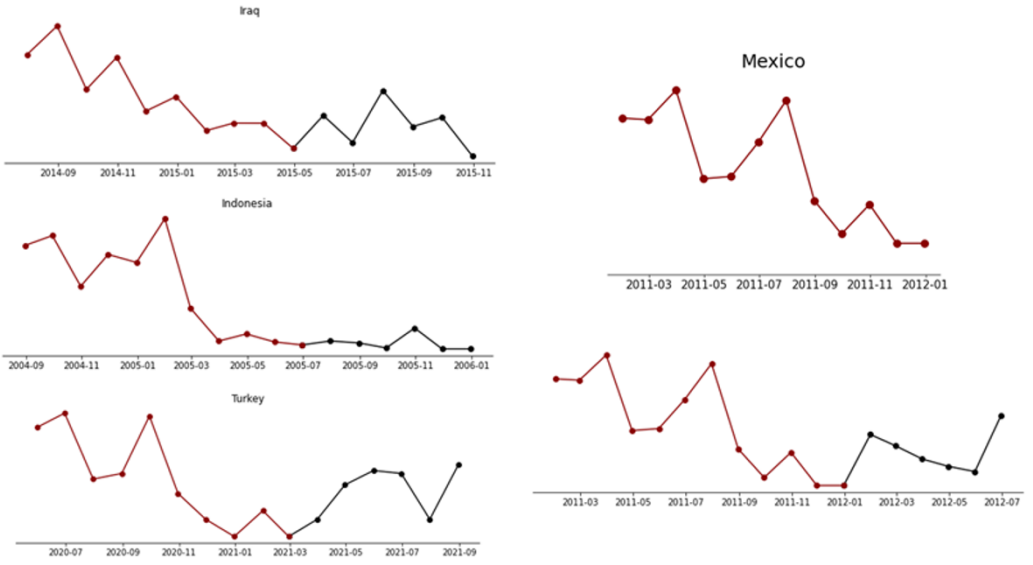


Figure 1. Illustration of prediction using a pattern similarity-based model. The top-right subplot shows the input pattern (red), drawn from the test set. The three subplots on the left display the closest historical matches (in red) along with their true future values (in black). The bottom-right subplot shows the actual future trajectory of the input pattern (black). This example highlights how high redundancy, i.e., repeated historical patterns, enables accurate prediction based on a small number of matched analogs, even when a classical autoregressive model would fail.

which simulate interactions between the atmosphere, oceans, land, and ice, are effective at predicting broad, long-term patterns such as global warming. Epidemiological models, while useful for predicting disease spread, also come with substantive uncertainty, as seen during the COVID-19 pandemic.

To illustrate the implications of redundancy for prediction, Figure 1 presents a concrete example of a recurring pattern in conflict fatalities. The red line shows a 12-month sequence from the test set characterized by a declining trend with intermittent rebounds. A classical time series model, such as ARIMA, trained on the full dataset, would likely extrapolate a continued decrease. However, our pattern similarity-based model identifies a small number of highly similar historical sequences (also in red), all of which are followed by an increase in conflict fatalities (black). The resulting prediction, based on this localized matching, successfully captures this trend reversal. The step-by-step illustration is developed in Appendix F. This example demonstrates how redundancy—i.e., the reappearance of similar temporal structures—can support accurate forecasts even with limited data, in contrast to global autoregressive approaches that average over many dissimilar sequences.

Each field faces unique predictive challenges. Earthquakes and conflicts have identified partial mechanisms for prediction, but only under limited conditions (e.g., yearly forecasts for conflicts or low-magnitude earthquakes) Banna *et al.* (2020). Financial markets' unpredictability stems from their strategic nature Hill and Motegi (2019), while yearly climatic patterns are highly predictable due to strong seasonality Papacharalampous *et al.* (2018), though short-term climate events can still be chaotic.

2. Data

We draw upon conflict casualty data sourced from the Uppsala Conflict Data Program's (UCDP) Georeferenced Event Dataset. This data collection provides an extensive overview of global conflict incidents, each accounting for at least one reported fatality. The dataset encompasses a worldwide

https://doi.org/10.1017/psrm.2026.10119 Published online by Cambridge University Press

scope, capturing daily occurrences from 1989 onwards. Our analysis focuses on the fluctuation of fatality numbers over time.[1] A replication of the main results using the Armed Conflict Location and Event Data (ACLED) Raleigh *et al.* (2010) produced similar observations. These results are presented in Appendix H.

To compare redundancy across various fields, we rely on a range of temporal data sources. Climate data, specifically country-averaged temperature readings, is obtained from ERA-5, a satellite data product of the ECMWF Integrated Forecast System (IFS) under the Copernicus European program Hersbach *et al.* (2020). For seismological data, we turned to the comprehensive earthquake catalog maintained by the United States Geological Survey (USGS). This catalog provides crucial information on seismic events, including location, magnitude, and depth. To ensure meaningful analysis, we extracted country-level earthquake data with a magnitude exceeding five, considering lower magnitude events to be generally predictable Banna *et al.* (2020). In the context of epidemiology, we used data detailing daily Covid cases per country.[2] Historical financial stock return data were obtained from Yahoo Finance and inflation data from the World Bank data, encompassing 192 countries and spanning from January 1970 to December 2022. Ha *et al.* (2023).

3. Methods

In this section, we describe how we extract and analyze temporal patterns in conflict fatalities. The workflow proceeds as follows (see also Fig. 2). First, we split each country's time series into fixed-length windows ("sequences"; e.g., 12 months). Second, we compare sequences using a shape-similarity measure (DTW, introduced below) and group similar sequences into "patterns," allowing them to unfold at different speeds. Third, we assess how informative these patterns are by examining each sequence's "future"—the six subsequent months—and summarizing outcome uncertainty with entropy. We apply this procedure across temporal resolutions and domains to quantify temporal redundancy. A sequence is thus a contiguous window from a country's time series; a pattern is the shared shape across multiple sequences; and a future is the six months that follow a sequence. Formal definitions are provided in Appendix E.

3.1. Pattern extraction

Our first step is the extraction of sequences from each of the larger time series. We refer to sequences as a subset of a time series of a given country that could last 10, 15, 30 weeks or 6, 12, or 18 months. The default length is 12 months. We used a yearly sliding window approach to extract sequences from each country's conflict data, from 1989 to 2023. This yields a total of 4,480 sequences for all 130 countries. For the other fields, we randomly extracted 300 sequences from each dataset. The sample size is limited to 300 because our goal is to compare it with conflict data, rather than conducting a detailed analysis of other variables. Analyzing additional fields would require significantly more computation time, without significantly affecting the entropy scores. Each sequence was normalized to

[1]We underscore that our study focuses on temporal patterns in conflict fatality figures, which imposes certain limitations on our conclusions. Conflict, as a sociopolitical phenomenon, is shaped by a myriad of observable and non-observable variables beyond just fatalities. We have chosen to focus on fatalities as a key measurable output of these multifaceted conflict dynamics. Though this focus might only give a partial view, fatality numbers, as a tangible expression of conflict severity, offer a starting point for our exploration of temporal redundancy and predictability. Our study is not an attempt to simplify conflict to its death toll but rather a starting point for a broader conversation about redundancy in conflict data.

[2]The data was sourced from Our World In Data for the period between March 2020 and April 2023 Mathieu *et al.* (2020).

https://doi.org/10.1017/psrm.2026.10119 Published online by Cambridge University Press

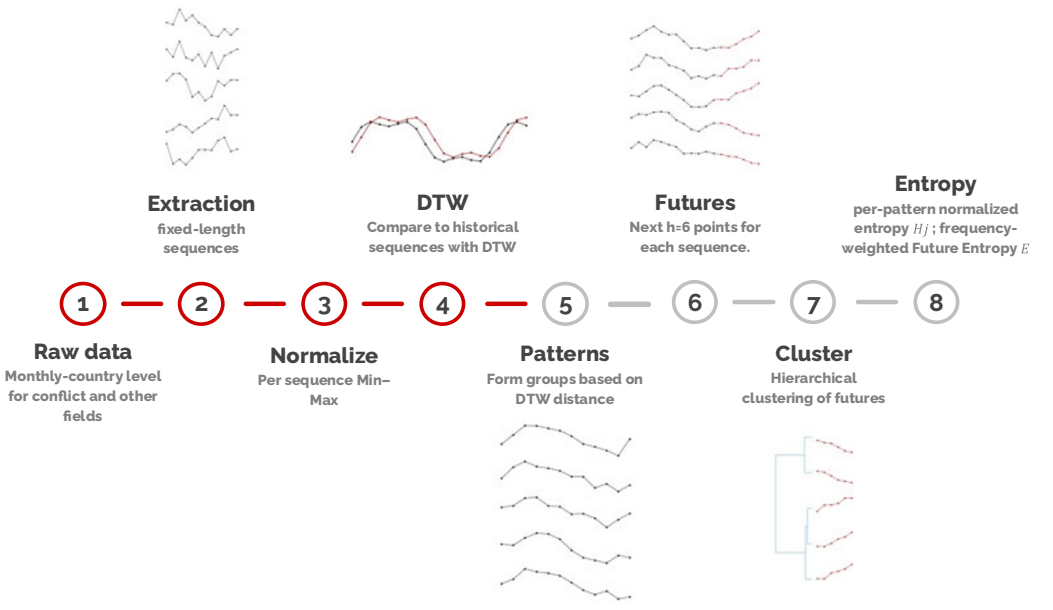


Figure 2. Methodological pipeline. Starting from raw data (1), we extract fixed-length sequences using a sliding window (2). Each sequence is normalized using Min-Max scaling (3). We then compare all sequences using Dynamic Time Warping (DTW) and group similar sequences into patterns based on a distance threshold (4–5). For each sequence, we extract its future (the following 6 months) (6). These futures are grouped using hierarchical clustering (7). Finally, we calculate per-pattern normalized entropy H_j and frequency-weighted Future Entropy E (8), which quantifies outcome predictability.

values between 0 and 1 to allow for comparison across different scales of conflict.[3] These normalized sequences form the basis for our subsequent analysis.[4]

Due to pronounced zero-inflation in the data—57% of the extracted sequences consist of only zeros—we conducted two analyses: one using the full dataset and another focused on sequences with at least one non-zero value (conflict-filtered). This approach helps to examine the dynamics within conflicts without being skewed by long periods of peace.[5]

Additionally, since patterns can vary significantly across different temporal scales, we examined temporal redundancy at multiple levels of aggregation, including weekly and monthly intervals. For

[3]The normalization is performed using the following formula: normalized value $= \frac{\text{original value} - \text{min value}}{\text{max value} - \text{min value}}$ where min value and max value are the minimum and maximum values in the specific 12-month sequence.

[4]Our choice to normalize each sequence independently prioritizes temporal shape over absolute magnitude. This design decision reflects a substantive focus: we aim to identify whether conflict dynamics follow recognizable *trajectories* of escalation, de-escalation, or stagnation, regardless of the absolute number of casualties involved. A low-intensity conflict that doubles in severity over six months and a high-intensity conflict that follows the same proportional increase share a common escalatory dynamic, even though their casualty counts differ by orders of magnitude. This perspective aligns with theories of conflict cycles and path-dependent escalation Kriesberg (1998); Harish and Little (2017), which emphasize temporal rhythms over levels. The normalization choice does entail a trade-off: it discards information about conflict scale. We address this limitation in two ways. First, we stratify results by magnitude (see Fig. 6), demonstrating that low- and high-intensity conflicts exhibit distinct entropy profiles. Second, we replicate key findings using z-normalization (Appendix G), which preserves relative differences within each sequence while still enabling shape-based comparison; results remain substantively unchanged.

[5]The skewed distribution also raises concerns for Min-Max normalization. Mostly flat sequences with a single small non-zero value could, in theory, be matched to sequences with very large values. However, Appendix Figure I.2 shows that transitions from 0 are overwhelmingly small: the modal next value is 1; 65% are below 10 and 80% below 25. This indicates that emergence from peace is typically gradual, not jumpy, meaning flat patterns with only one non-zero value align naturally with patterns of similar magnitude.

https://doi.org/10.1017/psrm.2026.10119 Published online by Cambridge University Press

example, in weather forecasting, predictability is higher on an annual scale—summers are generally warmer than winters, and certain regions consistently experience more rainfall during specific months. However, this predictability diminishes at the weekly scale due to numerous variables affecting day-to-day weather. Similarly, in conflict data, the dynamics may differ depending on the temporal scale. By adopting this multi-scale approach, we are able to explore and compare the predictability of patterns across various levels of temporal granularity.

Once extracted, the sequences are compared to each other using DTW, a metric that captures similar shapes by accommodating variations in the timing of events Keogh and Ratanamahatana (2005). In essence, DTW measures the similarity between two temporal sequences by allowing for non-linear alignments in time. This means that two sequences do not need to unfold at the same speed to be considered similar. For example, a sudden spike in conflict fatalities that occurs over a shorter time period in one conflict could still be matched with a more gradual increase in another, provided the overall shape of the sequences is similar. DTW accomplishes this by "warping" the time axis, aligning points in the two sequences that are similar in value, even if they occur at different times. Details of the DTW methods can be found in Appendix A.

The use of DTW in our analysis is motivated by the need to capture recurring patterns in conflict trajectories that may not align precisely in time. Traditional methods that compare sequences point by point (e.g., Euclidean distance) assume that escalation occurs at the same pace across cases, which is unlikely in real-world conflict dynamics. DTW, by contrast, allows for temporal distortion—matching sequences that are structurally similar but unfold at different speeds or include lags. This flexibility is essential when studying political conflict, where escalation and de-escalation may be triggered by endogenous cycles (e.g., election timing), exogenous shocks (e.g., international intervention), or strategic delays. DTW thus allows us to compare 12-month sequences with sequences ranging from 10 to 14 months, as similar sequences may unfold over different time spans across regions or periods Chadefaux (2022).[6]

Sequences that fall below a predetermined distance threshold are considered part of the same pattern.[7] A pattern is thus a collection of sequences grouped together based on shape similarity, representing a recurring temporal motif in the data (see Appendix E for definitions). While DTW is a non-parametric distance measure and does not rely on strong statistical assumptions (e.g., linearity, stationarity), its validity depends on meaningful temporal alignment and consistent preprocessing. We therefore normalize all time series before analysis and conduct sensitivity checks across distance thresholds to ensure our results are not artifacts of particular choices (see Figures C.2–C.3).

3.2. Assessing the repetition of conflict patterns

Our initial aim is to determine the extent of redundancy in conflict patterns. We use entropy, a concept from information theory, to measure the randomness or uncertainty in the dataset Rabiner (1989). Lower entropy indicates higher redundancy, meaning that specific patterns repeat frequently with similar outcomes, suggesting a more structured and predictable dataset. In contrast, higher entropy reflects greater variability and less predictability.

Maximum entropy occurs in a uniform distribution where all outcomes are equally likely, while minimum entropy occurs when an outcome is nearly certain. Thus, high-entropy patterns indicate

[6]The DTW algorithm measures similarity between temporal sequences of varying lengths using specific parameters. Euclidean distance was used as the base metric between individual points, with a Sakoe–Chiba band of 10% of sequence length to constrain the warping path and reduce computational complexity. To account for different sequence lengths, the DTW distance was normalized by dividing the raw DTW distance by the warping path length. A symmetric step sequence allowed equal movement in diagonal, horizontal, and vertical directions, with strict boundary conditions ensuring that the warping path started and ended at the beginning and conclusion of both sequences.

[7]We set this threshold at 0.3 to balance the similarity of sequences within patterns and the number of sequences identified. To address the issue of parameter sensitivity and clustering validity, we conducted a systematic robustness analysis, the results of which are now included in Appendix C (Figures C.1–C.3).

https://doi.org/10.1017/psrm.2026.10119 Published online by Cambridge University Press

greater unpredictability, whereas low-entropy patterns suggest more predictable patterns. Calculating entropy for temporal patterns in the conflict data allows us to assess and compare redundancy across fields, providing insights into the unique or shared dynamics of conflict relative to other complex systems. The primary assumption underlying this method is the validity of the clustering process, as entropy depends on the frequency and consistency of similar patterns across the dataset. To assess robustness, we conducted a sensitivity check on the clustering parameter (see Fig. C.1).

Entropy complements DTW by measuring the level of uncertainty in the outcomes that follow these matched sequences. Low entropy indicates that similar trajectories often lead to similar results, which is crucial for understanding how much predictive content is embedded in conflict patterns. Together, DTW and entropy allow us to assess both the recurrence of temporal motifs and their relevance for forecasting. To assess the power of the method using formal statistical analysis, two simulation-based power tests are detailed in Appendix D.

Traditional methods like Approximate Entropy (ApEn) and Sample Entropy (SampEn) Pincus (1991) are often used to assess short-term predictability but struggle to capture broader, long-term patterns. Conflict dynamics, influenced by political, economic, and social changes, often unfold over weeks, months, or years, making localized entropy measures insufficient. Our approach shifts the focus to clustering sequences that capture entire temporal structures, allowing us to identify patterns that span longer time frames. In conflict data, significant escalations or de-escalations are typically driven by long-term trends rather than short-term fluctuations, requiring a more comprehensive analysis of temporal structures. For a more detailed discussion of entropy, see Appendix D.

3.3. Evaluating the consistency of patterns' outcomes

Our second key question is whether a pattern consistently leads to the same outcome in the future. To assess the consistency of outcomes that follow specific patterns, we examine *"pattern futures"*, which are defined as the six values following each sequence. By employing hierarchical clustering, an algorithm that successively groups the most similar observations into progressively larger clusters, we are able to group similar "futures" based on the subsequent trajectory of conflict fatalities. This method allows us to discover the natural groupings of outcomes without requiring prior assumptions about how many distinct future scenarios exist. The resulting groups are "clusters." This approach helps identify potential future scenarios that could follow a given pattern and assigns probabilities to each.[8] We compute, for each pattern j, the normalized entropy of its distribution over future clusters:

$$H_j = \frac{-\sum_{i=1}^{p} P(x_i) \log P(x_i)}{\log(N_{\text{sub}}) + 2}, \tag{1}$$

where p is the number of future clusters for pattern j and N_{sub} is the number of sequences assigned to pattern j. The Future Entropy is the frequency-weighted mean of these per-pattern entropies:

$$E = \frac{\sum_{j=1}^{np} f_j H_j}{\sum_{j=1}^{np} f_j}, \tag{2}$$

[8]Hierarchical clustering offers several advantages over other techniques, such as k-means clustering or principal component analysis (PCA) in the context of our study. First, unlike methods such as k-means clustering, where the number of clusters must be predefined, hierarchical clustering does not require us to specify how many distinct clusters exist beforehand. This is particularly advantageous for our analysis of conflict patterns, as we do not have prior knowledge of the number of recurring "futures" that follow each pattern. Second, conflict dynamics, especially the future trajectories following a particular pattern, are often nonlinear and complex. Hierarchical clustering, with its ability to consider nested relationships between data points, captures these complexities more effectively than methods that assume linear relationships, such as PCA. Specifically, we perform hierarchical clustering on pairwise Euclidean distances and cut the dendrogram at $3 \times (\text{horizon/input length}) \times d$, where d is the DTW match cutoff from the similar-cases step. This keeps the clustering threshold consistent with the matching rule, allows slightly more spread for longer horizons, and adds a small buffer so near-duplicate futures are not split.

https://doi.org/10.1017/psrm.2026.10119 Published online by Cambridge University Press

where f_j is the frequency of pattern j and np is the number of distinct patterns. We add 2 to the denominator to avoid division by infinite values or zero (if N_{sub} is 0 or 1). The first equation calculates the normalized entropy for each individual pattern. The second equation, E (called the *Future entropy*) aggregates the entropy values and weighs them by the frequency f_j of each np pattern. Thus, E captures in a single number the average predictability of the future outcomes of the patterns.

Importantly, we note that our cross-domain comparisons focus on statistical properties, not on causal mechanisms. In particular, we do not suggest that conflicts share generative processes with earthquakes or financial markets. Rather, these systems provide well-documented benchmarks for different levels of temporal structure. Climate data exhibit strong seasonality, earthquakes are notoriously unpredictable, and financial markets occupy a strategic but noisy middle ground. Situating conflict within this spectrum helps characterize the statistical nature of its temporal evolution

4. Results

4.1. How often do patterns repeat?

The recurrence of specific patterns across different domains offers insight into the underlying complexity and predictability of each system. Figure 3 presents the frequency of pattern matches within a variety of datasets, ranging from conflict and earthquakes to economic and environmental data. A high concentration of sequences within a few patterns points to repetitive and relatively simple systems. Conversely, datasets like white noise, characterized by randomness, show little to no pattern repetition.

Unsurprisingly, we find that the distribution of conflict data is highly skewed and dominated by the "flat" pattern. This skewness is expected, as most regions of the world are at peace most of the time, leading to numerous sequences with zero conflict events. To gain a clearer understanding of within-conflict dynamics, it is more insightful to analyze a dataset that filters out the non-conflict periods. We refer to this dataset as "conflict-filtered." This filtered dataset focuses on sequences with at least one conflict event, thereby providing a more accurate representation of the patterns and trends within active conflict periods. The conflict-filtered distribution shares similarities with other datasets characterized by recurring events, such as rainfall and filtered earthquakes. Like conflict-filtered, these datasets exhibit patterns where certain sequences repeat with a degree of regularity. In the case of rainfall, the patterns are influenced by seasonal cycles and climatic conditions, while filtered earthquakes reflect the periodicity of seismic activities without the noise of minor tremors. The similarity between conflict-filtered and these natural phenomena suggests that conflicts, like rainfall and significant earthquakes, are governed by underlying factors that drive their recurrence.

4.2. Do patterns lead to predictable outcomes?

The distribution of patterns informs us about the likelihood of observing a specific pattern at a given time. However, a system with few patterns can still be unpredictable if its outcomes are random, while a system with many patterns can be highly predictable if each leads to consistent outcomes. For example, traffic flow on a few routes (low diversity) can be unpredictable due to random events, while a busy train schedule with many routes (high diversity) can remain predictable if trains run consistently on time. Figure 4 displays both the pattern distribution and the distribution of their future outcomes. The x-axis represents pattern frequency, indicating how often each pattern appears in the data. A point on the far right suggests that all sequences are identical, which is nearly the case for conflict data, where most sequences are "flat." Conversely, points on the left indicate patterns that occur only once—as seen with many white noise patterns, which appear infrequently, as expected in random draws. While short temporal patterns can occasionally repeat by chance, the expected frequency of repetition is limited, serving as a useful reference point.

https://doi.org/10.1017/psrm.2026.10119 Published online by Cambridge University Press

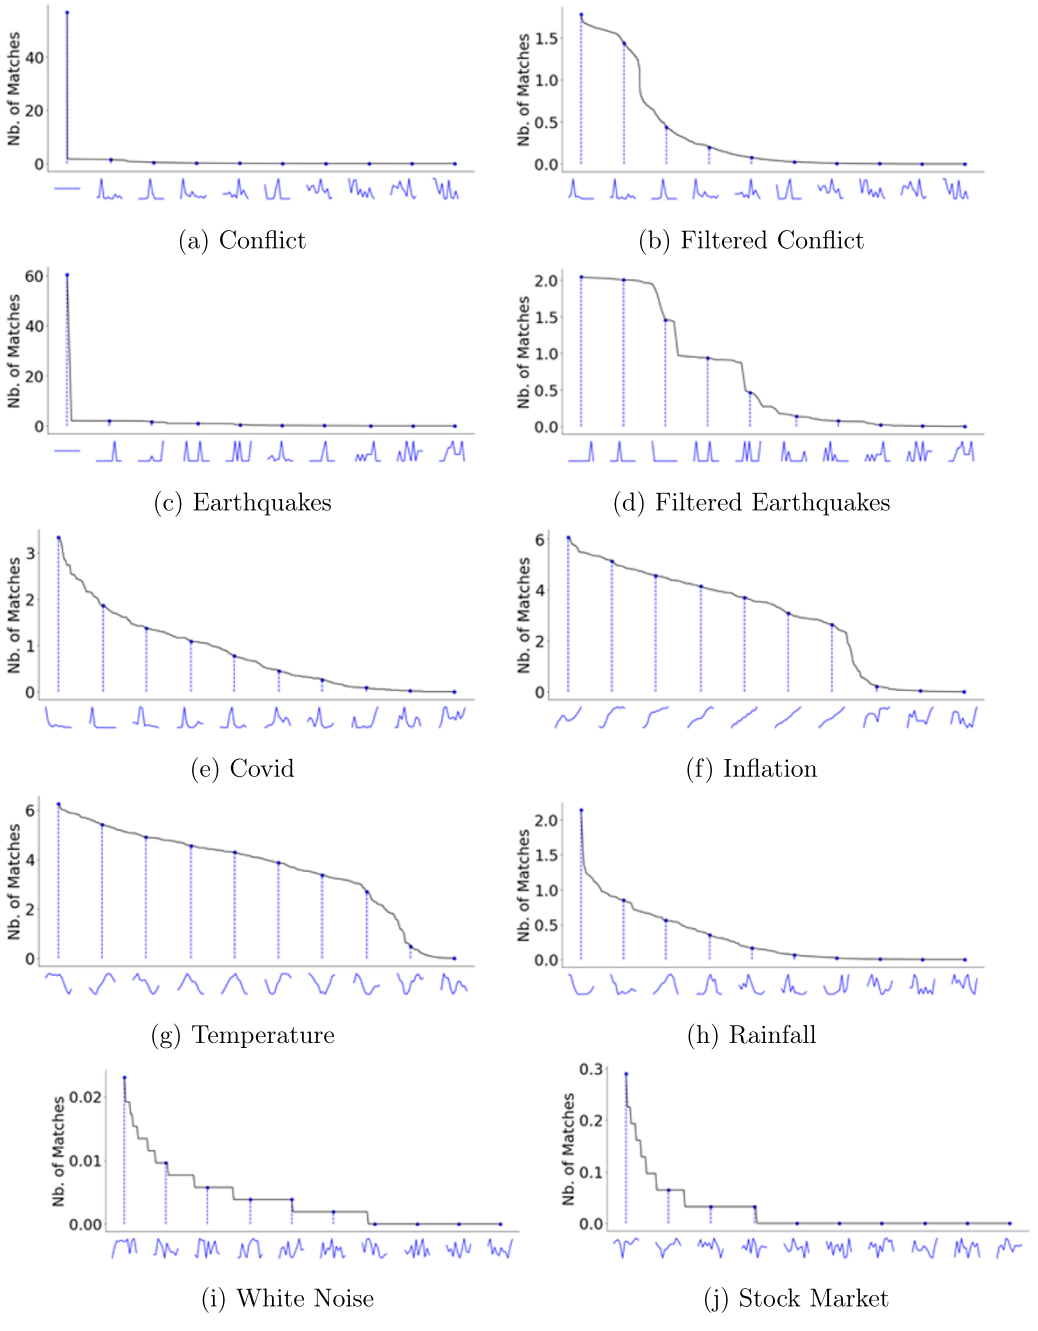


Figure 3. Frequencies of patterns for various datasets with a 12-month window length. The y-axis reports the number of sequences in each pattern. The x-axis displays sample patterns for illustration (there are many more patterns than displayed), ordered from most frequent on the left to least frequent on the right. In the "filtered" datasets, only sequences with at least one non-zero value were retained. (a) Conflict. (b) Filtered Conflict. (c) Earthquakes. (d) Filtered Earthquakes. (e) Covid. (f) Inflation. (g) Temperature. (h) Rainfall. (i) White Noise. (j) Stock Market.

On the *y*-axis, we plot the entropy of each pattern's future, which is calculated by collecting the future six values for each sequence and determining the entropy of their distribution (see above).

https://doi.org/10.1017/psrm.2026.10119 Published online by Cambridge University Press

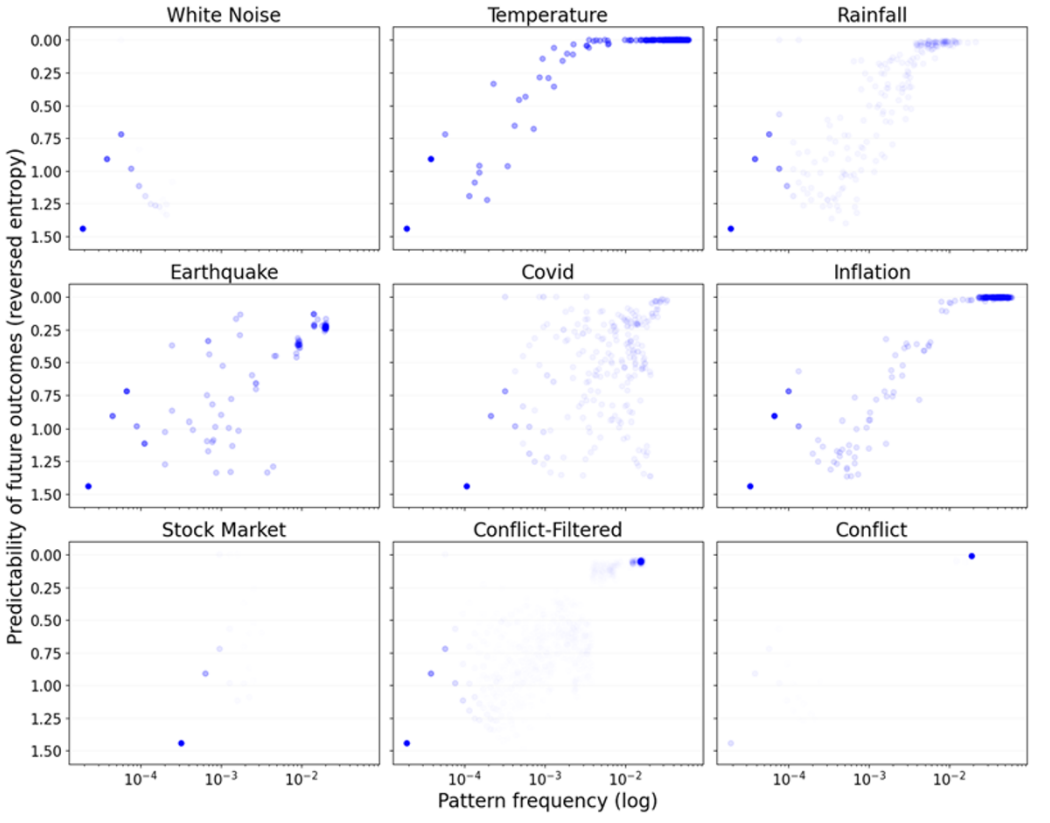


Figure 4. Density of patterns with normalized entropy (reversed scale) on the y-axis (predictability) and the log number of sequences per pattern on the x-axis (pattern frequency). Patterns in the top right show frequent and predictable outcomes, while those in the bottom left are rare and unpredictable.

High entropy on the *y*-axis indicates that these patterns lead to widely varying outcomes, while low entropy suggests more consistent and predictable outcomes.

We observe a clear correlation between the frequency of patterns and the predictability of their future states, which is expected given the finite size of the dataset.[9] However, differences across domains remain clear. For example, temperature data shows a distribution similar to earthquakes and inflation, with most observations clustered in the top right, along with some patterns exhibiting lower entropy along the diagonal. This top-right position corresponds to a high frequency of sequences per pattern and low entropy in their futures. Conflict data is highly concentrated in this corner, mainly due to the prevalence of flat patterns, resembling the pattern seen in earthquakes, though high-magnitude earthquakes occur more frequently than conflict fatalities. Inflation, meanwhile, shows clear and consistent patterns due to its "sticky" nature, while temperature sequences are influenced by seasonal trends, enhancing their predictability.

Additionally, the density distribution for conflict-filtered reveals mixed similar patterns characterized by two separate zones: one in the top-right corner similar to the top right corner of

[9]To see why, consider the extreme case where the dataset consists of a single repeating pattern—in this case, each pattern consistently leads to the same outcome, appearing in the top right of the figure. Conversely, if each pattern is unique (one sequence per pattern), their futures become unpredictable since no comparisons can be made, positioning them in the bottom left. This creates a natural constraint, clustering data along the diagonal.

https://doi.org/10.1017/psrm.2026.10119 Published online by Cambridge University Press

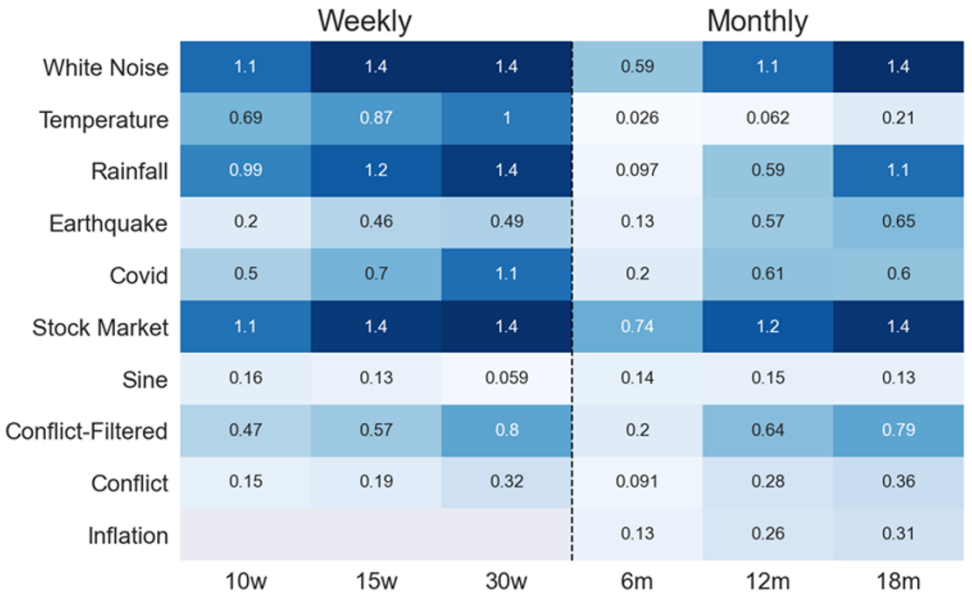


Figure 5. Future entropy of different fields for various temporal resolutions and window lengths. Conflict corresponds to the fatalities dataset extracted from UCDP, and conflict-filtered is the conflict dataset with filtered sequences.

Earthquakes, and another in the bottom-left corner similar to White Noise and Stock Market. The top-right density mass suggests good predictability, as it corresponds to patterns with one or two peak values—commonly found in the earthquake dataset. These patterns generally lead to flat subsequent points, typically representing only a few fatalities that do not escalate into long-term conflict, contributing to low entropy. The second density mass corresponds to patterns with high entropy and few sequences, as in the White Noise and Stock Market dataset (concentrated in the bottom-left density). High entropy suggests these patterns are associated with low predictability. In contrast, Covid data and Rainfall reveal a more dispersed density with spread entropy scores through the diagonal, reflecting the varied dynamics in those domains.

We do not observe recurrent patterns in the bottom-right corner, corresponding to unpredictable futures due to the nature of the data included in the study. Extremely recurrent patterns take two forms in our cases: flat patterns in the case of earthquakes and conflict, and highly seasonal patterns in the case of temperature, rainfall, and inflation (with gradual growth). Those two forms generate predictable futures, leading to low entropy.

We also examined the impact of two parameters: window size and data aggregation frequency. Window size, which determines the number of observations in each sequence (e.g., 6, 12, or 18), directly affects our ability to capture patterns. Smaller windows reveal short-term, detailed fluctuations but may introduce noise, reducing overall predictability. Larger windows, while better suited for capturing long-term trends, can obscure finer details. Data aggregation frequency, which controls how often data is grouped (e.g., weekly or monthly), similarly shapes our analysis. High-frequency data captures rapid changes but increases volatility and unpredictability, while lower-frequency data smooths out short-term noise, highlighting more stable, recurring trends and improving predictability.

Figure 5 presents the Future Entropy values (y-axis in Figure 4) for various fields across different data aggregation configurations, including weekly and monthly frequencies. White noise and stock market data consistently have the highest Future Entropy scores, reflecting their stochastic nature

https://doi.org/10.1017/psrm.2026.10119 Published online by Cambridge University Press

and inherent unpredictability. These results align with previous findings that highlight the random behavior of these time series. A Sine curve was added to establish a lower boundary on the Future Entropy scale, as it is highly predictable.[10] Even with minor random noise introduced, the entropy values remain low across all conditions, demonstrating their high level of predictability.

Conflict has low entropy scores, similar to those of the Sine patterns, mostly due to the prevalence of the flat pattern. Rainfall and temperature data have significantly higher entropy scores at weekly frequencies compared to monthly, highlighting the increased variability in short-term weather patterns. Finally, filtered conflict patterns display entropy scores higher than earthquake ones. In particular, the weekly entropy values for Conflict-filtered are close to their monthly values and similar to Covid scores, indicating a relatively low level of predictability.

4.3. Assessing the universality of conflict patterns: regional, temporal, and scale variations

We have so far analyzed patterns globally, but it is possible that significant variation exists across regions, time periods, and scales of conflict. Are these patterns universal, or do they change depending on the context? The answer matters because it affects how we study and model conflicts. In particular, can a single, overarching model capture conflict dynamics, or do we need region- or magnitude-specific approaches to better understand and predict conflict behavior? We therefore now break down our results for conflict-filtered by magnitude, period, and region.

Of particular interest is whether certain patterns are specific to particular contexts, such as regions, time periods, or conflict sizes. For example, are there patterns unique to Africa or the 1990s, or specific to large-scale conflicts? If so, we would expect sequences of fatalities from Africa to cluster with other African sequences, rather than those from other regions. Similarly, patterns from large wars should differ from those in smaller skirmishes. To test this, we calculated the probability that a pattern (e.g., extracted from a sequence from an African country) would match a sequence from another context (e.g., a country in Asia). If patterns are not specific to regions or periods, we would expect this probability to be consistent across regions. In other words, the likelihood of two sequences belonging to the same pattern would not depend on where or when they occurred. Figure 6 represents the distribution of these probabilities by conflict magnitude, decade, and region. For example, the top-right cell of Fig. 6(b) corresponds to:

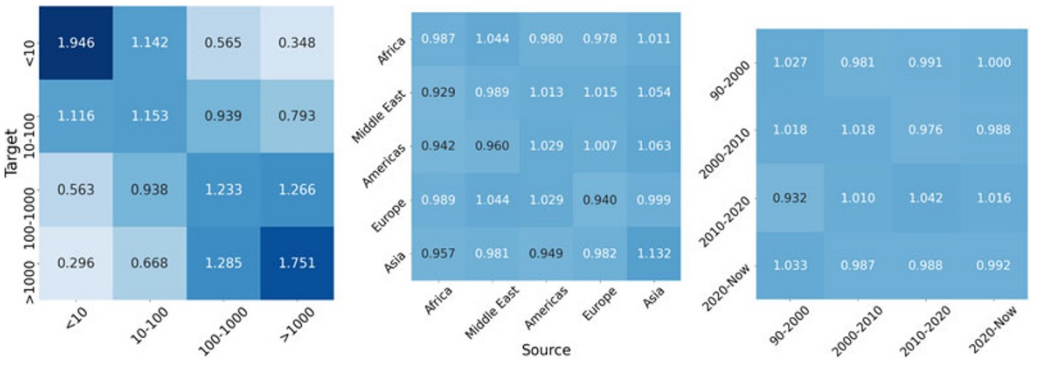
$$R_{top-right} = \frac{P(\text{pat}(i) = \text{pat}(j) \mid i \in \text{Asia}, j \in \text{Africa})}{P(j \in \text{Africa})} \tag{3}$$

where $pat(i)$ denotes the pattern to which sequence i belongs. This ratio expresses the probability that a sequence from Asia shares the same pattern as one from Africa, normalized by the probability that sequence j belongs to Africa. Figure 6(a) shows that the highest proportion of matching sequences is found within the same magnitude sequence category. This indicates that patterns derived from low-magnitude conflicts predominantly match with other low-magnitude conflict sequences. This finding suggests that distinct conflict patterns are specific to each magnitude category and are not generally shared across different magnitudes. On the contrary, in Fig. 6(b) and (c), the distribution is consistent regardless of the region or decade where the pattern is extracted. African and the 90s decade sequences are prevalent, aligning with the distribution of non-zero sequences. It suggests that conflict patterns are comparable across space and time. In addition, Fig. I.3 shows the distribution across conflict types, as defined by UCDP: state-based conflict, one-sided, and non-state. The main conclusion is that, as with time and region, patterns remain comparable across conflict types.

These results are supported by the distribution of entropies (of futures) presented in Fig. 7. Decadal and regional densities have similar multimodal distributions, aligning with the general

[10]The Sine dataset was generated by simulating sinusoidal functions with consistent frequency and adding random noise. The number of series matches the number of countries in the conflict dataset, and the length of each series corresponds to the number of months in the conflict dataset.

https://doi.org/10.1017/psrm.2026.10119 Published online by Cambridge University Press



(a) Scale (b) Region (c) Decade

Figure 6. Distribution of conditional probability ratios: for each category (rows), we compute the probability that a similar occurrence shares the same pattern (columns), and divide it by the unconditional probability of observing that category in similar occurrences. We classify patterns based on four fatality magnitude scales (left) (low: < 10, mid-low: 10-100, mid-high: 100-1000, and high: > 1000 per year), the five regions/continents coded in the UCDP dataset (center), and the available four decades in the dataset (right). Values greater than 1 indicate that the category is overrepresented relative to chance, whereas values below 1 indicate underrepresentation. For example, the top-right cell corresponds to cases where the row category is conflicts with fewer than 10 fatalities, and the column category is conflicts with more than 1,000 fatalities. Its value of 0.348 indicates that sequences from small conflicts are only about 35% as likely as expected by chance to share a pattern with sequences from very large conflicts. (a) Scale. (b) Region. (c) Decade.

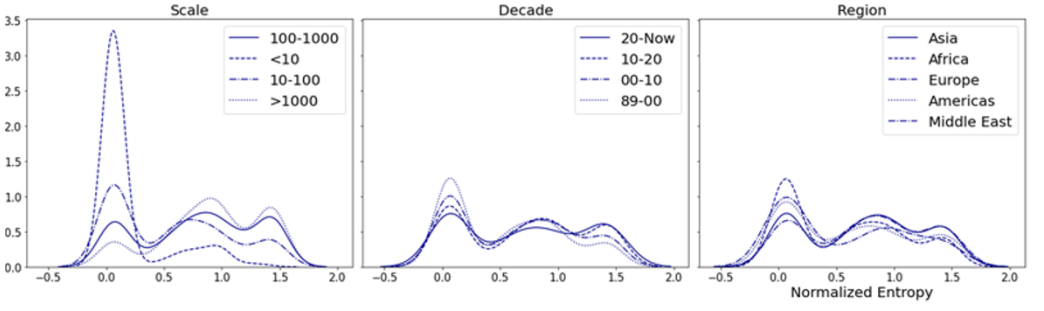


Figure 7. Density of normalized entropy of futures (Conflict-filtered data), per magnitude (left), decade (center), and region (right).

conflict-filtered density. In contrast, the scale category shows disparate results. Patterns extracted from low-intensity conflict sequences are characterized by low entropy rates, mainly due to the high number of sequences that only have non-zero values in one or two months, followed by zero values. These are similar in density and pattern shape to Earthquakes. Patterns from mid-high and high-magnitude sequences have comparable densities, concentrated around high entropy values. This could be partly explained by the small number of observations in history with a large number of fatalities, but also by the diversity of the few sequence futures that lead to uncertain following values. Average values for each magnitude subgroup are shown in the Appendix in Fig. I.1.

5. Discussion

These findings have significant implications in conflict studies. For conflict researchers, our findings highlight the need for different analytical approaches depending on the pattern features. Conflicts associated with rare or unique historical patterns, linked to high-intensity events, have high entropy

https://doi.org/10.1017/psrm.2026.10119 Published online by Cambridge University Press

values. In such cases, autoregressive models alone are insufficient, and additional real-time or contextual variables are necessary to improve forecasting accuracy. Researchers studying these conflicts should prioritize real-time data collection and analysis of rapidly changing situations. They might also explore the use of complex systems theory or chaos theory to better understand these highly volatile scenarios. In contrast, more predictable conflicts could be studied with an emphasis on long-term, structural factors. Here, researchers might employ longitudinal studies, focusing on gradual changes in socio-economic conditions, political structures, or demographic shifts that contribute to conflict patterns over time.

While our findings confirm that conflict dynamics are often difficult to predict, we stress that this unpredictability is not uniform. Some conflict sequences—particularly those involving low-intensity violence or recurring temporal motifs—exhibit high redundancy and are substantially more forecastable. Others, especially high-intensity or escalation-prone conflicts, display far greater entropy and lower pattern recurrence.

This heterogeneity is theoretically meaningful. It echoes long-standing arguments in conflict research that distinguish between slow-building, structurally driven violence and sudden-onset or strategically reactive escalation Fearon (2004); Goldstone *et al.* (2010). Our analysis provides empirical grounding for this distinction, offering a framework to assess when conflict trajectories follow recognizable paths and when they resemble "black swan" events (outcomes that are rare, impactful, and difficult to anticipate Taleb (2007)).

Moreover, by quantifying the temporal predictability of conflict events, our approach contributes to the growing literature on bounded rationality and limits of foresight in complex political environments Tetlock (2005). The ability to distinguish between forecastable and non-forecastable sequences may help refine the use of early warning systems, improve resource targeting, and generate testable hypotheses about the political and institutional factors that shape these temporal dynamics.

In this sense, our contribution is not just that conflict is difficult to predict in general—a point long recognized, at least theoretically—but that patterns of predictability vary systematically across conflict types. Recognizing and measuring this variation offers a path forward for both theoretical modeling and policy-relevant forecasting in political science.

Yet, important limitations that should be addressed. One is our focus on autoregressive patterns, without fully exploring the broader range of covariates—such as political, economic, and social factors—that may influence conflict dynamics. Future research could extend this study by incorporating such variables to develop more comprehensive models, particularly for the first class of unpredictable patterns. This could lead to a better understanding of the factors that drive sudden escalations in conflict fatalities and improve forecasting accuracy in such scenarios.

The quality and source of the data also present challenges. The reliance on news reports for conflict fatalities data introduces uncertainties, particularly regarding the accuracy of reported dates and death tolls. These potential discrepancies could obscure the detection of true patterns and may affect the robustness of our findings. Additionally, the timeframe of the study, beginning in 1989, may not fully capture long-term shifts in conflict dynamics or account for pre-existing trends.

A key limitation of our study is its reliance solely on conflict-related fatalities. While this choice excludes non-lethal forms of violence or protest, it enables us to focus on a consistently reported and substantively meaningful outcome: the severity of conflict. Fatalities reflect escalation in a way that is both observable and temporally structured, allowing for robust comparisons across space and time. Nonetheless, future extensions could incorporate broader event data, such as those from ACLED ACLED (2025), to capture a wider range of conflict dynamics, including protests, riots, or threats. Doing so would enrich the analysis but also introduce new challenges, such as greater heterogeneity in event types and reduced comparability in magnitude.

A key takeaway from our findings is the consistency of conflict patterns across time and space, suggesting that while conflicts may differ in scale, their underlying dynamics are often comparable. This offers promise for enhancing predictive accuracy by leveraging historical data to anticipate future

https://doi.org/10.1017/psrm.2026.10119 Published online by Cambridge University Press

trends in conflict regions. However, the unpredictability of high-variation patterns remains a significant challenge for forecasting efforts. This finding underscores the need for further refinement of models that can account for sudden escalations in conflict intensity.

Acknowledgements. The authors used ChatGPT to assist with grammar and spelling corrections during the preparation of this work. They reviewed and edited all content and take full responsibility for the final manuscript.

Data availability statement. The authors declare that they have no competing interests.

Declaration of interests. The authors declare that they have no competing interests.

Funding statement. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program under grant agreement No. 101002240.

Supplementary material. The supplementary material for this article can be found at https://doi.org/10.1017/psrm.2026.10119. To obtain replication material for this article, https://url.uk.m.mimecastprotect.com/s/k3QwCoQRRfXpN2x5h1fWfpBUho?domain=doi.org.

References

Acemoglu D and Wolitzky A (2014) Cycles of conflict: An economic model. *American Economic Review* **104**(4), 1350–1367.
ACLED (2025) ACLED: Armed Conflict Location & Event Data project (accessed 6 December 2025). https://acleddata.com/.
Al Banna MH, Taher KA Kaiser MS, Mahmud M, Rahman MS, Hosen AS and Cho GH (2020) Application of artificial intelligence in predicting earthquakes: State-of-the-art and future challenges. *IEEE Access* **8**, 192880–192923.
Cederman L-E, Wimmer A and Min B (2010) Why do ethnic groups rebel? New data and analysis. *World Politics* **62**(1), 87–119.
Chadefaux T (2014) Early warning signals for war in the news. *Journal of Peace Research* **51**(1), 5–18.
Chadefaux T (2017) Conflict forecasting and its limits. *Data Science* **1**(1-2), 7–17.
Chadefaux T (2022) A shape-based approach to conflict forecasting. *International Interactions* **48**(4), 633–648.
Collier P (2009) *Wars, Guns, and Votes: Democracy in Dangerous Places.* New York: HarperCollins.
Collier RB and Mazzuca S (2006) Does history repeat? In Goodin RE and Tilly C (eds), *The Oxford Handbook of Contextual Political Analysis*. Oxford, UK: Oxford University Press, pp. 472–489.
Fearon JD (2004) Why do some civil wars last so much longer than others? *Journal of Peace Research* **41**(3), 275–301.
Fearon JD and Laitin DD (2003) Ethnicity, insurgency, and civil war. *American Political Science Review* **97**(1), 75–90.
Geller RJ, Jackson DD, Kagan YY and Mulargia F (1997) Earthquakes cannot be predicted. *Science* **275**(5306), 1616–1616.
Goldstone JA, Bates RH, Epstein DL, Gurr TR, Lustik MB, Marshall MG, Ulfelder J and Woodward M (2010) A global model for forecasting political instability. *American Journal of Political Science* **54**(1), 190–208.
Goodin RE and Tilly C (eds.) (2006) *The Oxford Handbook of Contextual Political Analysis.* Oxford: Oxford University Press.
Harish SP and Little AT (2017) The political violence cycle. *American Political Science Review* **111**(2), 237–255.
Hegre H, Karlsen J, Nygård HM, Strand H and Urdal H (2013) Predicting armed conflict, 2010–2050. *International Studies Quarterly* **57**(2), 250–270.
Hersbach H, Bell B, Berrisford P, Hirahara S, Horányi A, Muñoz-Sabater J, Nicolas J, Peubey C, Radu R and Schepers D *et al.* (2020) The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society* **146**(730), 1999–2049.
Hill JB and Motegi K (2019) Testing the white noise hypothesis of stock returns. *Economic Modelling* **76**, 231–242.
Holsti KJ (1996) *The State, War, and The State of War.* Cambridge: Cambridge University Press.
Ha J, Kose MA and Ohnsorge F (2023) One-stop source: A global database of inflation. *Journal of International Money and Finance* 137, 102896.
Keogh E and Ratanamahatana CA (2005) Exact indexing of dynamic time warping. *Knowledge and Information systems* **7**, 358–386.
Kriesberg L (1998) *Constructive Conflicts: From Escalation to Resolution.* Lanham, MD: Rowman and Littlefield.
Kumar U, Kumar V and Kapur JN (1986) Normalized measures of entropy. *International Journal Of General Systems* **12**(1), 55–69.
Machiavelli N ([1532] 2003) *The Prince.* London: Penguin Classics.
Mahoney J and Rueschemeyer D (eds.) (2003) *Comparative Historical Analysis in the Social Sciences.* Cambridge: Cambridge University Press.
Marx K ([1852] 1963) *The Eighteenth Brumaire of Louis Bonaparte.* New York: International Publishers Co.
Mathieu E, Ritchie H, Rodés-Guirao L, Appel C, Giattino C, Hasell J, Gavrilov D, Macdonald B, Dattani S, Beltekian D, Ortiz-Ospina E and Roser M (2020) Coronavirus Pandemic (COVID-19). *Our World in Data.* Available at https://ourworldindata.org/coronavirus.

https://doi.org/10.1017/psrm.2026.10119 Published online by Cambridge University Press

Papacharalampous G, Tyralis H and Koutsoyiannis D (2018) Predictability of monthly temperature and precipitation using automatic time series forecasting methods. *Acta Geophysica* **66**(4), 807–831.

Pincus SM (1991) Approximate entropy as a measure of system complexity. *Proceedings of the National Academy of Sciences* **88**(6), 2297–2301.

Rabiner LR (1989) A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE* **77**(2), 257–286.

Raleigh C, Linke A, Hegre H and Karlsen J (2010) Introducing ACLED: An armed conflict location and event dataset. *Journal of Peace Research* **47**(5), 651–660.

Santayana G (1905) *The Life of Reason: The Phases of Human Progress.* New York: Charles Scribner's Sons.

Schincariol T and Chadefaux T (2025) Temporal patterns in migration flows evidence from South Sudan. *Journal of Forecasting* **44**(2), 575–588.

Schincariol T, Frank H and Chadefaux T (2025) Accounting for variability in conflict dynamics: A pattern-based predictive model. *Journal of Peace Research* **62**(6), 2052–2069.

Thucydides ([c. 400 BCE]1974) *History of the Peloponnesian War.* London: Penguin Classics.

Taleb NN (2007) *The Black Swan: The Impact of the Highly Improbable.* New York: Random House.

Tetlock PE (2005) *Expert Political Judgment: How Good Is It? How Can We Know?* Princeton, NJ: Princeton University Press.

Tversky A and Kahneman D (1974) Judgment under uncertainty: Heuristics and biases. *Science* **185**(4157), 1124–1131.

Ward MD, Greenhill BD and Bakke KM (2010) The perils of policy by p-value: Predicting civil conflicts. *Journal of Peace Research* **47**(4), 363–375.

Ward MD, Metternich NW, Dorff CL, Gallop M, Hollenbach FM, Schultz A and Weschle S (2013) Learning from the past and stepping into the future: Toward a new generation of conflict prediction. *International Studies Review* **15**(4), 473–490.

https://doi.org/10.1017/psrm.2026.10119 Published online by Cambridge University Press

Cite this article: Schincariol T and Chadefaux T (2026) Echoes of conflict: quantifying redundancy in fatality time series. *Political Science Research and Methods*, 1–18. https://doi.org/10.1017/psrm.2026.10119