Estimating Strategic Response in Sequential Data

Danny Klinenberg[†] Eli Berman[§] Esteban F. Klor[¶]

March 2026

Abstract

Social scientists often estimate how actors respond to each other. However, data is usually aggregated temporally, into days, weeks, or years, when response occurs at higher frequency. How does temporal aggregation affect standard regression estimates? We show that studying interactions aggregated (or disaggregated) into predetermined intervals, rather than the actual response, can distort estimates, generating attenuation, amplification and even reverse signs. We provide analytic derivations, simple examples, and empirical Monte Carlo simulations. We conclude by examining how temporal aggregation can distort our understanding of the Israel-Gaza conflict from 2007 to 2017.

[†]IGCC; dklinenberg@ucsd.edu [§]UC San Diego, IGCC and NBER; elib@ucsd.edu [¶]The Hebrew University of Jerusalem, INSS, and CEPR; eklor@huji.ac.il

JEL Codes: C15, C32, D74

[1]We thank seminar participants at ASSA National Security Seminar, Antoine Deeb, Spencer Sween, Sarah Papich, Richard Startz, Jamel Saadaoui, and Daniele Paserman for valuable feedback and suggestions. We acknowledge support of the Minerva Research Institute (FA9550-23-1-0437). Views expressed are of the authors and do not reflect the official policy or position of the Department of War or U.S. Government.

1. Introduction

> *And have no doubt — we will hold all those responsible to account at a time and in a manner [of] our choosing.*
>
> • Former President Biden discussing the United States' response to missile and drone attacks targeting United States military installations resulting in three casualties (2024).

How do political actors respond to one another during ongoing strategic interactions such as violent conflict, trade retaliation, cyber operations, or campaign competition? A large empirical literature studies these dynamics using time-series models estimated on data aggregated into fixed calendar intervals—typically days, weeks, or months. Vector autoregressions (VARs), in particular, have become a standard tool for estimating retaliation, escalation, and deterrence across a wide range of political settings.[1]

Yet strategic behavior rarely unfolds according to fixed calendar intervals. Actors respond when opportunities arise, when capabilities permit, or when political incentives demand it. As a result, responses often occur at irregular intervals that may be much shorter—or occasionally longer—than the temporal units used in empirical analysis. Researchers, therefore, face a fundamental mismatch between the timing of strategic actions and the timing of the data used to measure them. While earlier studies were often constrained to temporally aggregated data, the increasing availability of high-frequency event data makes it possible, and necessary, to model strategic interaction at the level at which it occurs.

This paper shows that this mismatch can fundamentally distort empirical inference about strategic interaction. When actions occur in action time—that is, as a sequence of moves and countermoves between actors—but are analyzed using data aggregated in calendar time, standard regression approaches can produce misleading estimates. We demonstrate that

[1] See, for example, Enders and Sandler (1993) on counterterrorism policy; Freeman et al. (1989) and Freeman (1989) on political relationships; Box-Steffensmeier et al. (2009) on campaign dynamics; Barberá et al. (2019) on legislative agenda-setting; and Jaeger and Paserman (2008) and Haushofer et al. (2010) on conflict dynamics.

temporally aggregated data can cause estimated responses to appear attenuated, amplified, or even reversed in sign relative to the underlying strategic reaction functions. Moreover, models estimated using aggregated data can produce impulse response functions that suggest prolonged retaliation cycles even when actors respond only to the most recent action.

The intuition is straightforward. Strategic interactions are sequential: one actor responds to another's previous move. When researchers aggregate these actions into fixed calendar intervals, multiple actions may be combined within a single observation period, while other actions may span multiple periods. Temporal aggregation therefore obscures the order and timing of moves that define strategic interaction. The resulting empirical estimates reflect the statistical properties of the aggregation process rather than the underlying strategic behavior. Hence, researchers may infer extended retaliation dynamics or delayed responses that are artifacts of aggregation rather than features of the strategic process.

We develop this argument in three steps. First, we provide an analytical framework that distinguishes between action time, which records the sequence of strategic moves, and calendar time, which records outcomes aggregated into fixed intervals. We show formally that estimating standard vector autoregressions using temporally aggregated data produces coefficients that are linear combinations of the underlying action-level reaction parameters. These coefficients may therefore be distorted in magnitude or even sign.

Second, we demonstrate the implications of this distortion through Monte Carlo simulations. In simulated data where actors respond only to the immediately preceding action—a simple Markovian strategic process—VAR models estimated on daily data frequently (mis)identify multiple statistically significant lagged responses. The optimal lag length selected by common criteria such as the Bayesian Information Criterion closely tracks the average waiting time between actions rather than the true structure of strategic interaction. Likewise, impulse response functions derived from aggregated data often display prolonged retaliation effects despite the underlying process containing only immediate responses.

Third, we illustrate the empirical relevance of these distortions using high-frequency data on Israeli and Gazan attacks between 2007 and 2017. The underlying data records attacks at five-minute intervals, allowing us to reconstruct strategic behavior at the action level. When the data are aggregated to the daily level—following common practice in the conflict literature—standard VAR analysis suggests retaliation cycles lasting several weeks. Yet the underlying action-level data reveal that most episodes of violence consist of short bursts of activity lasting hours or days, separated by periods of calm. Simulations calibrated to the empirical data indicate that these prolonged impulse responses arise primarily from temporal aggregation, particularly when actions span multiple days.

Our findings contribute to two literatures. First, we extend the econometric literature on temporal aggregation, which has primarily examined how aggregation affects estimates of relationships between economic indicators (Working, 1960; Zellner and Montmarquette, 1971; Brewer, 1973; Tiao and Wei, 1976; Geweke, 1978; Wei, 1978; Freeman, 1989; Marcellino, 1999; Jordá, 1999; Silvestrini and Veredas, 2008). We show that temporal aggregation poses additional challenges in strategic settings, where actors choose both the magnitude and the timing of their responses. Second, we contribute to the empirical study of political conflict and other strategic interactions by demonstrating that commonly used time-series methods may mischaracterize the dynamics of retaliation when applied to temporally aggregated data. In particular, we show how strategic behavior and temporally aggregated observations may affect our understanding of the Israel-Palestinian conflict. Previous work on this conflict used VAR on daily units of observation (Jaeger and Paserman, 2008), then proceeded to study different damage metrics (Haushofer et al., 2010; Golan and Rosenblatt, 2011), and different functional forms (Asali et al., 2017). This paper goes a step further and studies how the temporal unit of observation affects empirical findings, using our high-frequency data on the conflict.

More broadly, our results highlight the importance of matching empirical models to the temporal structure of strategic behavior. Today's wars are documented nearly in real time allowing researchers to map actions to reactions.[2] Researchers should prioritize leveraging this additional information by modeling interactions in action time rather than calendar time. When only aggregated data are available, researchers should interpret estimated retaliation dynamics with caution, as these patterns may reflect artifacts of temporal aggregation—attenuation, amplification or sign reversals—rather than genuine strategic responses.

The remainder of the paper proceeds as follows. Section 2 motivates the problem using high-frequency data from the Israeli–Gaza conflict. Section 3 introduces the analytical framework linking action time and calendar time. Section 4 derives the distortions that arise when strategic interactions are estimated using temporally aggregated data. Section 5 presents Monte Carlo simulations illustrating these distortions. Section 6 applies the framework to the Israeli–Gaza conflict. Section 7 concludes.

2. Motivating Example: Strategic Interaction in the Israel-Gaza Conflict

Empirical studies of violent conflict frequently estimate retaliation dynamics using time-series models applied to temporally aggregated data. A common approach aggregates attacks into daily or weekly observations and estimates vector autoregressions (VARs) to evaluate how violence by one actor affects subsequent violence by another. While these methods have become standard in the conflict literature, the use of fixed calendar intervals may obscure the timing of strategic responses.

To illustrate this issue, we begin with high-frequency data from the Israeli–Gaza conflict. Between 2007 and 2017 the United Nations recorded attacks across the Israeli–Gazan border at five-minute intervals (Berman et al., 2024). These reports include injuries and fatalities as well as munitions launched by both sides, including Israeli airstrikes, shelling, small-arms fire, and incursions, and Gazan rockets, mortars, and small-arms attacks.

[2] This is especially relevant during the US-Israel-Iran war of 2026. For example, the Israelis conducted an attack on the Iranian gas and oil sites in South Pars gas field and Asaluyeh oil refinery. The Iranians broadcast they would respond by attacking Saudi Arabian oil fields (Associated Press, 2026).

Figure 1 plots the number of Israeli and Gazan attacks per day over the sample period. The magnified portion highlights December 2013, a randomly selected month. Even at the daily level, the data display pronounced bursts of activity followed by periods of relative calm. These bursts often consist of multiple attacks occurring within a short period of time, reflecting sequences of moves and countermoves between the two sides. The grey bars highlight days in which neither Israel nor Gaza attacked. There were no attacks during about a quarter of the days in the sample.

Three features of these data highlight the challenges of studying conflict dynamics using daily observations.

First, multiple attacks often occur within a single day. Over the full sample period, there were about 3.97 attacks per day, with no attacks on 24.5% of days and only one attack on 23.7%. Conditional on an attack occurring, there were on average 5.26 attacks per day. Aggregating to the daily level, therefore, compresses sequences of attacks that may represent distinct strategic moves and countermoves.

Second, response times vary substantially. Both Israel and Gaza display a large variation in the time between attacks relative to their mean response times. On average, Israel reacted to the previous attack after 40 minutes, and Gazan militants responded after about 50 minutes. That said, the Israeli response standard deviation is 0.62 days, and the Gazan is 0.49 days. Both sides reacted within a minute in about 0.2% of cases; Israel took longer than a day to respond 8% of cases, whereas Gazans did so 5% of cases.

Third, violence is interspersed between periods of calm. When an attack occurred, the violence ended within 24 hours about half the time and was followed by at least 6.6 days of calm. Statistical approaches that impose fixed calendar intervals implicitly average over this variation, potentially combining multiple strategic moves into a single observation period, while other actions may span multiple periods. As we show below, this aggregation obscures the timing of strategic responses and substantially distorts standard time-series estimates of retaliation dynamics.

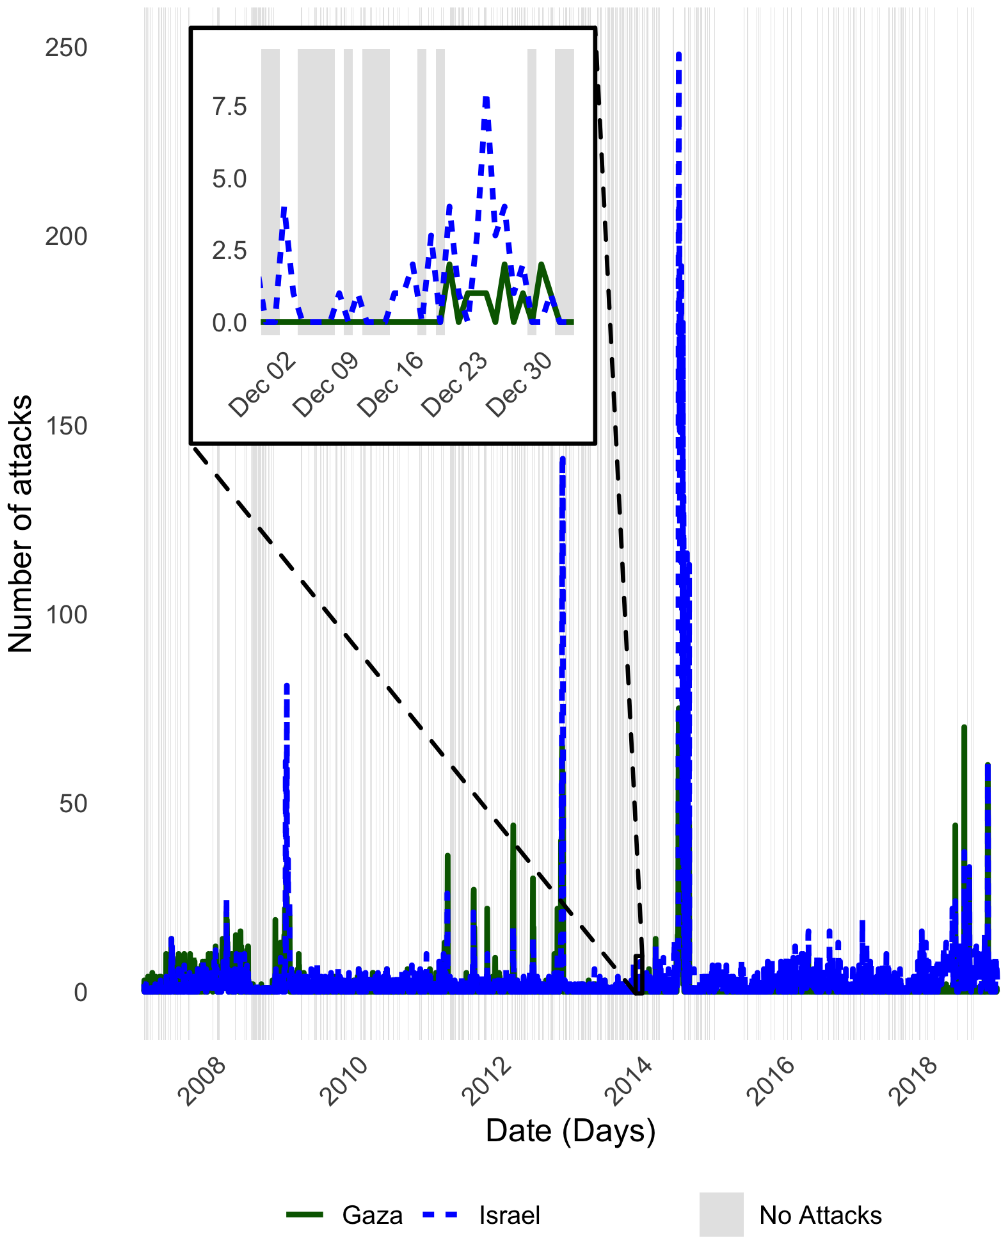

Figure 1: Israeli and Gazan Daily Attacks
Note: Data from Berman et al. (2024).

2.1 Estimating Impulse Response Functions

We estimate impulse response functions employing the methods from Jaeger and Paserman (2008) using data from Berman et al. (2024). Following their work, we first temporally aggregate the data to the daily level. While Jaeger and Paserman (2008)'s analysis focuses on fatalities, subsequent research shows that both fatalities and projectile launches reflect strategic behavior (Haushofer et al., 2010). We therefore construct an index that converts munitions into expected casualties. Specifically, we regress fatalities and injuries on munitions with no intercept and use the fitted values as a measure of intended damage.[3] This procedure yields two time series—Israeli expected damage and Gazan expected damage—measured in casualty equivalents.

We then estimate the orthogonal impulse response functions from a VAR with 14 lags.[4] Figure 2 plots the resulting impulse response functions for Israel and Gaza with 95% confidence intervals. Each response traces the evolution of expected violence for 60 days after an initial attack.

The results suggest substantial asymmetry in retaliation dynamics. A Gazan attack appears to generate sustained responses lasting nearly 30 days, during which Israeli damage inflicted on Gaza remains above its mean level. In contrast, Gazan responses to an initial Israeli attack appear weaker and shorter-lived. Consistent with this pattern, Granger causality tests indicate that damage inflicted on Israel weakly predicts subsequent damage inflicted on Gaza, while the reverse relationship is not statistically significant.

Taken at face value, these estimates imply that responses to a single attack persist for 30 to 40 days. Yet this interpretation appears inconsistent with basic features of the underlying data: 99% of all violent episodes in the dataset end within 30 days, and most episodes are far shorter.[5] On average, violent activity lasts approximately 0.4 days before at least one day

[3] Specifically, we regress deaths $+ 0.2 \times$ injuries on munitions type with no intercept and use the fitted values as our damage measure.

[4] Fourteen lags are selected by the Bayesian Information Criterion (BIC). The full reduced-form VAR results are reported in Table A1 in Appendix A.

[5] Episodes are defined following Berman et al. (2024): a sequence of actions bookended by at least 48

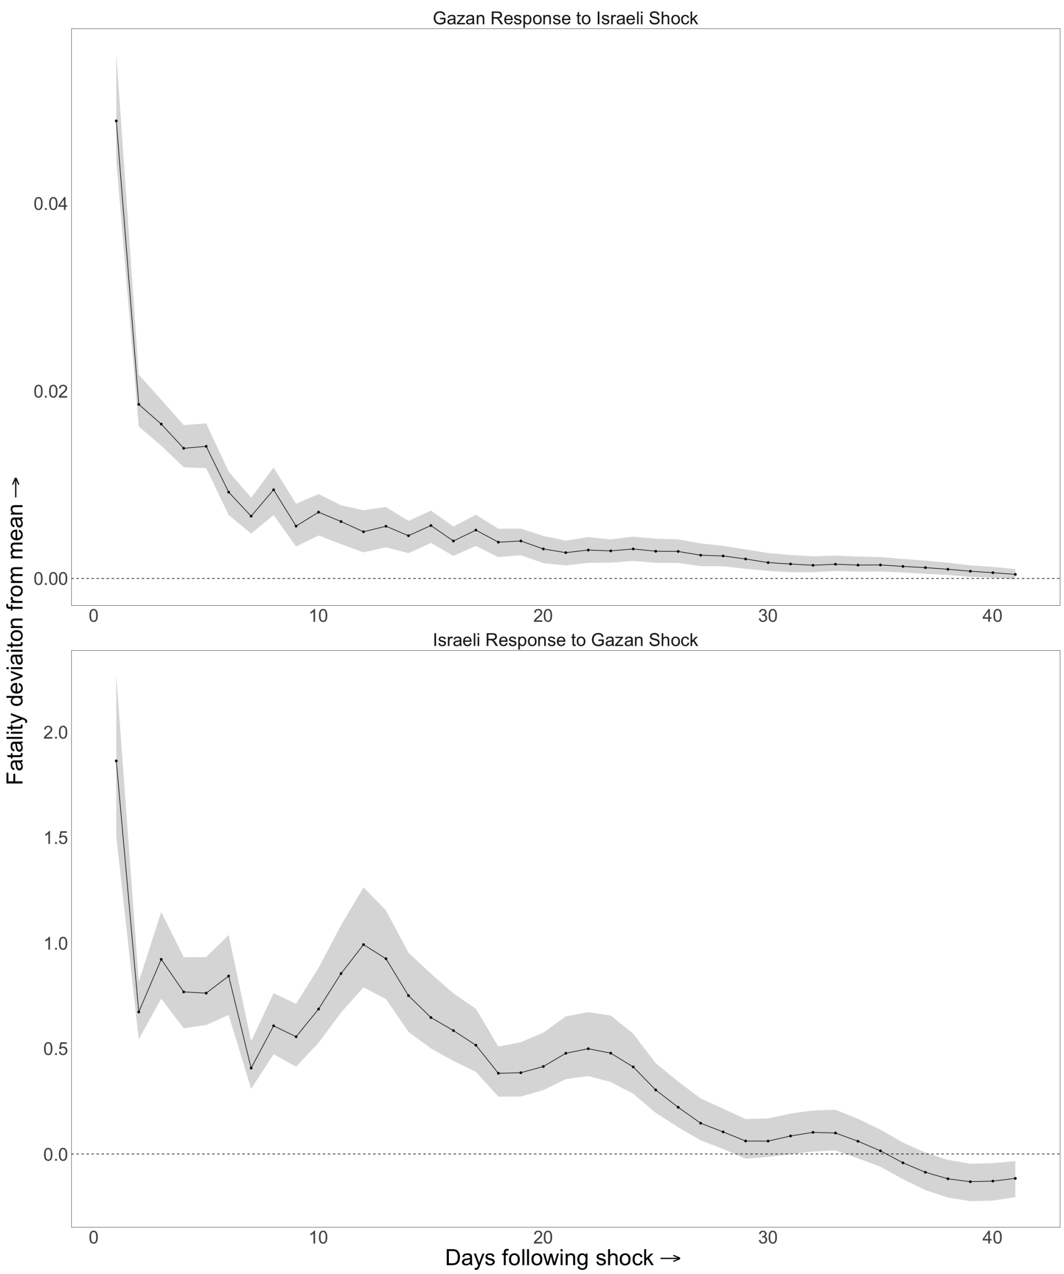

Figure 2: Daily-Level Impulse Response Functions for Expected Fatalities

Note: Orthogonal impulse response functions for expected Israeli and Gazan damages based on VAR(14) estimates.

of calm occurs.

This discrepancy raises an important question: why do time-series models estimated on daily data suggest retaliation dynamics that last far longer than the violent episodes observed in the data? In the remainder of the paper, we argue that this apparent contradiction arises from temporal aggregation. When strategic actions occur sequentially at irregular intervals but are analyzed using data aggregated into fixed calendar periods, standard time-series models may produce impulse responses that do not correspond to the underlying strategic behavior. The following sections develop a theoretical framework that formalizes this mechanism and demonstrate how temporal aggregation can distort empirical estimates of retaliation dynamics.

3. Analytical Framework: Action Time and Calendar Time

Empirical studies of strategic interaction often observe outcomes in fixed calendar intervals such as days, weeks, or months. The strategic process itself, however, typically unfolds in a different temporal structure: actors move sequentially, and both the magnitude and timing of each response are chosen endogenously. To clarify the distinction, this section introduces a simple framework that separates action time, in which strategic moves occur, from calendar time, in which researchers often observe the data. The framework provides the basis for the aggregation results developed in the next section.

3.1 Strategic Interaction in Action Time

Consider two actors, A and B, engaged in repeated sequential interaction. Let $i = 1, \ldots, I$ index actions rather than calendar periods. We refer to this representation as action time: the unit of observation is a strategic move, not a fixed temporal interval.[6] Each action inflicts damage on the opposing actor. Let d_i^j denote the damage generated by action i on actor j,

hours of calm.

[6]This general framework follows a marked process, first popularized within financial econometrics by Engle and Russell (1998) and Engle (2000).

and let z_i indicate which actor moved:

$$z_i = \begin{cases} 1 & \text{if actor } A \text{ moves at action } i \\ 0 & \text{if actor } B \text{ moves at action } i \end{cases}$$

The central behavioral assumption is that each actor's action depends on the damage inflicted by the opponent's previous action.[7]. When actor A moves, damage depends on the previous action by B, that is $d_i^B = R^A(d_{i-1}^A)$; when actor B moves, damage depends on the previous action by A, $d_i^A = R^B(d_{i-1}^B)$. We model these responses using linear reaction functions:

$$R^A(d_i^A) = d_i^B = \alpha_0^A + \alpha_1^A d_{i-1}^A + \epsilon_i \quad \text{if } z_i = 1 \tag{1}$$

$$R^B(d_i^B) = d_i^A = \alpha_0^B + \alpha_1^B d_{i-1}^B + \epsilon_i \quad \text{if } z_i = 0 \tag{2}$$

where ϵ_i is a mean zero disturbance term with variance σ_i^2.

Equivalently, the two equations can be combined into a single specification:

$$R(d_i|d_{i-1}, z_i) = d_i = \alpha_0^B + (\alpha_0^A - \alpha_0^B)z_i + \alpha_1^B d_{i-1} + (\alpha_1^A - \alpha_1^B)d_{i-1}z_i + \epsilon_i \tag{3}$$

This representation makes clear that strategic behavior in action time is governed by a small set of underlying parameters: the intercept and slope of each actor's reaction function. The slope parameters, α_1^A and α_1^B, capture how strongly each actor responds to the opponent's previous action. Positive values imply escalation; negative values imply de-escalation.

In matrix form, let D be the $I \times 1$ vector of action-level damages, Z the $I \times 1$ vector of mover indicators, L_I the lag matrix in action time, and $L_I D$ the lagged damage vector in

[7]While we focus our attention on an AR(1) response curve, the framework can accommodate more lags, higher order polynomials, and moving averages. This specification translates to a Markov one process mixed strategy employed in many sequential games, a common modeling choice outside conflict analysis (Noel, 2007). Berman et al. (2024) also show that this specification empirically captures this conflict's dynamics.

action time. Then equation (3) can be written as

$$R(D|Z) = X_I\alpha + \epsilon, \tag{4}$$

where $X_I = [\mathbf{1}, Z, L_ID, (L_ID) \odot Z]$, $\alpha = (\alpha_0^B,\ \alpha_0^A - \alpha_0^B,\ \alpha_1^B,\ \alpha_1^A - \alpha_1^B)^\top$, and $\odot$ denotes element-by-element multiplication. We assume $E[\epsilon|Z, L_ID] = 0$.

Equation (4) is the action-time object the researcher would ideally estimate if the sequence of strategic moves were directly observed. The specification allows us to clearly highlight distortion caused from studying conflict at predetermined intervals. This setup follows the logic of Berman et al. (2024), but here the framework is used to study a different question: not how actors respond in action time per se, but how these underlying responses are distorted when the data are observed only after temporal aggregation.

3.2 The Timing of Actions

Strategic interaction involves not only how actors respond, but also when they respond. Timing, therefore, influences the decision process as well. Wait too long, and the antagonist may misinterpret the protagonist's response as the beginning of a new conflict. Respond too predictably, and the antagonist can parry the protagonist's response. Let $w_i \geq 0$ denote the elapsed time between the end of action $i-1$ and the occurrence of action i. These waiting times may vary across actions and actors.

We allow waiting times to follow a conditional distribution $f_w(w_i|d_i, z_i)$ which may depend on both the size of the response and the identity of the actor. This formulation captures the idea that timing is itself strategic. Actors may respond immediately, delay their response, or remain inactive for extended periods. The sequence $\{d_i, w_i, z_i\}_{i=1}^I$ characterizes the strategic process in action time.

3.3 From Action Time to Calendar Time

Researchers often do not observe the strategic process in action time. Instead, they observe damage aggregated into exogenously defined calendar intervals such as days. To formalize this distinction, let $t = 1, \ldots, T$ index calendar periods and let D_t^A and D_t^B denote the total damage inflicted by actors A and B, respectively, during calendar period t.

Mapping action-time behavior into calendar time requires two ingredients. First, an action must be assigned to one or more calendar periods according to when it occurs. Waiting times determine the calendar placement of each action. Second, some actions may be effectively instantaneous, while others may span multiple calendar periods. To allow for both cases, let $p_{t,i}$ denote the share of action i's total damage assigned to calendar period t, where $0 \leq p_{t,i} \leq 1$ and $\sum_{t=1}^{T} p_{t,i} = 1$ for every i (this approach follows Jordá 1999 and Jordá and Marcellino 2000).

Collect these shares in a $T \times I$ aggregation matrix P, where the (t, i) entry gives the fraction of action i observed in calendar period t. Using P, calendar-time damage for each actor can be written compactly as

$$D^A = P \operatorname{diag}(Z)\, D \quad \text{and} \quad D^B = P \operatorname{diag}(1 - Z)\, D, \tag{5}$$

where $\operatorname{diag}(\cdot)$ is a diagonal matrix with vector $(\cdot)$ along the diagonal.

These equations define the relationship between the strategic process and the data observed by the researcher. The matrix P summarizes three distinct sources of temporal mismatch. First, multiple actions may occur within a single calendar period, so several strategic moves are combined into one observed outcome. Second, waiting periods may generate empty calendar intervals, so the temporal distance between observed outcomes need not correspond to the strategic distance between actions. Third, a single action may span multiple calendar periods, so one strategic move may be split across several observations.

These three features imply that calendar-time observations need not preserve the sequence of strategic responses present in action time. The same underlying reaction process may therefore produce very different empirical patterns once the data are aggregated into fixed intervals.

3.4 Implications

The distinction between action time and calendar time is central to the paper's argument. In action time, the strategic process is governed by a parsimonious reaction function in which each actor responds to the previous move. In calendar time, however, the researcher observes a transformed version of that process, shaped by the timing of actions and the aggregation rule P.

The key implication is that standard time-series models estimated on calendar-time data need not recover the underlying action-time response parameters. Instead, the coefficients of those models generally reflect both the strategic reaction functions and the aggregation process that maps actions into observed periods. The next section derives this result formally and shows how temporal aggregation can generate attenuation, amplification, and even sign reversals in estimated response dynamics.

4. Temporal Aggregation and Misestimation in Calendar Time

Section 3 defined the strategic process in action time and the aggregation rule that maps actions into calendar-time observations. We now show the main consequence for empirical analysis: when researchers estimate standard time-series models using calendar-time data, the resulting coefficients generally do not recover the underlying action-time reaction parameters. Instead, they reflect a transformed version of those parameters induced by the aggregation process.

4.1 A Calendar-Time VAR

Suppose the researcher observes calendar-time damage inflicted by actors A and B, denoted D_t^A and D_t^B, for $t = 1, \dots, T$. A standard empirical approach is to estimate a k-lag vector autoregression. Focusing on actor A's equation, the model is

$$D_t^A = \beta_0^A + \sum_{j=1}^{k} \beta_{A,j}^A D_{t-j}^A + \sum_{j=1}^{k} \beta_{B,j}^A D_{t-j}^B + \eta_t \tag{6}$$

and in matrix form,

$$D^A = X_T \beta^A + \eta, \tag{7}$$

where D^A is the $T \times 1$ vector of calendar-time damage inflicted by actor A, X_T is the design matrix containing an intercept and k lags of D^A and D^B, and β^A is the corresponding parameter vector.

This is the object typically estimated in empirical studies of retaliation, escalation, and deterrence using daily or weekly data. The question is whether β^A recovers the strategic response parameters defined in action time.

4.2 The Mapping from Action-Time Parameters to Calendar-Time Estimates

From Section 3.3, observed calendar-time damage inflicted by actor A is

$$D^A = P \, \mathrm{diag}(Z) \, D \tag{8}$$

and action-time damage is generated by

$$D = X_I \alpha + \varepsilon. \tag{9}$$

Substituting (9) into (8) yields

$$D^A = P\,\mathrm{diag}(Z)\,X_I\,\alpha + P\,\mathrm{diag}(Z)\,\varepsilon. \tag{10}$$

Under the usual projection formula, the population coefficient vector from the calendar-time VAR is

$$\beta^A = E[X_T' X_T]^{-1} E[X_T' D^A]. \tag{11}$$

Using (10), this becomes

$$\beta^A = E[X_T' X_T]^{-1} E[X_T' P\,\mathrm{diag}(Z)\,X_I]\,\alpha. \tag{12}$$

Define

$$\Upsilon^A = E[X_T' X_T]^{-1} E[X_T' P\,\mathrm{diag}(Z)\,X_I]. \tag{13}$$

Then

$$\beta^A = \Upsilon^A \alpha. \tag{14}$$

An analogous expression holds for actor B's equation.

Equation (14) is the central result of the paper. The coefficients estimated in calendar time are not, in general, the underlying action-time reaction parameters. Rather, they are weighted combinations of those parameters, where the weights are determined by the aggregation process and by the covariance structure of the observed calendar-time regressors.

4.3 Why Recovery Fails

The calendar-time VAR recovers the action-time parameters only under highly restrictive conditions. Recovery would require the regressors in the estimated VAR, X_T, to align exactly with the action-time regressors after aggregation.[8] In most strategic settings, that condition

[8] Formally, recovery requires $\Upsilon^A = I$, which in turn requires $E[X_T' P\,\mathrm{diag}(Z)\,X_I] = E[X_T' X_T]$. This holds only if the calendar-time design matrix perfectly spans the projection of the aggregated action-time regressors.

fails for three reasons.

First, multiple actions may occur within a single calendar period. In that case, one observation in calendar time combines several distinct strategic moves, so the observed lag structure no longer corresponds to the one-action-at-a-time response process defined in Section 3. Second, waiting periods generate irregular gaps between actions. A one-period lag in calendar time need not correspond to one strategic lag in action time. Depending on the timing of actions, the relevant response may appear at different calendar lags across observations. Third, individual actions may span multiple calendar periods. When a single action unfolds across days, weeks, or months, one strategic move is split across several observations, creating serial dependence that is not present in the underlying action-time reaction function.

These features imply that the relationship between β^A and α depends not only on strategic behavior, but also on how actions are distributed across calendar periods. The estimated VAR coefficients, therefore, reflect both the underlying reaction functions and the mechanics of temporal aggregation.

The next section illustrates potential distortions using simulated data and shows that even simple action-time processes can generate misleading calendar-time estimates.

5. Simulation Evidence on Temporal Aggregation Distortions

This section quantifies how temporal aggregation distorts standard time-series estimates of strategic interaction. We proceed in two steps. First, we present a simple illustration that isolates the mechanisms generating bias. We then implement Monte Carlo simulations to evaluate how these distortions emerge under more general conditions.

5.1 A Simple Illustration of Aggregation Bias

We begin with a stylized example that highlights how temporal aggregation alters estimated response dynamics. Suppose that actions occur in continuous time and follow a simple first-order process in action time. Both actors share the same action-time response function, responding only to the immediately preceding action, and waiting times are exponentially distributed:[9]

$$w_i \mid d_i, z_i \sim \exp(1) \tag{15}$$

$$d_i \mid z_i, d_{i-1} \sim N(\alpha_0 + \alpha_1 d_{i-1},\ \sigma^2) \tag{16}$$

In action time, this process is first-order Markov: each action depends only on the immediately preceding action.

The econometrician does not observe actions directly, but instead observes damage aggregated into daily intervals, as defined in Section 3. Using these daily data, the researcher estimates a VAR with five lags, ensuring that the true action-time response is always captured within the lag structure. Figure 3 reports the resulting VAR estimates under three parameterizations of the action-time response function.

The figure illustrates three distinct distortions induced by temporal aggregation. In Panel A, the first three lagged estimates for side B are statistically larger than the action-level slope coefficient $\alpha_1 = 0.1$. We refer to this as *amplification*: estimated coefficients may exceed the underlying action-time response parameters. This is caused by multiple actions being bundled into the same calendar period or when the same strategic effect is distributed across several lagged regressors in a way that inflates individual estimates. Panel B shows the opposite effect whereby VAR coefficients are all smaller than the true reaction curve slope (again $\alpha_1 = 0.1$). We refer to this as *attenuation*. Here, aggregation disperses the effect of a single action across multiple calendar periods, weakening the estimated response. Panel C shows an example where the time-space slope coefficients for side B are all positive. This would imply that side A escalates a conflict. However, the true reaction curve slope coefficient is negative, $\alpha_1 = -0.1$, which is a de-escalatory position. We refer to this as *sign reversal*. Because

[9] The functional forms are common in the ultra-high frequency literature. See Engle and Russell (1998) for more details.

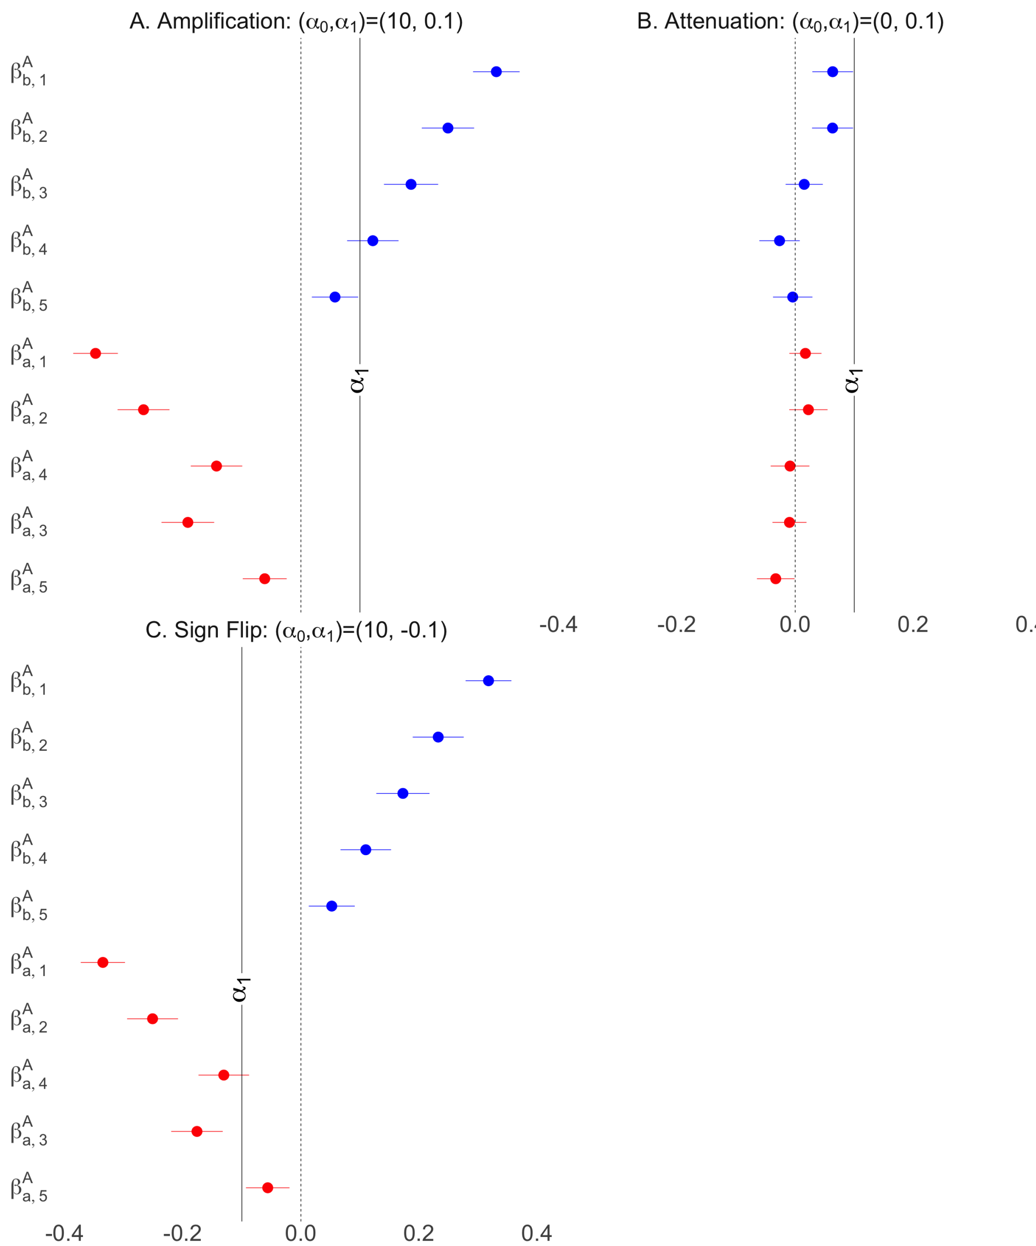

Figure 3: Actor A VAR Slope Estimates and 95% CI Under Different Action-Time Response Curves

each calendar-time coefficient is a linear combination of several action-time parameters, the estimated sign of a response may differ from the sign of the true underlying reaction function. A de-escalatory action-time process may therefore appear escalatory in aggregated data, or vice versa.

These distortions are substantively important. In conflict settings, for example, a researcher may infer prolonged retaliation, delayed escalation, or persistent own-lag dependence, even in a simple environment in which actors respond only to the most recent move. The reason is that a one-action lag does not map cleanly into a one-period lag in calendar time. If actions occur faster than the observation interval, several moves may be compressed into one period. If actions occur more slowly, the effect of one move may appear across several lagged periods. And if actions span multiple periods, one move can mechanically induce serial correlation across calendar-time observations.

5.2 Simulation Setup

We now turn to Monte Carlo simulations that incorporate asymmetry between actors and more realistic parameter values. The goal is to evaluate how temporal aggregation affects standard VAR estimates across a range of waiting times.

Our data generation process consists of simulated action-time data using the framework developed in Section 3. Damage evolves according to

$$d_i \mid z_i, d_{i-1} \sim N\Big(\alpha_0^B + \alpha_1^B d_{i-1} + (\alpha_0^A - \alpha_0^B) z_{i-1} + (\alpha_1^A - \alpha_1^B) z_i d_{i-1},\ \sigma^2\Big)$$

and waiting time follows an exponential distribution

$$w_i \mid d_i, z_i \sim \exp\Big(\frac{1}{\lambda}\Big).$$

We assume that actions are instantaneous, so each action is assigned to a single calendar period. This allows us to isolate the role of waiting times in generating aggregation distortions.

We simulate 15,000 actions per run using the parameter values

$$\left\{\alpha_0^B,\ \alpha_1^B,\ \alpha_0^A,\ \alpha_1^A,\ \sigma^2\right\} = \{0.0056,\ 0.1354,\ 0.0078,\ -0.7404,\ 0.1\}\,.$$

These values generate asymmetric response behavior across actors while maintaining a simple linear structure.

Simulated action-time data are converted into daily observations using the aggregation rule defined in Section 3. This produces calendar-time series for both actors that mirror the structure of empirical conflict data. We evaluate the performance of calendar-time VAR models along three dimensions:

1. **Optimal lag length**: determined using the Bayesian Information Criterion (BIC), allowing for up to 14 lags.

2. **Coefficient estimates and significance**: estimated from the VAR using the selected lag length.

3. **Impulse response functions (IRFs)**: tracing responses up to 21 days following a shock.

We vary the average waiting time parameter $\lambda \in \{0.5,\ 1,\ 2,\ \ldots,\ 7\}$, corresponding to mean response times ranging from half a day to one week. For each value of λ, we run 1,000 simulations.

5.3 VAR Lag Selection and Coefficient Estimates

We first examine how temporal aggregation affects the lag structure selected by standard model selection criteria. Table 1 reports the BIC-optimal lag length across simulations for different values of λ. As the average waiting time increases, the selected lag length increases nearly one-for-one. When actors respond more slowly, calendar-time models require more lags to capture the delayed appearance of responses across periods.

This result highlights an important point: even though the true process is first-order in action time, the estimated VAR frequently selects multiple lags. The additional lags do not reflect more complex strategic behavior, but rather the temporal dispersion created by aggregation.

Figure 4 reports the fraction of simulations in which each lagged coefficient is statistically significant at the 5% level. The left-hand side plots the coefficients when the outcome is actor A's action, while the right-hand side plots the coefficients for actor B's action. Each row showcases a different simulation specification. The top row assumes each side waits, on average half a day before responding, the middle assumes a full day, and the bottom assumes a full week.

The figure shows that when waiting times are short, only the first few lags are significant. As waiting times increase, a larger number of lagged coefficients become significant, often extending well beyond the true response horizon. Moreover, both own-lag and cross-lag coefficients become significant at longer horizons, even though the underlying process depends only on the immediately preceding action. These results demonstrate that temporal aggregation can generate spurious persistence and multi-period dynamics in estimated VAR models.

5.4 Impulse Response Functions

We next examine how temporal aggregation affects impulse response functions. Figure 5 plots these functions over 1,000 Monte Carlo simulations.

When average waiting times are short, impulse responses decay rapidly and become statistically insignificant within a few periods. In contrast, when waiting times are longer, impulse responses remain statistically significant over extended horizons, often lasting several weeks. Notably, these prolonged responses emerge even though the true action-time process features only immediate responses. Temporal aggregation spreads the effect of a single action across multiple calendar periods, creating the appearance of low-level, persistent retaliation dynamics.

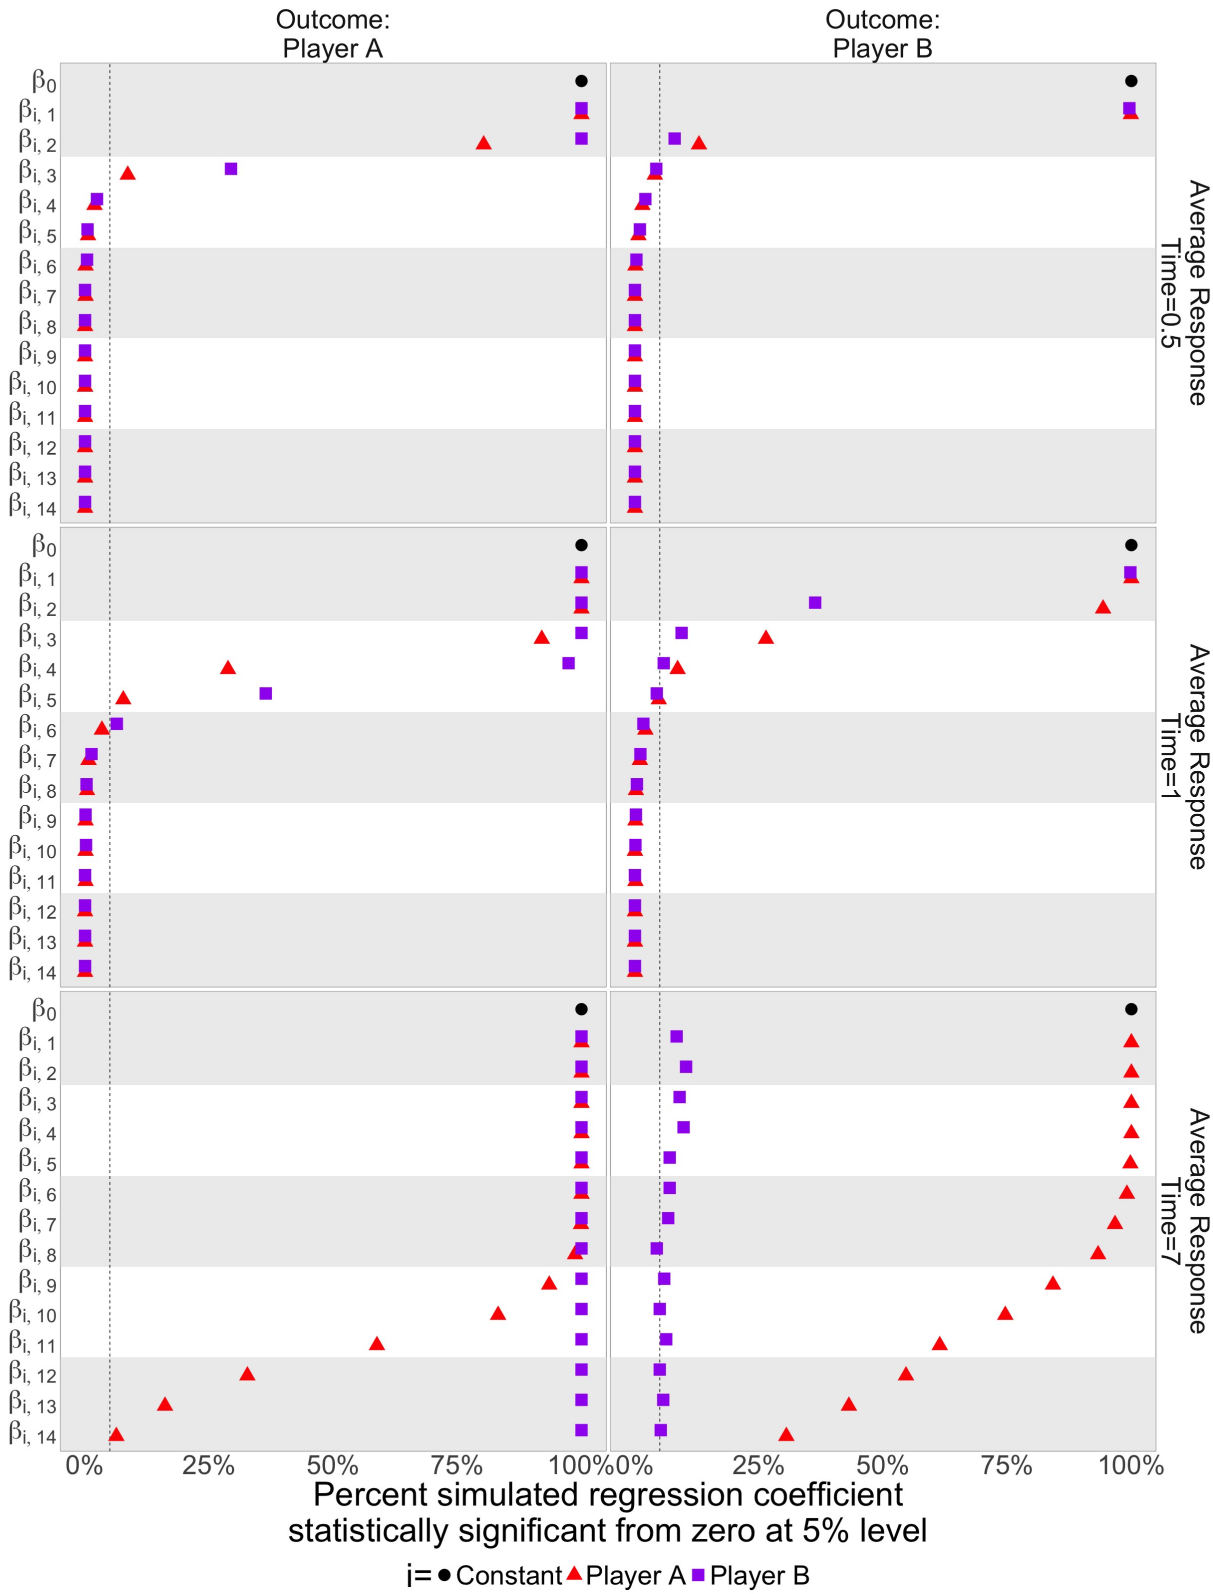

Figure 4: VAR Coefficients Over 1,000 Monte Carlo Simulations for Actors A and B

Note: Each dot represents the percent of lagged coefficients statistically significant at the 5% level.

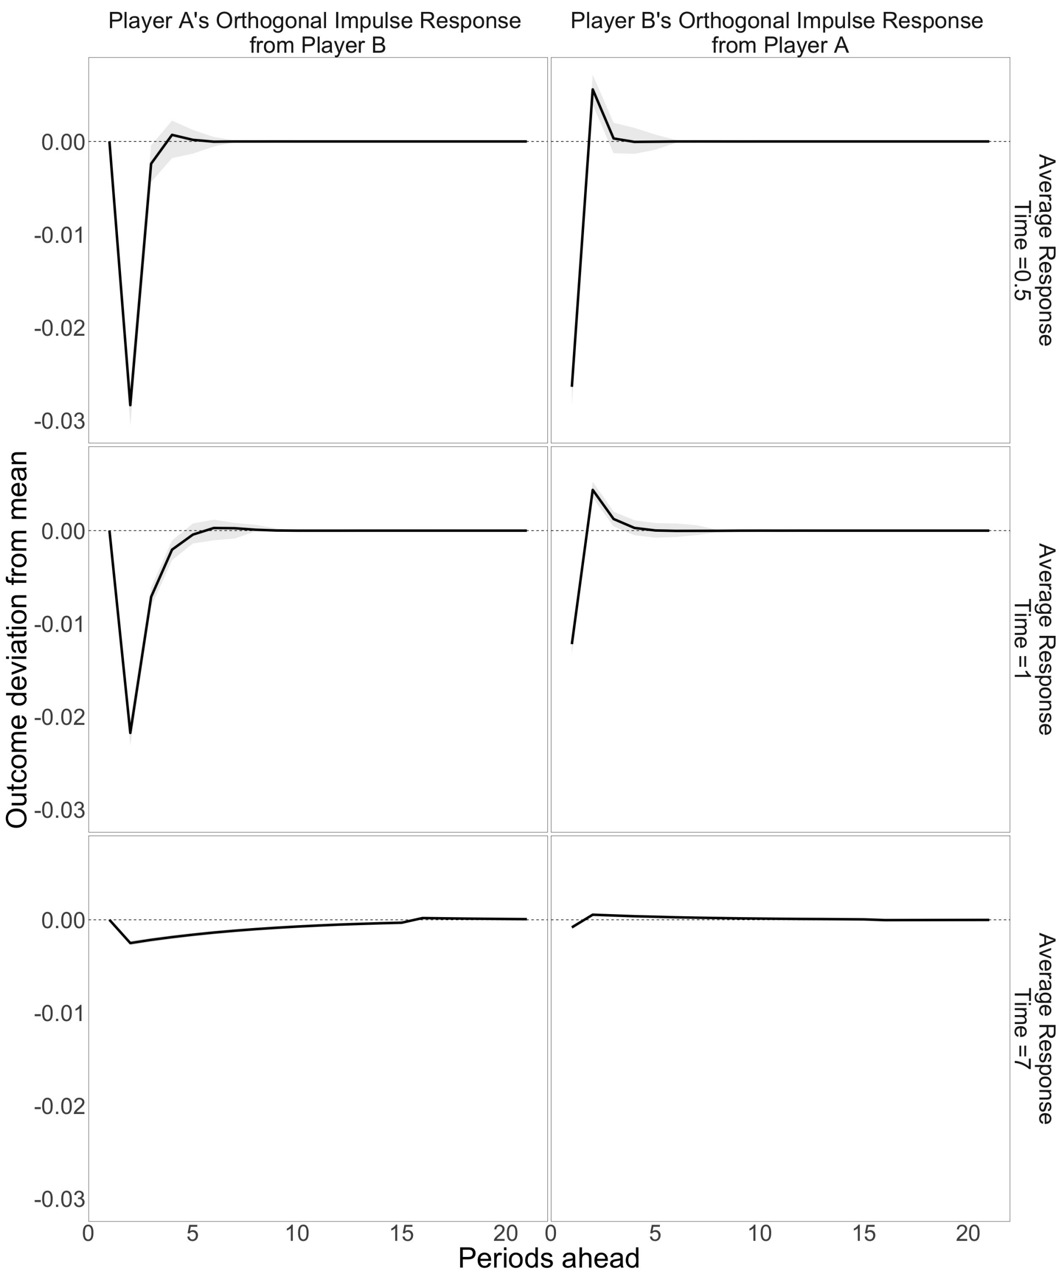

Figure 5: Simulated Orthogonal Impulse Response Functions

5.5 Summary

The simulations highlight three key findings. First, temporal aggregation causes standard lag selection criteria to identify spurious higher-order dynamics. Second, it generates statistically significant coefficients at multiple lags, even when the true process is first-order. Third, it produces impulse response functions that suggest prolonged retaliation cycles, despite the absence of such dynamics in the underlying strategic process.

Taken together, these results demonstrate that temporally aggregated data can substantially distort empirical inference about strategic interaction. The next section examines how these distortions manifest in the Israeli–Gaza conflict.

6. Empirical Application: Temporal Aggregation in the Israeli-Gaza Conflict

The preceding sections showed that temporal aggregation can distort inference about strategic interaction even when the underlying process is simple. We now examine the empirical relevance of this problem in the Israeli–Gaza conflict. The goal is to determine how much of the prolonged impulse responses in Section 2 can be explained by temporal aggregation alone.

To do so, we proceed in three steps. First, we reconstruct the conflict in action time using five-minute attack reports. Second, we estimate action-level response functions and recover the empirical aggregation process that maps actions into daily observations. Third, we simulate daily data using these empirically estimated objects and compare the resulting VAR dynamics across different forms of aggregation. This exercise allows us to isolate whether the long response patterns observed in daily data arise from strategic behavior itself or from the way actions are recorded in calendar time.

6.1 Reconstructing Action-Time Data

We begin with the five-minute attack reports described in Section 2 and convert them into a sequence of actions using a simplified version of the coding procedure in Berman et al. (2024). Their approach transforms the UN data into action-level observations by grouping attacks according to timing, munitions, and location. For our purposes, we adopt a simplified version designed to recover the sequence of strategic moves while preserving the timing structure most relevant to aggregation.

Attacks are first ordered chronologically, abstracting from location. An action is then defined as a sequence of attacks by one side that satisfies two conditions: first, the sequence is not interrupted by an attack from the opposing side; and second, no more than 48 hours elapse between consecutive attacks within the sequence. These coding rules reflect the idea that both sides may carry out coordinated offensive activity over time and across munitions, and that each side is capable of responding to the other within roughly two days.

To capture strategic inaction, we also add zero-damage actions. Specifically, for every 48-hour period of calm, we introduce two zero-damage actions, one for each side. These observations allow the absence of response to enter the data as a strategic choice rather than disappearing from the sequence entirely.

The resulting dataset has actions, rather than days, as the unit of observation. For each action, we observe the actor initiating the action, the expected damage generated by that action, the set of calendar days over which the action occurs, and the share of total damage allocated to each day.

Table 2 summarizes the resulting action-level data. Because the interaction is coded sequentially, the sample contains 2,952 Israeli actions and 2,953 Gazan actions, for a total of 5,905 actions. Roughly one-third of actions are zero-damage actions. Both sides respond, on average, in less than one day, and fewer than ten percent of responses occur more than a day later. Israeli actions are both longer and more damaging on average: they span more attacks per action and generate substantially more serious expected damage than Gazan actions. The average Israeli action caused 0.21 units of expected damage (measured as fatalities $+ 0.2 \times$ casualties), while the Gazan actions cause on average 0.03 units of expected damage.

VAR analyses compress a much richer strategic process. Actions occur at irregular intervals, many are short-lived, and some unfold across multiple days. The key question is whether these timing features are sufficient to generate the prolonged retaliation dynamics observed in calendar-time estimates.

6.2 Estimating Action-Time Response Functions and Aggregation

We next estimate the action-time response function using the framework introduced in Section 3. Specifically, we estimate

$$R(d_i|d_{i-1}, z_i) = d_i = \alpha_0^G + (\alpha_0^I - \alpha_0^G)z_i + \alpha_1^G d_{i-1} + (\alpha_1^I - \alpha_1^G)d_{i-1}z_i + \epsilon_i \tag{17}$$

where $z_i = 1$ if Israel moves at action i, and $z_i = 0$ otherwise. To keep the empirical exercise tightly connected to the theoretical framework, we restrict attention to a first-order linear response process. This is not intended as a complete behavioral model of the conflict, but rather as a parsimonious approximation that allows us to isolate the consequences of temporal aggregation.

Table 3 reports the estimated action-time response functions. The intercept is positive for both sides, indicating a baseline level of offensive activity even when the previous action generated little or no damage. The slope estimates are asymmetric. Israeli actions respond positively and significantly to prior Gazan damage, whereas the estimated Gazan slope is small and statistically indistinguishable from zero. In other words, at the action level, we find evidence that Israeli actions respond systematically to the previous Gazan action, while the corresponding linear relationship is much weaker for Gaza.

We then recover the empirical aggregation process directly from the data. For each action, we observe when it begins, how long it lasts, and how its damage is distributed across days. This allows us to construct the aggregation matrix P, where each element p_{ti} denotes the share of action i's damage assigned to day t. For example, if an action spans three days and generates expected damage of 0.1, 0.2, and 0.3 on those days, then the corresponding shares are $\{1/6, 2/6, 3/6\}$.

Combining the estimated reaction function with the empirical aggregation matrix yields calendar-time damage series of the form

$$D^I = P\,\mathrm{diag}(Z)\,D \quad \text{and} \quad D^G = P\,(I - \mathrm{diag}(Z))\,D, \tag{18}$$

where D^I denotes daily damage inflicted by Israel and D^G denotes daily damage inflicted by Gaza.

This decomposition is central to the empirical application. The estimated response function captures behavior in action time, while the matrix P captures how those actions are transformed into the daily observations used in standard conflict analyses.

6.3 Simulation Design

We now use the estimated action-time response function and the empirical aggregation process to simulate daily conflict data. The aim is to ask a counterfactual question: if actors behaved according to the estimated action-time response functions, how much persistence would appear in daily VAR estimates purely because of temporal aggregation?

For each simulation, we generate action-level damage using the estimated coefficients from Table 3 and draw disturbances from the empirical residual distribution. We then draw aggregation patterns from the observed data and map the simulated actions into daily observations using the corresponding P matrix. Finally, we estimate a daily VAR with up to 14 lags, select the lag length using the BIC, and compute the corresponding impulse response functions. We repeat this procedure 1,000 times for each of three aggregation environments.

First, we simulate using all observed actions. This reproduces the full empirical timing structure of the conflict, including both within-day actions and multi-day actions. Second, we simulate using only actions completed within a single day. This isolates the role of strategic waiting and within-day sequencing while removing the possibility that one action is mechanically spread across multiple daily observations. Third, we simulate using only multi-day actions. This isolates the role of actions whose damage is distributed across calendar periods and therefore have the greatest potential to induce serial dependence in daily data.

This design allows us to distinguish between two potential sources of distortion: variation in the timing of responses across actions, and the splitting of individual actions across multiple days.

6.4 Results

Figure 6 reports the fraction of simulations in which each lagged VAR coefficient is statistically significant at the five percent level. The top row uses all observed actions, the middle row uses only same-day actions, and the bottom row uses only multi-day actions. The green squares represent coefficients for lagged damage to Israel (Gazan actions) while the blue triangles represent the lagged damage to Gaza (Israeli actions).[10]

Using the full set of actions, the simulated daily VAR produces substantial persistence. Lagged Gazan damage is frequently significant in both the Israeli and Gazan equations, and significant effects often extend several days beyond the true one-action response horizon. This reproduces the main qualitative feature of the daily estimates in Section 2: calendar-time analysis suggests a longer dynamic process than the action-level model would imply.

The source of this persistence becomes clear once the simulations are separated by aggregation type. When the simulations use only actions completed within a single day, significant lagged effects diminish sharply. The resulting dynamics are much shorter-lived and concentrated in the earliest lags. By contrast, when the simulations use only multi-day

[10]Appendix Figure A8 reports the magnitude of simulated impulse responses.

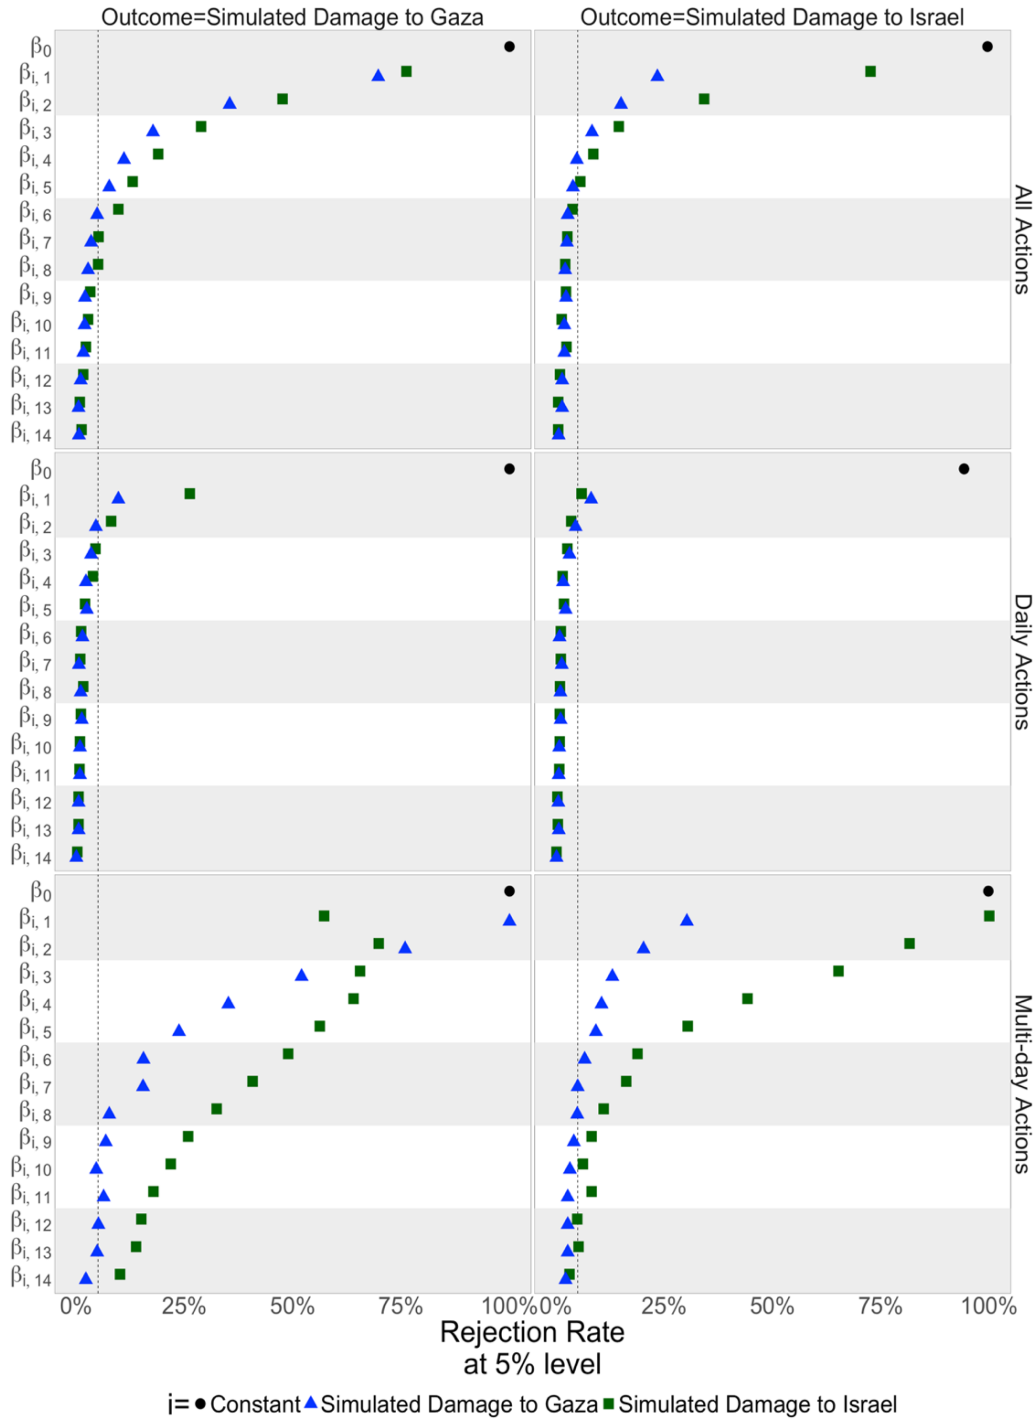

Figure 6: Share of 1,000 Monte Carlo simulations in which each lagged VAR coefficient is statistically significant at the 5% level. Each row represents a different lag with constants shown at top. Left column: coefficients in the Gazan damage equation. Right column: coefficients in the Israeli damage equation. Blue triangles: coefficients on lagged Gazan damage. Green squares: coefficients on lagged Israeli damage. The three row-panels differ by which actions are used to generate the simulated daily data: all actions (top), only within-day actions (middle), only multi-day actions. (bottom)

actions, the persistence becomes even stronger than in the full sample. Significant lagged coefficients remain common over longer horizons in both equations.

These results indicate that, in this setting, the main source of distortion is not simply irregular waiting times between actions. Rather, it is the fact that many actions are distributed across multiple calendar days. Once a single strategic move is split across several daily observations, standard VAR models attribute that temporal pattern to lagged strategic response rather than to the aggregation process itself.

Figure 7 reports the fraction of simulations in which the orthogonal impulse response is statistically significant at each horizon. When all actions are included, the simulated Israeli response (left-hand side) remains significant over much of the 25-day horizon, closely resembling the prolonged impulse responses observed in the actual daily data. The simulated Gazan response (right-hand side) is shorter-lived, but still exhibits persistence beyond the one-action response structure estimated in action time.

Again, the persistence is driven primarily by multi-day actions. Simulations based only on within-day actions produce strong but short-lived responses that fade quickly. Simulations based on multi-day actions, by contrast, generate the most persistent impulse responses. In other words, the long retaliation patterns seen in daily conflict data can arise even when actors respond only to the previous action, provided that some actions unfold across several days.

The empirical application yields two conclusions. First, the prolonged retaliation cycles estimated from daily data in the Israeli–Gaza conflict should not be interpreted mechanically as evidence that actors respond strategically over several weeks. A substantial share of that persistence can be generated by temporal aggregation alone.

Second, not all forms of aggregation are equally important. In this setting, the primary source of distortion is multi-day actions, not merely variation in the delay between one action and the next. This distinction matters substantively. It suggests that the long impulse responses in daily VARs reflect, to a considerable extent, the way sustained military operations

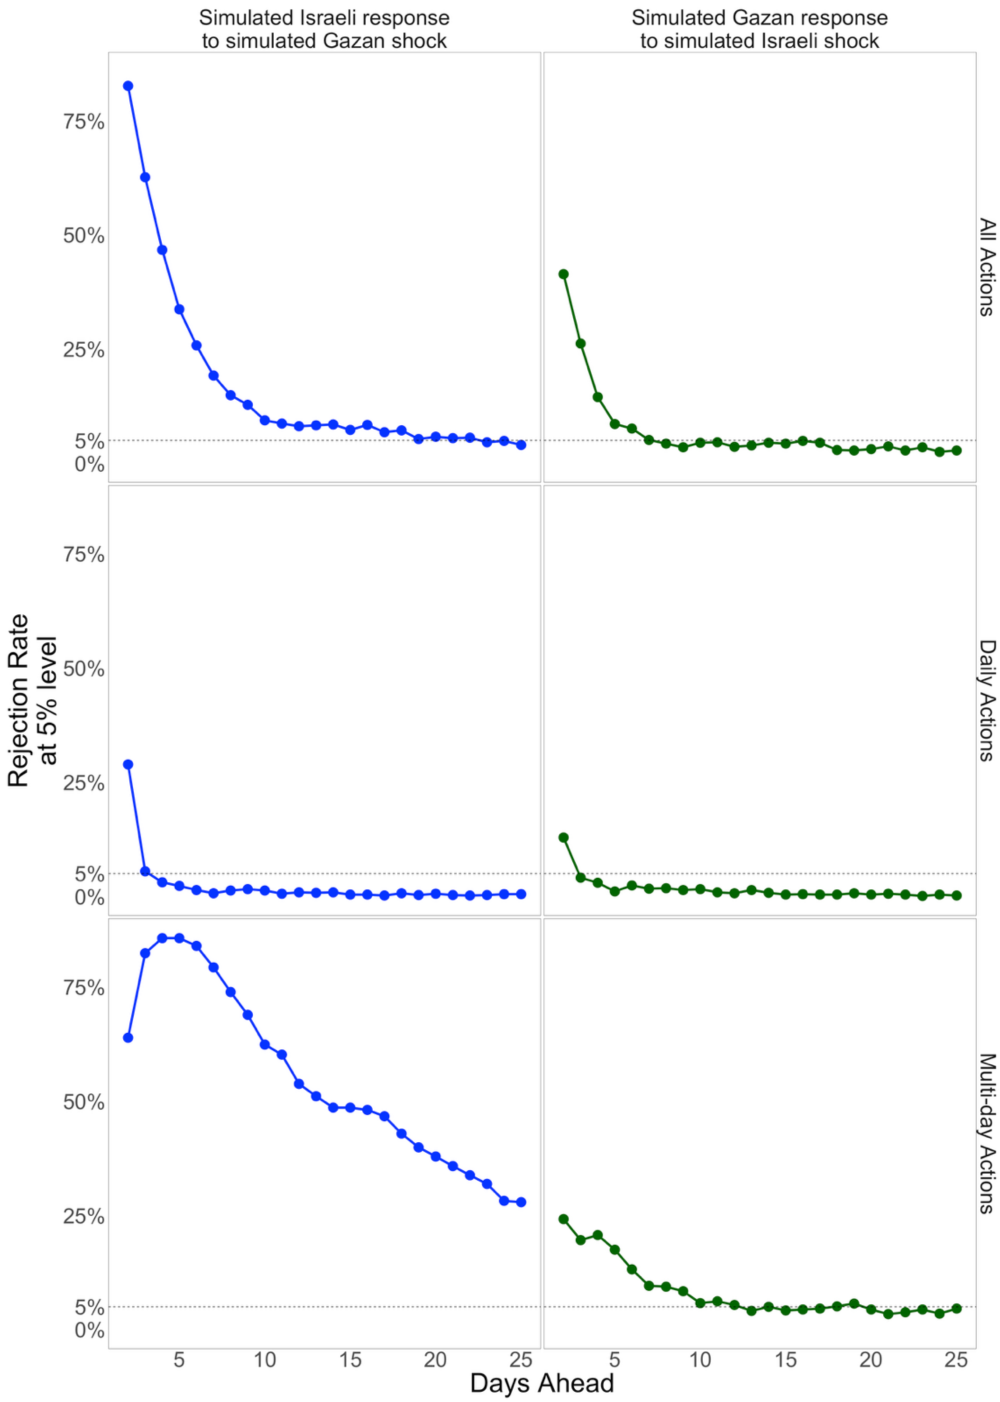

Figure 7: Share of 1,000 Monte Carlo simulations in which the orthogonal impulse response function is statistically significant at the 5% level. Left column: simulated Israeli response to a simulated Gazan shock. Right column: simulated Gazan response to a simulated Israeli shock. The three row-panels differ by which actions are used to generate the simulated daily data: all actions (top), only within-day actions (middle), and only multi-day actions (bottom). The horizontal reference line marks the 5% rejection rate expected under the null of no response.

are recorded in calendar time rather than the existence of long chains of delayed retaliation.

More broadly, the Israel–Gaza application shows that the concern raised in this paper is not only theoretical. Temporal aggregation can materially alter the conclusions researchers draw from conflict data, including the apparent duration and asymmetry of retaliation dynamics. Analyses based on daily data may therefore overstate the persistence of strategic responses unless they explicitly account for the action-time structure underlying the observed events.

7. Conclusion

Estimating how actors respond to one another is central to the study of political conflict and strategic interaction more broadly. This paper shows that when such interactions are analyzed using temporally aggregated data, standard empirical methods can produce misleading inferences about the magnitude, persistence, and even direction of responses.

The core issue is a mismatch between the timing of behavior and the timing of measurement. Strategic interactions unfold in action time, as a sequence of moves and countermoves, while empirical analyses are typically conducted in calendar time, using data aggregated into fixed intervals. We show that this mismatch induces a systematic distortion: coefficients estimated from calendar-time models are linear combinations of the underlying action-time response parameters. As a result, temporally aggregated data can generate attenuation, amplification, and sign reversals, as well as impulse responses that suggest prolonged dynamics even when actors respond only to the most recent action.

We demonstrate these distortions analytically, in simulations, and in an empirical application to the Israeli–Gaza conflict. Using high-frequency data, we reconstruct the sequence of actions and estimate response functions in action time. We then show that the prolonged retaliation patterns observed in daily data can largely be reproduced by temporal aggregation alone. In this setting, the primary source of distortion is not simply variation in response timing, but the fact that many actions unfold across multiple calendar periods, mechanically inducing persistence in aggregated data.

These findings have broader implications for empirical research on strategic interaction. First, they suggest that estimates of retaliation, escalation, and deterrence derived from temporally aggregated data should be interpreted with caution. Apparent persistence or delayed responses may reflect the aggregation process rather than the underlying strategic behavior. Second, they highlight the importance of modeling interactions at the level at which decisions are made. When high-frequency data are available, researchers should prioritize approaches that recover or approximate action-time dynamics.

More generally, the growing availability of high-resolution event data makes this issue increasingly salient. Earlier work was often constrained to data aggregated at fixed intervals. Today, researchers frequently have access to detailed temporal information, making it possible to distinguish between behavior and measurement. Our results suggest that failing to exploit this information can lead to systematic misinterpretation of strategic dynamics.

Taken together, these findings underscore a simple but important principle: the temporal unit of analysis is not innocuous. Matching empirical models to the timing of strategic behavior is essential for accurately understanding how actors respond to one another.

References

Asali, Muhammad, Aamer S. Abu-Qarn, and Michael Beenstock, "The Cycle of Violence in the Second Intifada: Causality in Nonlinear Vector Autoregressive Models," *Journal of Applied Econometrics*, 2017, *32* (6), 1197–1205.

Associated Press, "Defense minister says Israel hit South Pars petrochemical plant at Asaluyeh after Iran reported attack there," WTOP News, https://wtop.com/world/2026/04/defense-minister-says-israel-hit-south-pars-petrochemical-plant-at-asaluyeh-after-ira April 2026. Accessed: April 19, 2026.

Barberá, Pablo, Andreu Casas, Jonathan Nagler, Patrick J. Egan, Richard Bonneau, John T. Jost, and Joshua A. Tucker, "Who Leads? Who Follows? Measuring Issue Attention and Agenda Setting by Legislators and the Mass Public Using Social Media Data," *American Political Science Review*, 2019, *113* (4), 883–901.

Berman, Eli, Prabin B. Khadka, Danny Klinenberg, and Esteban Klor, "Deterrence Through Response Curves: An Empirical Analysis of the Gaza-Israel Conflict," 2024. NBER Working Paper 33273.

Box-Steffensmeier, Janet M., David Darmofal, and Christian A. Farrell, "The Aggregate Dynamics of Campaigns," *Journal of Politics*, 2009, *71* (1), 309–323.

Brewer, K. R. W., "Some Consequences of Temporal Aggregation and Systematic Sampling for ARMA and ARMAX Models," *Journal of Econometrics*, 1973, *1* (2), 133–154.

Enders, Walter and Todd Sandler, "The Effectiveness of Antiterrorism Policies: A Vector-Autoregression-Intervention Analysis," *American Political Science Review*, 1993, *87* (4), 829–844.

Engle, Robert F., "The Econometrics of Ultra-High-Frequency Data," *Econometrica*, 2000, *68* (1), 1–22.

_ **and Jeffrey R. Russell**, "Autoregressive Conditional Duration: A New Model for Irregularly Spaced Transaction Data," *Econometrica*, 1998, *66* (5), 1127–1162.

Freeman, John R., "Systematic Sampling, Temporal Aggregation, and the Study of Political Relationships," *Political Analysis*, 1989, *1*, 61–98.

_ **, John T. Williams, and Tse min Lin**, "Vector Autoregression and the Study of Politics," *American Journal of Political Science*, 1989, *33* (4), 842–877.

Geweke, John, "Temporal Aggregation in the Multiple Regression Model," *Econometrica*, 1978, *46* (3), 643–661.

Golan, David and Jonathan D. Rosenblatt, "Revisiting the Statistical Analysis of the Israeli–Palestinian Conflict," *Proceedings of the National Academy of Sciences*, 2011, *108* (15).

Haushofer, Johannes, Anat Biletzki, and Nancy Kanwisher, "Both Sides Retaliate in the Israeli–Palestinian Conflict," *Proceedings of the National Academy of Sciences*, 2010, *107* (42), 17927–17932.

Jaeger, David A. and M. Daniele Paserman, "The Cycle of Violence? An Empirical Analysis of Fatalities in the Palestinian-Israeli Conflict," *American Economic Review*, 2008, *98* (4), 1591–1604.

Marcellino, Massimiliano, "Some Consequences of Temporal Aggregation in Empirical Analysis," *Journal of Business & Economic Statistics*, 1999, *17* (1), 129–136.

Noel, Michael D., "Edgeworth Price Cycles: Evidence from the Toronto Retail Gasoline Market," *Journal of Industrial Economics*, 2007, *55* (1), 69–92.

Óscar Jordá, "Random-Time Aggregation in Partial Adjustment Models," *Journal of Business & Economic Statistics*, 1999, *17* (3), 382–395.

_ **and Massimiliano Giuseppe Marcellino**, "Stochastic Processes Subject to Time-Scale Transformations: An Application to High-Frequency FX Data," 2000. SSRN Working Paper 224788.

Silvestrini, Andrea and David Veredas, "Temporal Aggregation of Univariate and Multivariate Time Series Models: A Survey," *Journal of Economic Surveys*, 2008, *22* (3), 458–497.

Tiao, G. C. and W. S. Wei, "Effect of Temporal Aggregation on the Dynamic Relationship of Two Time Series Variables," *Biometrika*, 1976, *63* (3), 513–523.

Wei, William W. S., "The Effect of Temporal Aggregation on Parameter Estimation in Distributed Lag Model," *Journal of Econometrics*, 1978, *8* (2), 237–246.

Working, Holbrook, "Note on the Correlation of First Differences of Averages in a Random Chain," *Econometrica*, 1960, *28* (4), 916–918.

Zellner, Arnold and Claude Montmarquette, "A Study of Some Aspects of Temporal Aggregation Problems in Econometric Analyses," *Review of Economics and Statistics*, 1971, *53* (4), 335–342.

Table 1: Optimal BIC Over Monte Carlo Simulations with Varying Response Time

Response Time (λ)	Average BIC	Minimum BIC	Median BIC	Maximum BIC
0.5	1.0	1	1	1
1.0	1.9	1	2	2
2.0	3.1	2	3	4
3.0	4.2	3	4	6
4.0	5.2	4	5	7
5.0	6.0	5	6	8
6.0	6.8	5	7	9
7.0	7.5	6	8	10

Note: The simulation is repeated 1,000 times per response time. The BIC-optimal lag length is selected from a maximum of 14 lags. Response time λ is the mean of the exponential waiting-time distribution (in days).

Table 2: Action-Level Summary Statistics

	Gaza	Israel
Number of Actions	2,952	2,953
Proportion Zero-Damage Actions	0.36	0.30
Average Days Between Actions	0.65	0.65
Proportion Days Between Actions $>$ 1 Day	0.10	0.10
Average Length of Actions (days)	0.53	0.88
Average Number of Attacks per Action	1.30	2.35
Proportion Actions Longer Than One Day	0.11	0.16
Predicted Damage (fatalities $+0.2\times$ injuries)	0.03	0.21

Note: Actions are constructed from five-minute UN attack reports following the coding procedure described in Section 6.1. A zero-damage action represents strategic inaction coded for each 48-hour period of calm. Predicted damage is computed as fatalities $+ 0.2 \times$ injuries, using fitted values from a regression of casualties on munitions type with no intercept. The sample covers the Israeli–Gaza conflict from 2007 to 2017.

Table 3: Estimated Action-Time Response Curves

	(1)	(2)
α_0^G (Gaza intercept)	0.034^{***}	0.034^{***}
	(0.001)	(0.001)
α_1^G (Gaza slope)		0.003
		(0.002)
α_0^I (Israel intercept)	0.214^{***}	0.191^{***}
	(0.008)	(0.009)
α_1^I (Israel slope)		0.634^{***}
		(0.148)
Observations	5,905	5,901
Standard Errors	HAC	HAC

Note: Estimates of equation (17). The outcome is expected damage, computed as fatalities $+ 0.2 \times$ injuries using fitted values from a regression of casualties on munitions type. Column (1) reports intercepts only; column (2) adds slope coefficients. The first action in the dataset and the first action following each major operation have no lagged damage observation and are excluded from column (2). Standard errors are HAC-robust. $^{+}p < 0.1$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

A. Additional VAR Results

Table A1 reports the reduced-form VAR results corresponding to Figure 2 in Section 2.

Table A1: Reduced-Form VAR Findings Associated with Figure 2

	Pred. Damage to Gaza	Pred. Damage to Israel
Pred. Damage to Gaza, Lag 1	0.345^{***}	0.003^{***}
	(0.017)	(0.0005)
Pred. Damage to Israel, Lag 1	1.953^{**}	0.380^{***}
	(0.619)	(0.017)
Pred. Damage to Gaza, Lag 2	0.326^{***}	-0.002^{***}
	(0.018)	(0.0005)
Pred. Damage to Israel, Lag 2	3.787^{***}	0.187^{***}
	(0.658)	(0.018)
Pred. Damage to Gaza, Lag 3	0.107^{***}	0.001^{+}
	(0.019)	(0.001)
Pred. Damage to Israel, Lag 3	-1.504^{*}	0.073^{***}
	(0.668)	(0.018)
Pred. Damage to Gaza, Lag 4	0.052^{**}	-0.000
	(0.019)	(0.001)
Pred. Damage to Israel, Lag 4	0.505	0.089^{***}
	(0.670)	(0.018)
Pred. Damage to Gaza, Lag 5	0.092^{***}	0.002^{***}
	(0.019)	(0.001)
Pred. Damage to Israel, Lag 5	-1.258^{+}	-0.046^{*}

Continued on next page

Table A1 continued

	Pred. Damage to Gaza	Pred. Damage to Israel
	(0.672)	(0.018)
Pred. Damage to Gaza, Lag 6	-0.170^{***}	-0.001
	(0.018)	(0.001)
Pred. Damage to Israel, Lag 6	-2.005^{**}	-0.030
	(0.671)	(0.018)
Pred. Damage to Gaza, Lag 7	0.027	-0.002^{***}
	(0.019)	(0.001)
Pred. Damage to Israel, Lag 7	2.053^{**}	0.074^{***}
	(0.668)	(0.018)
Pred. Damage to Gaza, Lag 8	0.098^{***}	-0.002^{**}
	(0.019)	(0.001)
Pred. Damage to Israel, Lag 8	-3.761^{***}	-0.040^{*}
	(0.668)	(0.018)
Pred. Damage to Gaza, Lag 9	0.112^{***}	0.002^{***}
	(0.018)	(0.001)
Pred. Damage to Israel, Lag 9	0.618	0.065^{***}
	(0.670)	(0.018)
Pred. Damage to Gaza, Lag 10	0.113^{***}	0.001^{*}
	(0.018)	(0.001)
Pred. Damage to Israel, Lag 10	0.447	0.001
	(0.671)	(0.018)
Pred. Damage to Gaza, Lag 11	0.094^{***}	0.001^{*}

Continued on next page

Table A1 continued

	Pred. Damage to Gaza	Pred. Damage to Israel
	(0.019)	(0.001)
Pred. Damage to Israel, Lag 11	2.296^{***}	-0.018
	(0.669)	(0.018)
Pred. Damage to Gaza, Lag 12	-0.090^{***}	-0.002^{**}
	(0.019)	(0.001)
Pred. Damage to Israel, Lag 12	-0.974	0.009
	(0.668)	(0.018)
Pred. Damage to Gaza, Lag 13	-0.176^{***}	-0.002^{***}
	(0.017)	(0.0005)
Pred. Damage to Israel, Lag 13	3.506^{***}	-0.003
	(0.658)	(0.018)
Pred. Damage to Gaza, Lag 14	-0.097^{***}	0.000
	(0.017)	(0.0005)
Pred. Damage to Israel, Lag 14	-1.170^{+}	0.017
	(0.619)	(0.017)
Intercept	-0.005	0.004^{***}
	(0.033)	(0.001)
Observations (Days)	3,822	3,822
Average Outcome	0.54	0.02
R^2	0.776	0.497
Adj. R^2	0.774	0.493

Note: Reduced-form VAR estimates using daily expected damage. Lag length of 14 selected by BIC from a maximum of 14 lags. $^{+}p < 0.1$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Table A2 and Figure A1 show the VAR and orthogonal IRF findings using daily fatalities.

Table A2: Reduced-Form VAR Findings Using Daily Fatalities

	Israelis Killed	Gazans Killed
Israelis Killed, Lag 1	0.032^{+}	3.965^{***}
	(0.017)	(0.650)
Gazans Killed, Lag 1	0.002^{***}	0.240^{***}
	(0.0004)	(0.016)
Israelis Killed, Lag 2	0.058^{***}	3.970^{***}
	(0.017)	(0.651)
Gazans Killed, Lag 2	-0.001	0.163^{***}
	(0.0004)	(0.017)
Israelis Killed, Lag 3	-0.017	1.403^{*}
	(0.017)	(0.653)
Gazans Killed, Lag 3	0.001	0.130^{***}
	(0.0004)	(0.017)
Israelis Killed, Lag 4	0.048^{**}	0.395
	(0.017)	(0.652)
Gazans Killed, Lag 4	-0.001^{*}	0.043^{*}
	(0.0004)	(0.017)
Israelis Killed, Lag 5	-0.002	3.295^{***}
	(0.017)	(0.650)
Gazans Killed, Lag 5	-0.000	0.010
	(0.0004)	(0.017)
Israelis Killed, Lag 6	0.018	1.575^{*}

Continued on next page

Table A2 continued

	Israelis Killed	Gazans Killed
	(0.017)	(0.650)
Gazans Killed, Lag 6	-0.001	-0.062^{***}
	(0.0004)	(0.017)
Israelis Killed, Lag 7	-0.027	-0.523
	(0.017)	(0.650)
Gazans Killed, Lag 7	0.001^{*}	0.022
	(0.0004)	(0.017)
Israelis Killed, Lag 8	0.012	-2.286^{***}
	(0.017)	(0.650)
Gazans Killed, Lag 8	0.003^{***}	0.070^{***}
	(0.0004)	(0.017)
Israelis Killed, Lag 9	-0.029^{+}	-2.033^{**}
	(0.017)	(0.651)
Gazans Killed, Lag 9	-0.001^{**}	0.139^{***}
	(0.0004)	(0.017)
Israelis Killed, Lag 10	-0.043^{**}	3.323^{***}
	(0.017)	(0.650)
Gazans Killed, Lag 10	0.003^{***}	0.071^{***}
	(0.0004)	(0.017)
Israelis Killed, Lag 11	-0.032^{+}	2.023^{**}
	(0.017)	(0.652)
Gazans Killed, Lag 11	-0.000	-0.012

Continued on next page

Table A2 continued

	Israelis Killed	Gazans Killed
	(0.0004)	(0.017)
Israelis Killed, Lag 12	-0.030^{+}	-3.096^{***}
	(0.017)	(0.649)
Gazans Killed, Lag 12	-0.002^{***}	0.035^{*}
	(0.0004)	(0.017)
Israelis Killed, Lag 13	-0.012	5.405^{***}
	(0.017)	(0.650)
Gazans Killed, Lag 13	-0.001^{**}	-0.102^{***}
	(0.0004)	(0.016)
Intercept	0.007^{**}	0.054
	(0.002)	(0.090)
Observations (Days)	3,823	3,823
Average Daily Fatalities	0.009	0.842
R^2	0.064	0.442
Adj. R^2	0.058	0.438

Note: Reduced-form VAR estimates using daily fatality counts. Lag length of 13 selected by BIC from a maximum of 14 lags. $^{+}p < 0.1$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

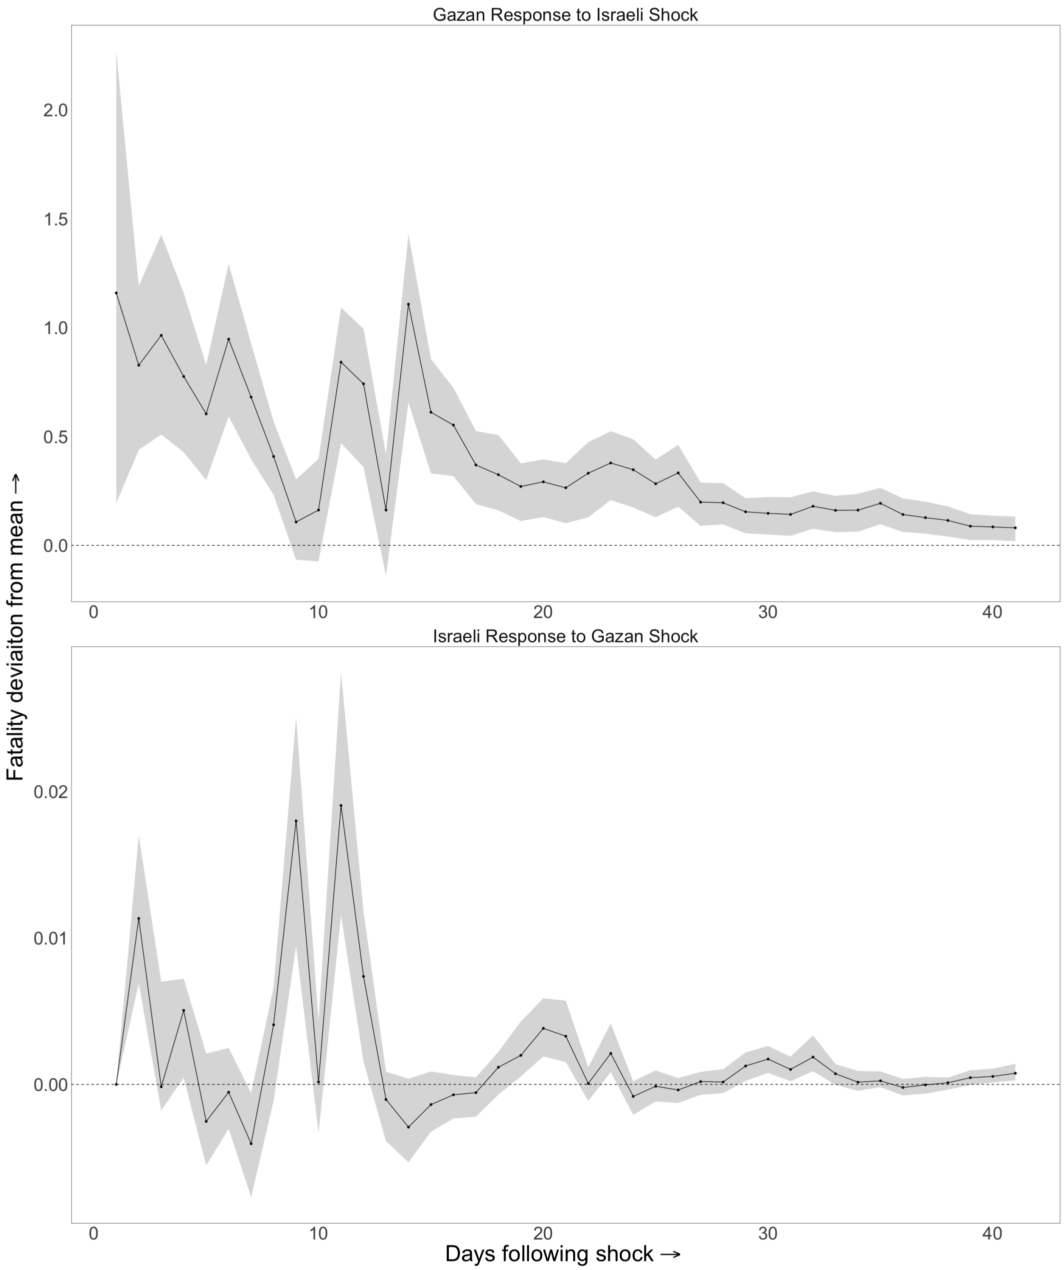

Figure A1: Orthogonal IRFs Using Daily Fatalities

Finally, Figure A2 plots the impulse response functions for expected fatalities generated using the local projection method. A similar prolonged response is seen for Gaza using this method.

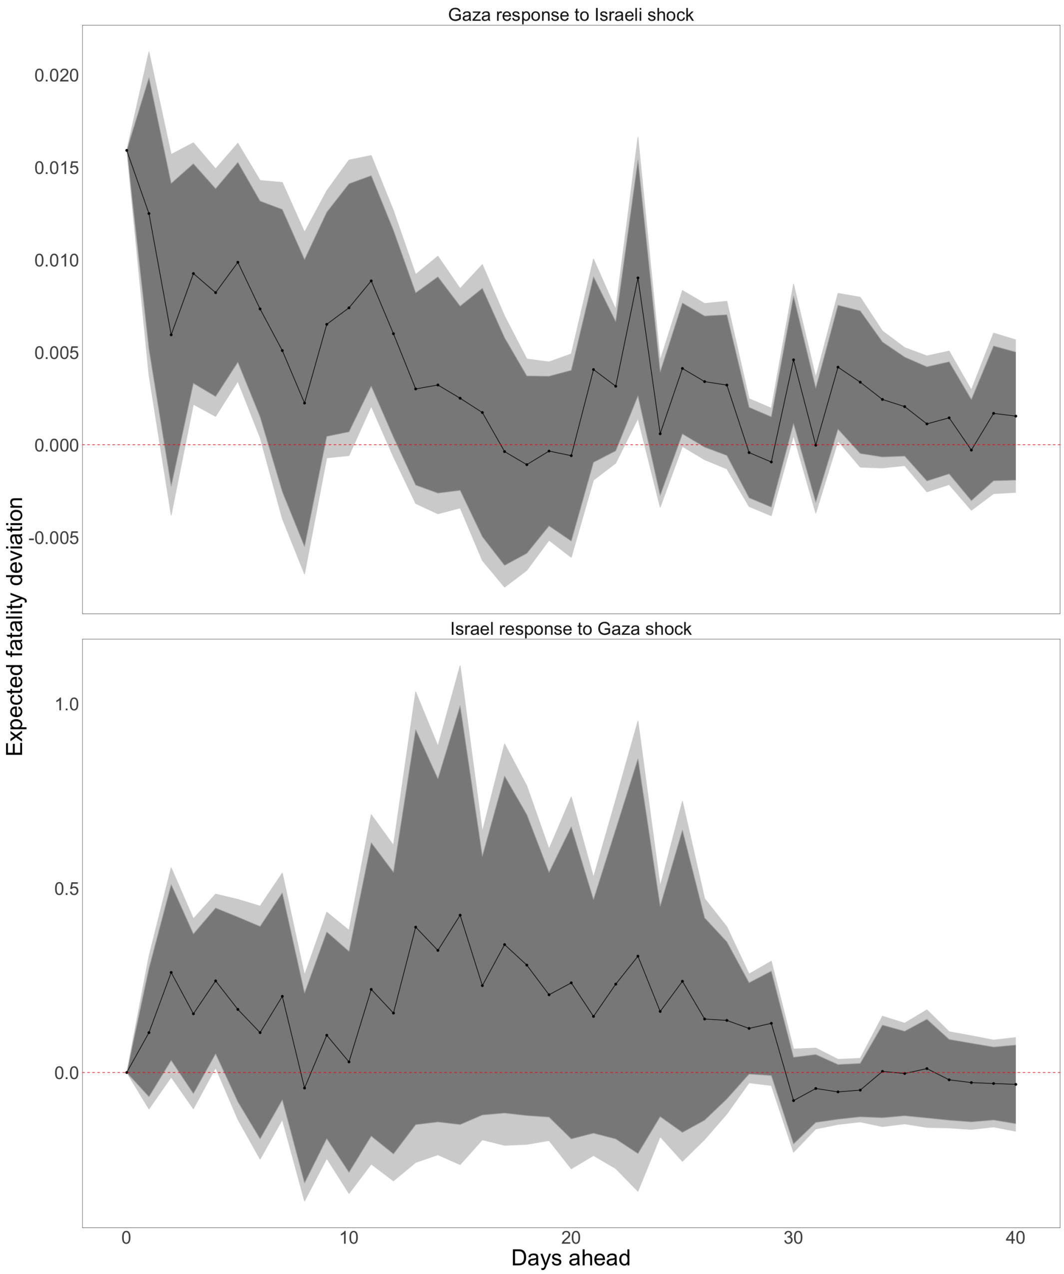

Figure A2: IRFs Generated Using Local Projection Method for Expected Fatalities

Note: The dark shaded area shows 90% confidence intervals; the light shaded area shows 95% confidence intervals.

B. An Example for Section 3.2

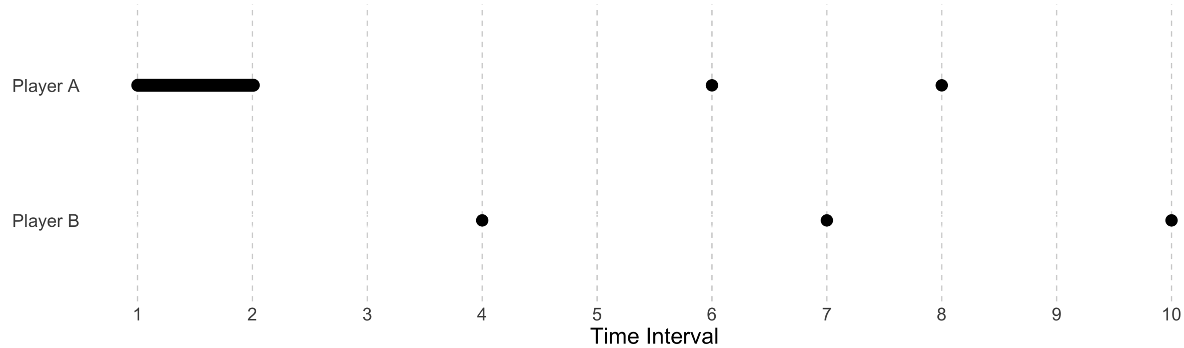

Figure A3: Graphical Example of Actions Recorded at Pre-Specified Time Intervals

Figure A3 provides a graphical illustration of six instantaneous actions over ten days when the time interval is always shorter than the response time.[11] Suppose Player A and Player B are observed over 10 time-units (e.g. days). Within each unit, a player may perform an action causing positive damage or wait to perform an action. The black circle represents an action, and empty slots represent waiting. In this example, Player A first inflicts positive damage to Player B over two days ($t = 1$ and $t = 2$). Player B then waits two time periods and attacks Player A ($t = 4$). Player A then attacks inflicting positive damage at $t = 6$, leading to a retaliation from Player B followed by further attacks from Player A.

The example highlights that studying the sequential game in arbitrary time intervals can cause the lag structure to differ across actions. Even though Player B's first and second actions (e.g. $t = 4$ and $t = 7$) are reacting to Player A's previous action, Player B is reacting to the third lagged time in their first action ($t = 1$) and first lagged time period in the second action ($t = 6$). We refer to the mapping between recording actions at the action-level and time-interval level as the *data aggregation process*.

We present the example mappings in matrix form of Player A:

[11]Mathematically, $w_i \geq 1 \ \forall i$.

$$P\,\mathrm{diag}(Z) = \begin{bmatrix} \alpha & 0 & 0 & 0 & 0 & 0 \\ 1-\alpha & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \mathbf{1} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \mathbf{1} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

where the ones are bolded for easier reading. The waiting periods lead to many rows of only zeroes. This is one of three data distortions from aggregating to arbitrary time intervals. The second is many actions occurring in one day, represented by rows having multiple actions.

Finally, an action may be spread over multiple days. This is represented by a column having two nonzero entries that add to unity. For example, side A's i-th action (like a shelling campaign) may begin on day t and conclude on $t+1$. If they dropped a third of their bombs on t and the rest on $t+1$, then the $[t, i]$ entry for PZ is $\frac{1}{3}$ and $\frac{2}{3}$ for $[t+1, i]$.

C. Additional Parameter Decomposition for Section 4.2

Equation (12) shows that each β is a linear combination of α, where each scalar term is unbounded. Formally and without loss of generality,

$$\beta^A_{b,j} = \sum_{j=1}^{4} \upsilon_{(1+k+j),p}\,\alpha_p.$$

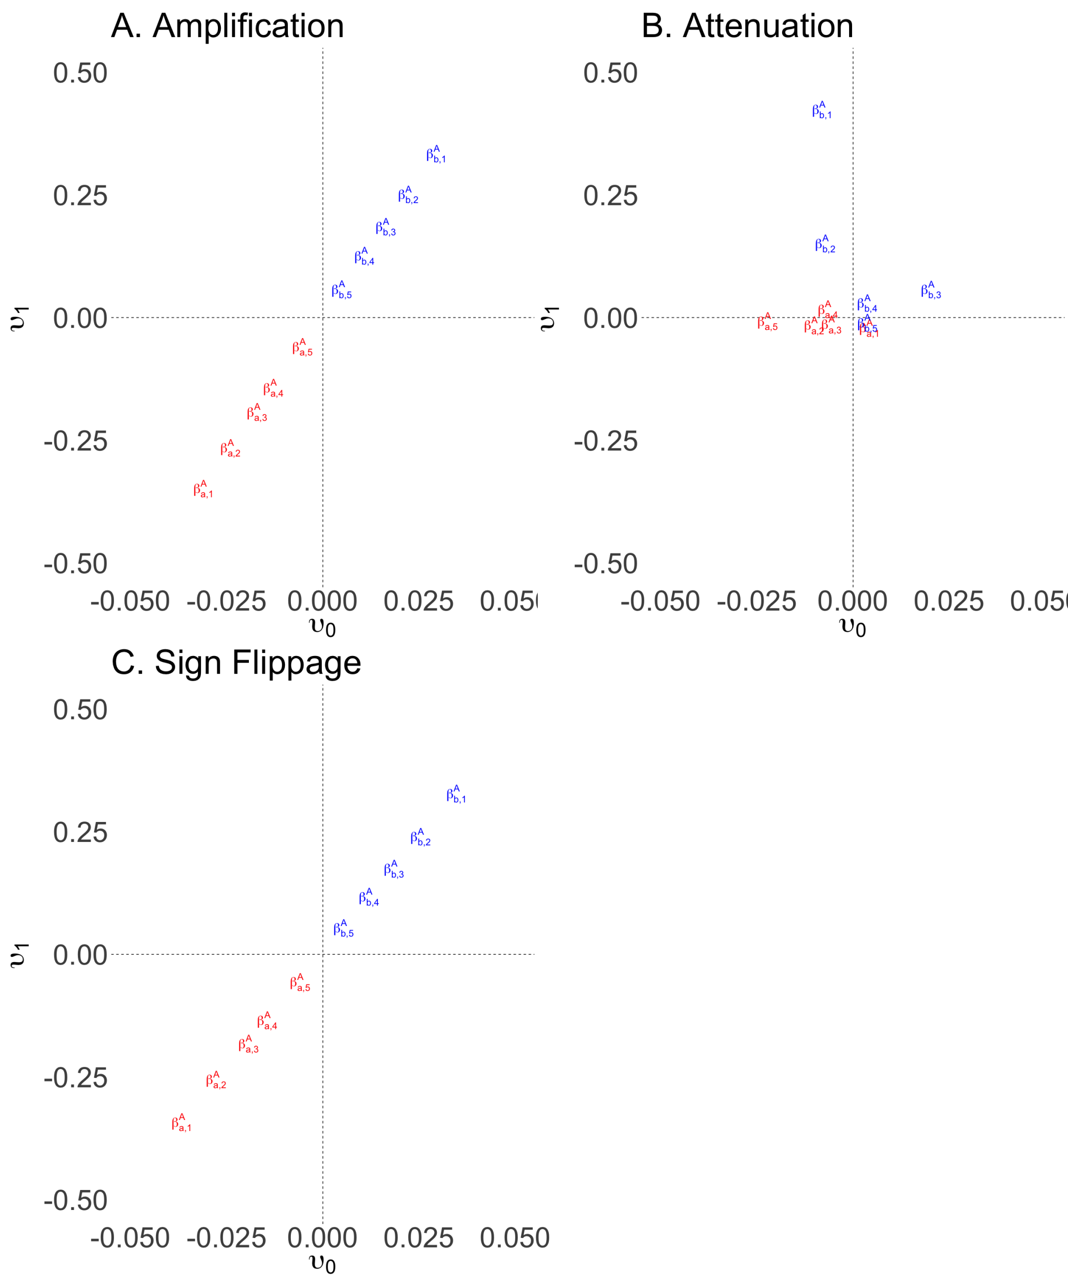

Figure A4: VAR Decomposition Under Different Action-Level Response Curves

Figure A4 decomposes each β slope coefficient between the Υ values. The x-axis plots the υ value multiplied by α_0 for all the β coefficients, while the y-axis does so for α_1. The panels correspond accordingly to the panels in Figure 3. Most slope coefficients are a combination of both the slope and the intercept. Previous B action scalars (depicted in blue) are positive in Panels A and C, but less than one. The scalars become smaller with more lags, highlighting the diminishing effects of previous lags. Conversely, previous A actions tend to receive negative weights on both the slope and intercept in Panels A and C. The coefficients in Panel B tend to have larger slope scalars than intercept scalars, in line with the intercept being null.

The decomposition highlights two points. First, the time-space β slope parameters are a linear combination of all the action-space α parameters. The VAR estimates capture a mix of the slope and intercept for both sides A and B. Second, the scaling terms are, without further assumptions, unbounded. Therefore, the β slope parameters may be amplified, attenuated or even switch signs. Sign flippage can even occur if all the action-space α coefficients are positive because the scaling terms need not be positive.

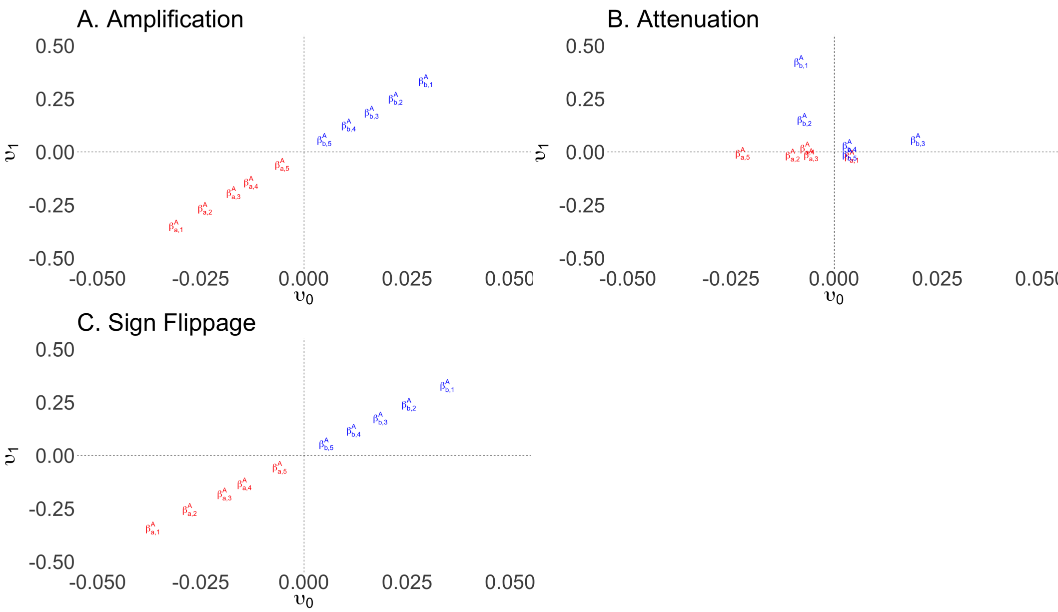

Figure A5: Additional Parameter Decomposition (continued)

D. Additional Estimation for Section 4.2

Figure A6 plots the average VAR coefficients over the simulations. The optimal lag length is identified using BIC for each simulation, then estimated. The points show the average coefficient estimate per lagged value. The 95% confidence intervals use the average standard error over the average response time simulations.

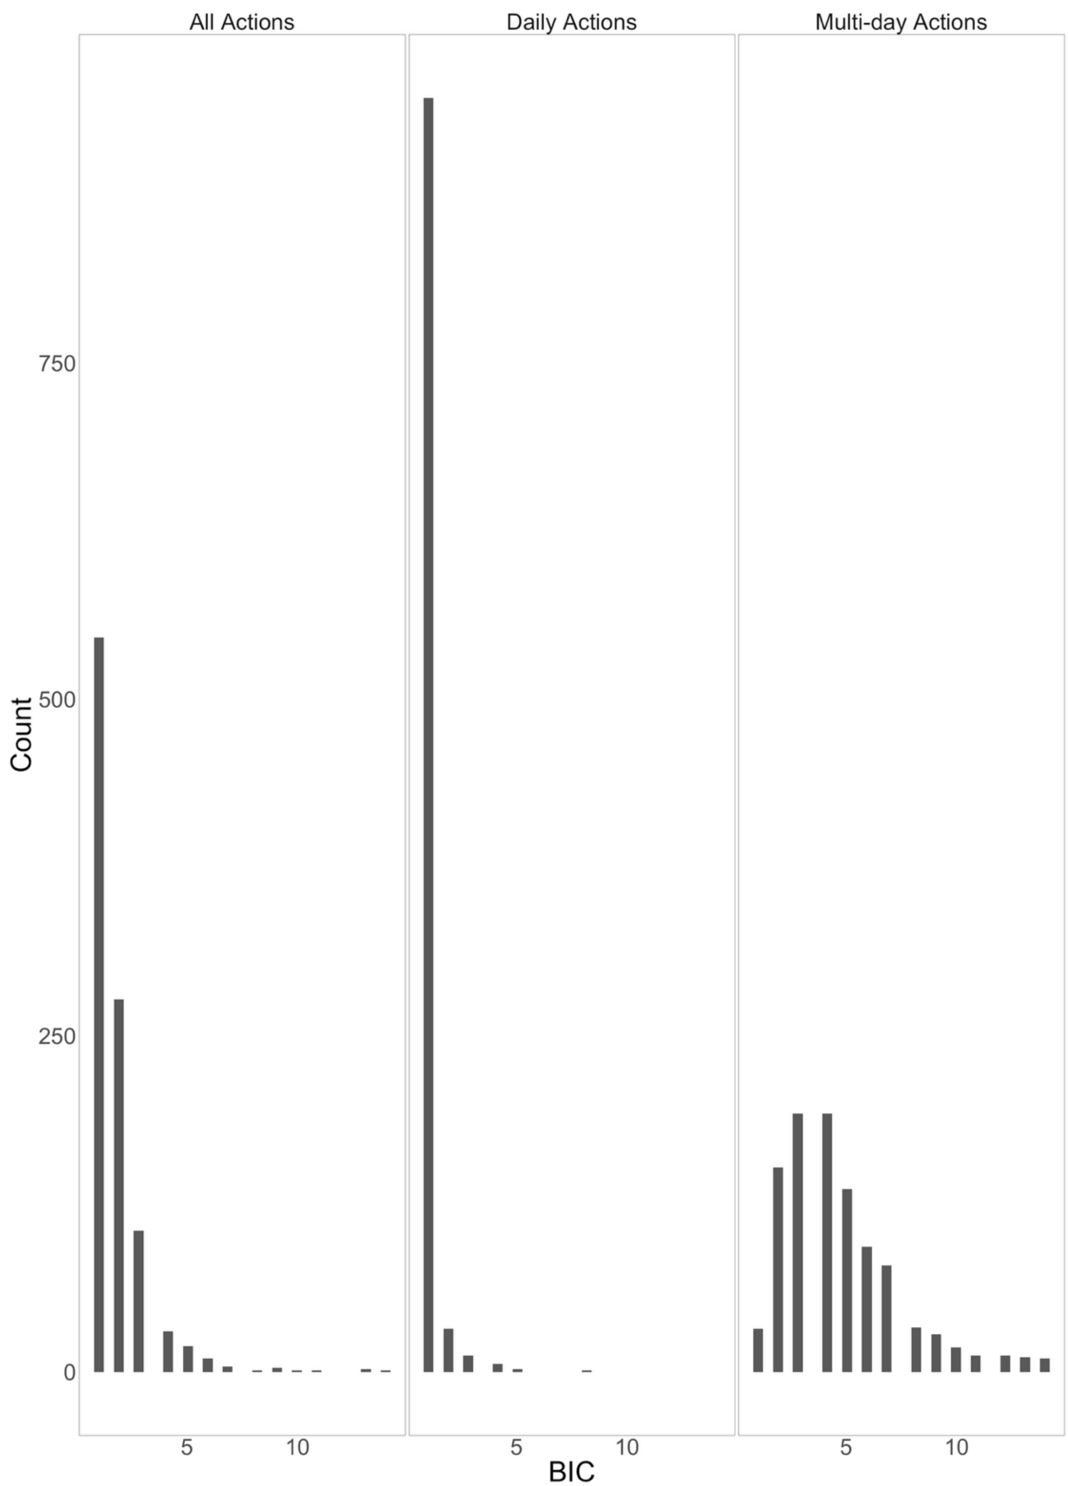

Figure A6: VAR Coefficients Over 1,000 Monte Carlo Simulations for Players A and B

Recall Player A's action-level response curve is $0.01 + (-0.74)\, d_{i-1}$, implying a negative slope. In all three iterations, the coefficients on Player B's past actions are nonpositive. When the response time is shorter (i.e. 1 day), the lagged coefficients beyond four periods tend to be economically and statistically insignificant. Additionally, the magnitudes tend to be much smaller than the action-level slope. Player B's lagged actions also remain statistically significant well beyond the average response times. The time distortion may lead a researcher to erroneously attribute Player A's previous actions in determining future reactions.

Conversely, Player B's action-level response curve had a positive slope. In all three simulation iterations, Player A's lagged daily actions are nonnegative. Player B's own lagged daily actions tend to be negative, but statistically insignificant. Together, these simulations highlight how longer response times create a facade of prolonged reactions.

E. Additional Facts and Figures for Section 5

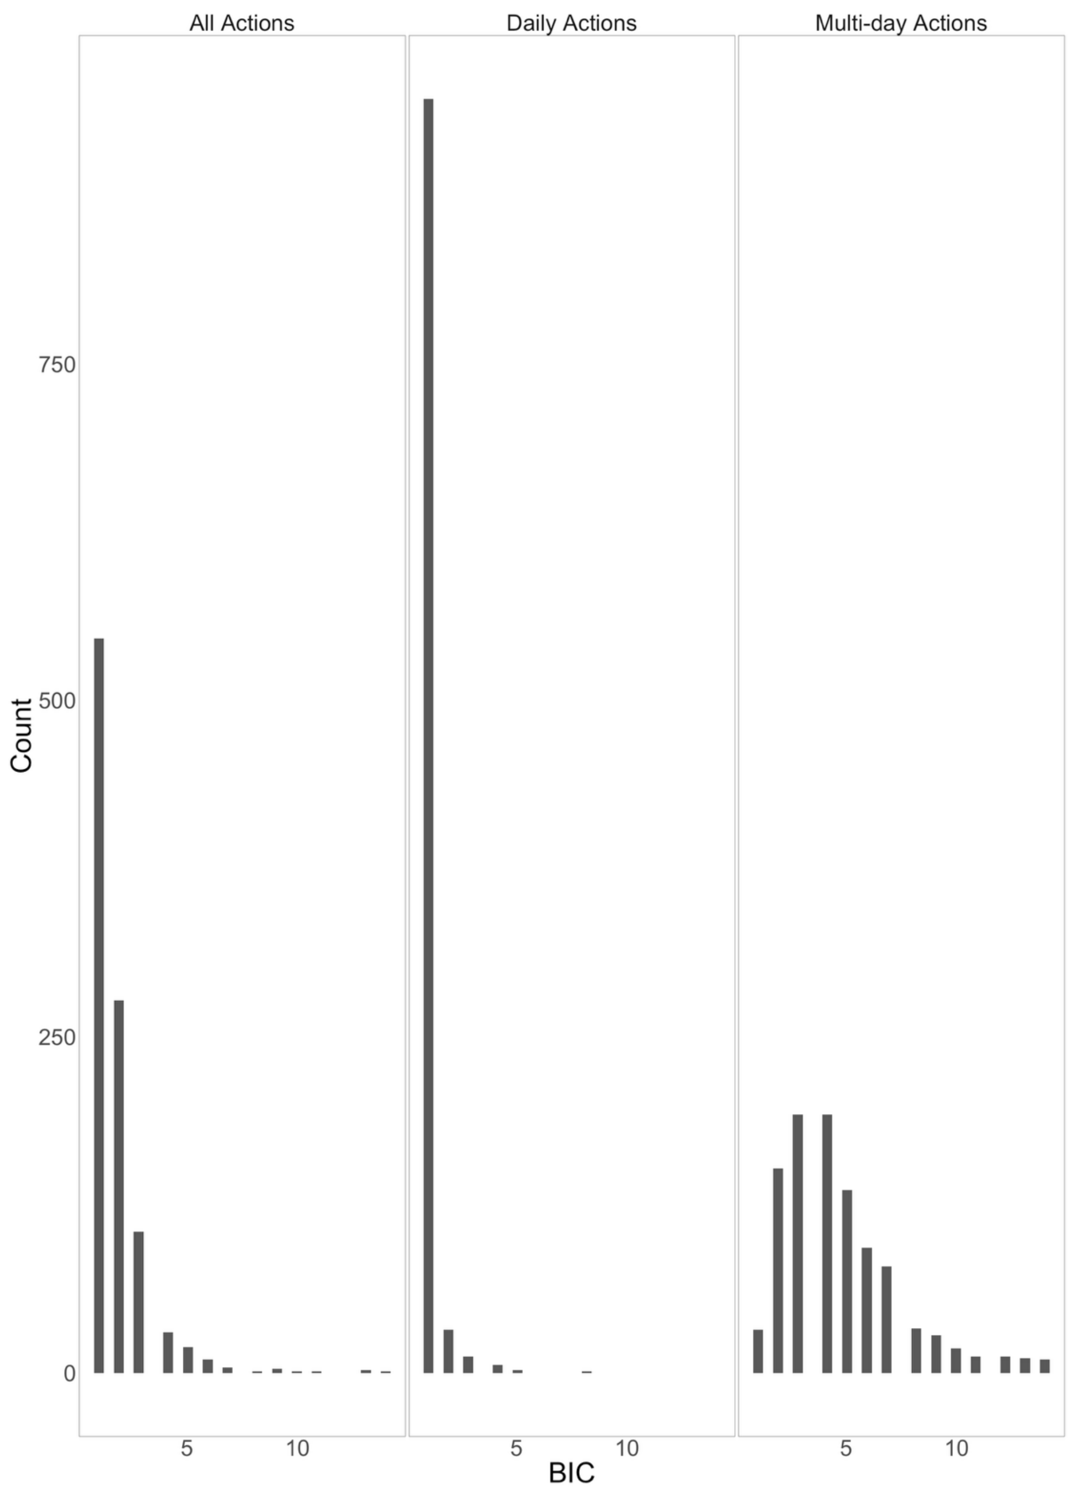

Figure A7: Optimal VAR Lag Length Over 1,000 Monte Carlo Simulations Using Bayesian Information Criterion

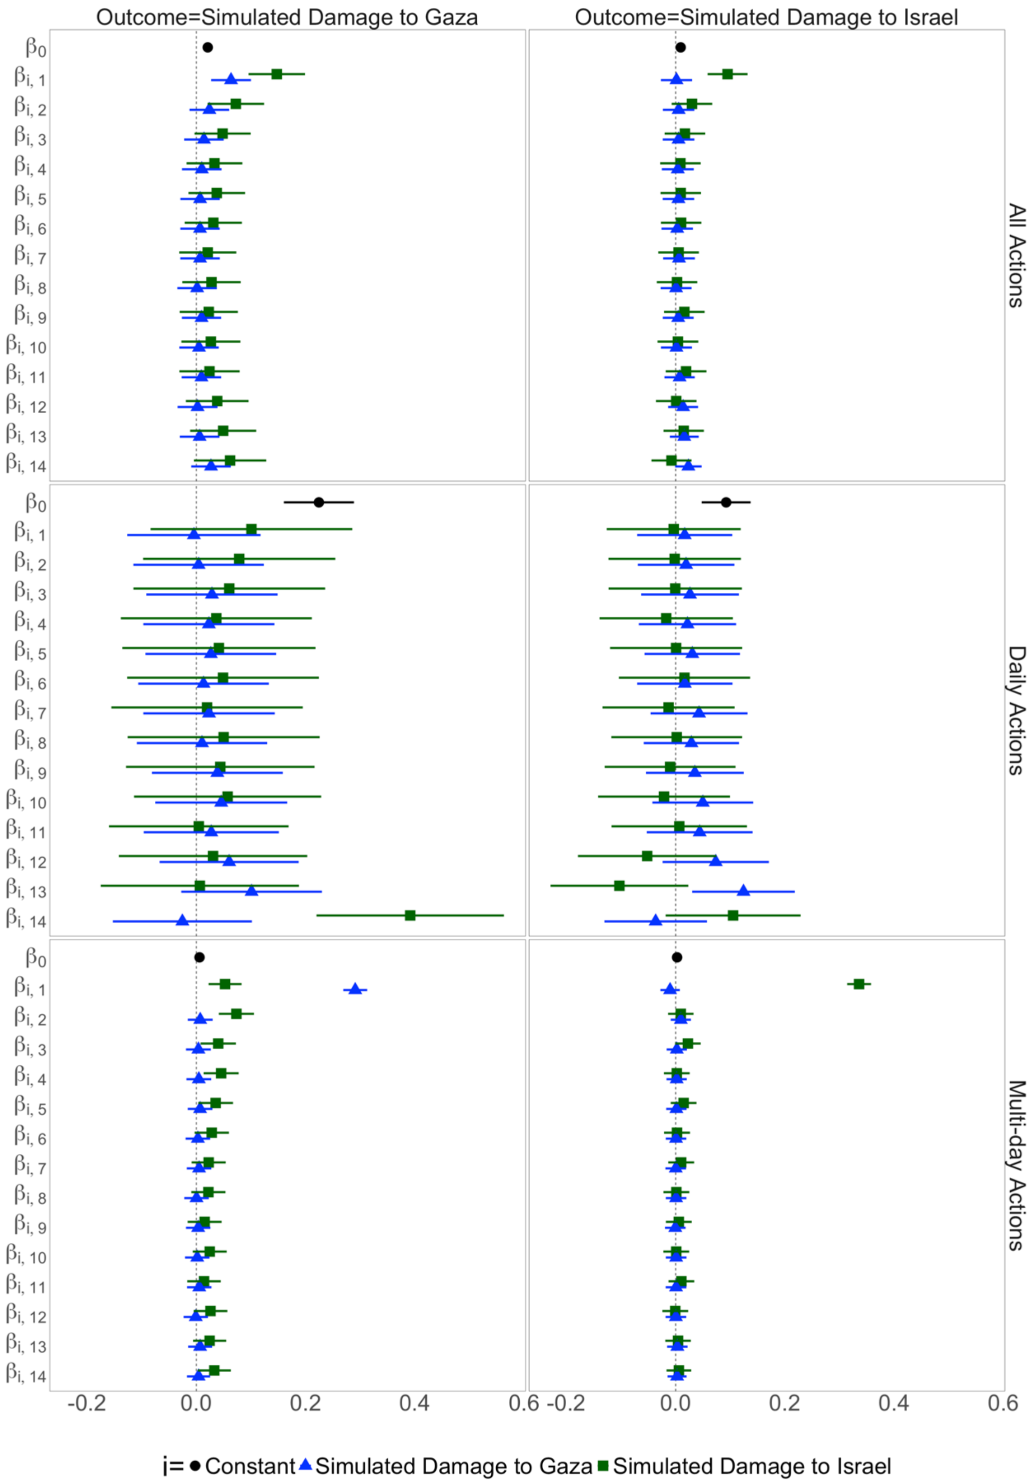

Figure A8: Coefficients from 1,000 Monte Carlo Simulations

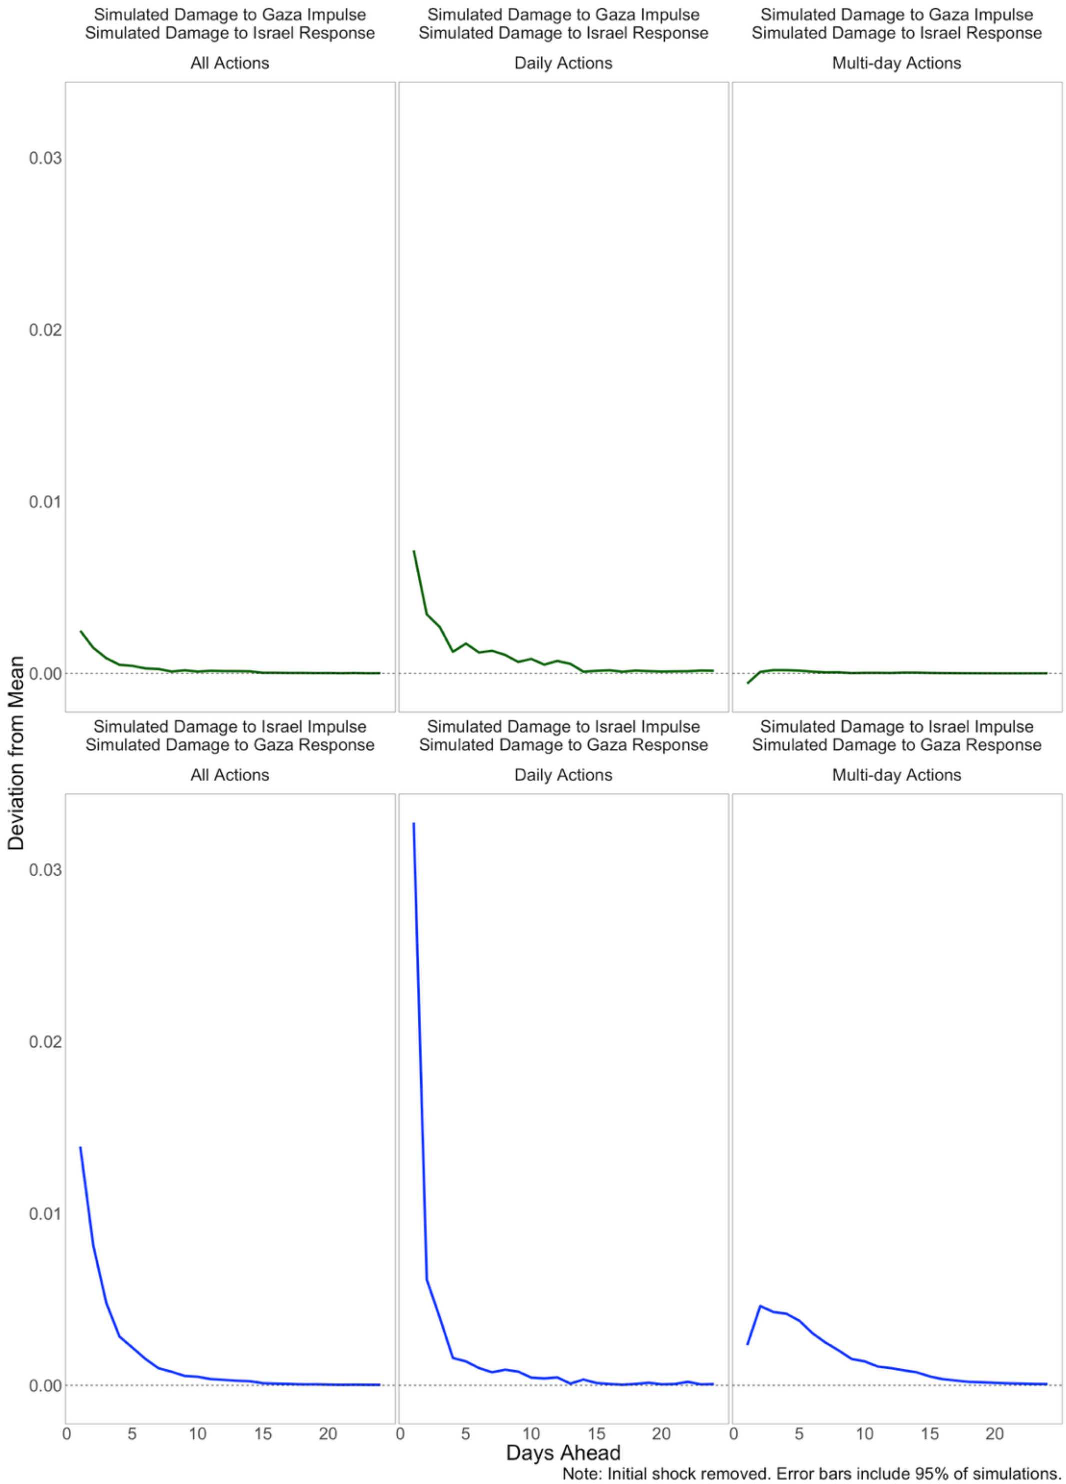

Figure A9: Average IRFs from 1,000 Monte Carlo Simulations

Note: Initial shock removed from graph. Initial one standard deviation shock varies based on estimated standard error.