UC INSTITUTE ON GLOBAL CONFLICT AND COOPERATION

WORKING PAPER NO. 2
APRIL 2025

Timing is Everything: A New Way to Estimate Strategic Behavior

Danny Klinenberg, Eli Berman, and Esteban Klor

Abstract

Many applied economic studies aim to estimate strategic behavior through reaction curves. Examples include two-sided conflicts, or economic trade wars, and algorithmic pricing between firms. Analysis is usually performed at a prespecified time interval, such as days, weeks, months, or years, using a vector autoregression (VAR). Yet sides may respond within a day to one action, but wait a month after another. If data is recorded in arbitrary time intervals, then the researcher may mistake waiting to act for inaction. We analytically show that VAR analyses do not recover true reaction curves if the timing of reaction is not accurately recorded. This misspecification can cause the sign of the VAR coefficient to reverse and misspecified standard errors leading to erroneous inference. We discuss an alternative structural approach rooted in game theory to estimate reaction curves and investigate its usefulness in a Monte Carlo simulation.

Keywords: conflict, game theory, response time, econometrics, vector autoregression, reaction curve

Authors

Danny Klinenberg
Nonresident Fellow, IGCC
Email: dklinenberg@ucsd.edu

Eli Berman
Professor, UC San Diego
Email: elib@ucsd.edu

Esteban Klor
Professor, Hebrew University of Jerusalem
Email: eklor@huji.ac.il

Acknowledgements

Special thanks to Richard Startz and Spencer Sween for advice on the project.

Suggested Citation

Klinenberg, Danny, Eli Berman, and Esteban Klor. 2025. *Timing is Everything*. IGCC Working Paper No 2. escholarship.org/uc/item/50087695

> *"And have no doubt—we will hold all those responsible to account at a time and in a manner [of] our choosing."*
>
> —President Joe Biden discussing the United States' response to missile and drone attacks targeting United States military installations resulting in three casualties.[1]

1. Introduction

A common inquiry in economics is quantitatively measuring how sides react to one another. For example, Jaeger and Paserman (2008) first document the dynamics of violence as cyclical in the Palestinian-Israeli conflict during the Second Intifada. Applying a standard vector autoregression (VAR) using daily deaths, it fails to find evidence that the Palestinians and Israelis engage in a predictable "tit for tat" cycle of violence. We highlight an unexplored challenge in this type of analysis: data being recorded at arbitrary time intervals.

Conflicts can be thought of as a sequential game, where each player performs actions (e.g., airstrikes, mortar fire).[2] We show that if players are reacting to each other's actions but wait a varying amount of time before doing so, then analyzing a conflict using a VAR with data recorded at some time unit, such as the daily level, can result in misspecified lag length, biased point estimates, and biased empirical responses. Put simply, ignoring the waiting dimension of a best response function creates severe misspecification in a VAR analysis. Moreso, we show analytically that this misspecification bias can lead to the VAR coefficient sign reversal.

We provide an approach to address this, stemming from the underlying microstructure of the problem. Our recommendation is to organize the data at the *action* level rather than a time unit like days or weeks whenever possible. This approach is inherently structural and requires the researcher to identify inaction from waiting to act. We analytically derive what happens when the researcher's probability of success departs from an oracle.

---

1 The White House (2024). "Statement from President Joe Biden on Attack on U.S. Service Members in Northeastern Jordan Near the Syria Border." January 28, 2024. https://jo.usembassy.gov/statement-from-president-joe-biden-on-attack-on-u-s-service-members-in-northeastern-jordan-near-the-syria-border/.

2 Maskin and Tirole (1988) make a similar argument when modeling pricing strategies.

We compare the effectiveness and pitfalls of VAR and our structural approach through an empirical Monte Carlo exercise. As our theory predicts, VAR fails to estimate the parameters of the reaction curves. The optimal lag length is longer and closely related to the maximum number of waiting periods between actions. The estimated impulse response functions also produce statistically significant effects for too many periods into the future. Finally, the point estimates are severely biased compared to the reaction function in magnitude and statistical significance.

Our structural solution performs somewhat better. If the researcher had perfect foresight to remove waiting periods, then the rule-based approach captures the reaction curve. However, removing inaction or failing to remove a waiting period can lead to misleading estimates. We investigate the approach's sensitivity to being overly cautious—leaving waiting periods in—and overly aggressive, removing periods of inaction. A small literature has developed from Jaeger and Paserman (2008) studying the Israel-Gaza conflict with VARs. Haushofer, Biletzki, and Kanwisher (2010) extend the analysis to include nonlethal acts of retaliation (e.g., Qassam rocket fire) and conclude that Israeli military actions against Palestinians lead to escalation rather than incapacitation and that the Palestinians are in fact reacting to Israeli behavior.[3] Finally, Asali, Abu-Qarn, and Beenstock (2017) revisit Jaeger and Paserman (2008) focusing on modeling the problem nonlinearly. However, none view the conflict as a sequential game and investigated the implications of strategic behavior in the econometric estimation.

Our findings more generally apply to any strategic interaction where the unit of observation differs from the unit of action. This phenomenon has been noted in previous literature. Noel (2007) writes that "*the 'true' length of a period t [time to respond] as determined by gasoline stations is unlikely to be identical to the length of a period chosen by the econometrician when collecting data*." A large econometric literature has developed around time series analysis' sensitivity to aggregation (e.g., Working 1960; Zellner and Montmarquette 1971). Additional work shows that when data is analyzed at a coarser time interval than the data generating process, there may be biases in the impulse response functions, trend-cycle decomposition, and forecasting (e.g., Brewer 1973; Tiao and Wei 1976; Geweke 1978; Wei 1978; Freeman 1989; Marcellino 1999). These works do not address the issue of sampling at a finer time interval than the data generating process because this was not a common data issue for the time. Given the increasing access researchers have to high frequency data, we expect our findings to solve an ever-growing issue in estimating reaction curves.

3 See Golan and Rosenblatt (2011) for a comment on this work.

Our explanation for why VAR fails in this setting is inspired by the decomposition of ordinary least squares (OLS) coefficients in the presence of heterogeneity common in the causal inference literature. Angrist (1998) decomposes the pooled OLS estimate into a weighted average of subgroups. We extend this train of thought to VAR estimates where the heterogeneity is unknown to the econometrician. In doing so, we take the intuition from the causal inference literature and apply it to time series methods used to study strategic behaviors.

We focus on developing a solution closely tied to the game-theoretic underpinnings of our setting. An alternative approach would be to develop VAR-like estimators that take into account irregularly spaced data. A subfield has developed from Engle and Russell (1998) on such problems in financial econometrics. Estimation is done at the transaction level jointly estimating the time between transactions and characteristics (e.g., Engle 2000). Ait-Sahalia and Mykland (2003) also show that ignoring the randomness in event arrivals leads to a substantial "cost of discreteness." While potentially relevant to our setting, this is a purely reduced-form estimation strategy that does not take into account the strategic behavior of either.

The remainder of the paper is organized as follows: Section 2 provides the setting and analytically derives the potential bias from analyzing actions in time intervals. Section 3 proposes the theoretical method used to recover the reaction curves: reformatting the data to the action level. Section 4 presents a Monte Carlo simulation investigating each approach, and Section 5 concludes.

2. Econometric Theory

We first assume two players are competing in a sequential game, similar to Berman et al. (2023). For simplicity, we assume each side reacts to the last side's action following a Markov one process. Intuitively, this means that the last action is a sufficient statistic for the state of the game.

Suppose the two sides are labeled Y and X. Without loss of generality, Y is the first mover. Assuming a sequential game, player Y will not make a move on even turns and player X will not make a move on odd turns. The action data is indexed as $\{Y_1, X_2, Y_3, X_4, \ldots\}$ to emphasize the nature of the game, where $Y_{a,} X_a \in \mathbb{R}$.

We set the data generating process to be

$$Y_a = g(X_a{-}1) + \epsilon_a \qquad (1)$$

where ϵ_a are independent idiosyncratic shocks, $\mathbb{E}[\epsilon_a|X_{a-1}] = 0$ and $\mathbb{E}[\epsilon_a^2|X_{a-1}] = \sigma_a^2$. The randomness translates to a Markov one process mixed strategy employed in many sequential games, and a common modeling choice in industrial organizations (Noel 2007). Notice that the index a says nothing about the time interval that occurs between actions. There could be a millisecond between Y_a and $X_a{-}1$ and a year between $X_a{-}1$ and $Y_a{-}2$. Finally, we assume $Pr(Y_a = 0|X_{a-1}),\ Pr(X_a = 0|Y_{a-1}) > 0$. This equates to both sides playing a mixed strategy where they choose to not respond some of the time. This is common in conflict, as motivated in Section 1. Similar behavior may be observed in price wars due to menu costs and customers anchoring their expectations on price. Without this assumption on strategic behavior, the main point of our paper would simplify to the trivial advice of "remove the zeros before performing econometric analysis."

2.2 Actions Recorded in Time

Assume the researcher observes the data at a specified time interval, such as days. Furthermore, assume that at most one action occurs per time interval and actions do not span multiple time intervals. Unbeknownst to the researcher, let l be the time interval between actions with $l = \{0, 1, \ldots, L\}$. Notice l is allowed to differ between actions. Because no actions occur between $t-1$ and $t-l$, the researcher observes these "waiting" periods where $\{Y_k, X_k\}_{k=t-1}^{t-l} = \{(0,0)\}_{k=t-1}^{t-l}$. This captures the dual decision the players make: whether to react and, if so, how long to wait before reacting. The econometrician is unable to discern whether the observation $(Y_k, X_k) = (0, 0)$ generated from waiting or nonaction.

Figure 1. Graphical example of actions recorded at prespecified time intervals. Black points represent time intervals in which a side attacked causing nonzero damage. Open points represent time intervals in which a side caused zero damage (performed a nonaction). Time intervals with no markers represent waiting, which are also recorded as zeros.

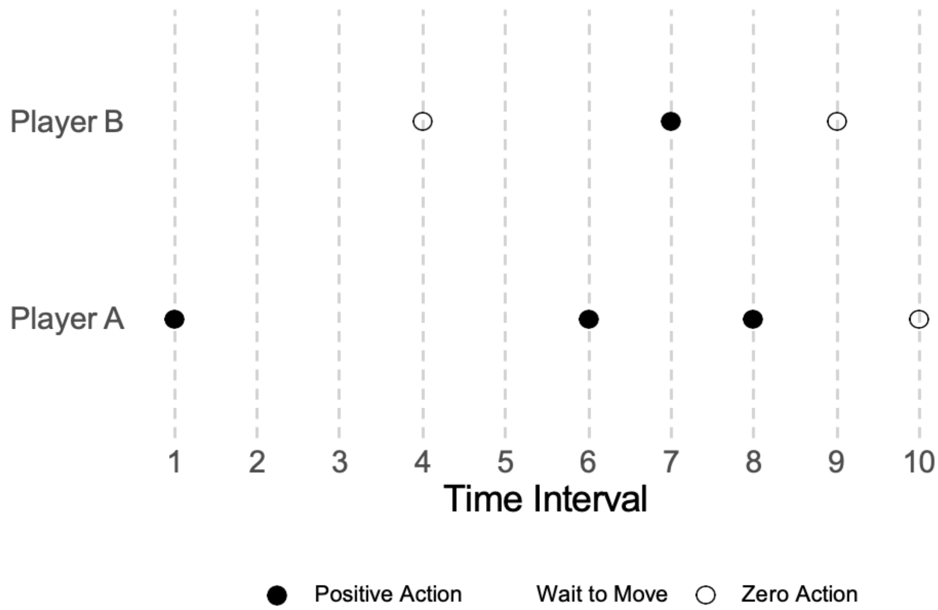

Figure 1 provides a graphical illustration.[4] Suppose Player A and Player B are observed over 10 time intervals (e.g., days). Within each day, a player may perform an action causing positive damage, perform an action causing no damage (inaction), or wait to perform an action. The black circle represents an action that causes positive damage, the open circle represents nonactions (no damage), and empty slots represent waiting. In this example, Player A first inflicts positive damage on Player B (t = 1). Player B then waits three time periods and decides not to respond to 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 that ends with both sides performing nonactions (t = 9 and t = 10).

[4] We take heavy inspiration in this graphic from Brown and MacKay (2023) explaining the reaction of algorithmic pricing.

Studying the sequential game in arbitrary time intervals causes 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 period 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*.

Formally, let St be a variable unobserved by the econometrician equal to the number of lags caused by the data aggregation process. This equates to the number of recorded time periods a side waits before responding to an action. Let S_t = L + 1 if an actor is "waiting" in period t. Then the observed data can be partitioned by the lag length such that:

$$Y_t = \begin{cases} g(X_{t-l}) + \epsilon_t & S_t = l \text{ for } l \in \{1, 2, 3, \dots, L\} \\ 0 & S_t = L + 1 \end{cases} \tag{2}$$

At the action level, Equation (1) is a continuous data generating process (e.g., $\mathbb{E}[Y_a|X_a]$ = $g(X_a)$ $\forall a$). Although the data generating process is continuous, Equation 2 shows that the data aggregation process is piecewise based on the length of response time.

2.3 Pitfalls of VAR

Given time-level data, a common estimation strategy is a reduced-form vector autoregression with order $p \geq l + 1$:

$$Y_t = \beta_0 + \sum_{j=1}^{p} \beta_{y,j} Y_{t-j} + \sum_{j=1}^{p} \beta_{x,j} X_{t-j} + \epsilon_t \tag{3}$$

How does β relate to the action-level reaction function given the difference in response times and waiting observations?

Theorem 1. *Suppose Yt comes from the data aggregation process defined in Equation (2). $\beta_{x,l}$ from Equation (3) can be decomposed as*

$$\beta_{x,l} = \sum_{k=1}^{L+1} w_k^l \beta_{x,l}^k + \left(1 - \sum_{k=1}^{L+1} w_k^l\right) \ddot{\beta}_{x,l} \tag{4}$$

Where

i) $\beta_{x,l}^{k}$ is the least squares parameter using observations when $S_t = k$,
ii)

$$\beta_{x,l} = \frac{cov\,(\mathbb{E}[Y_t\,|S_t], \mathbb{E}[\tilde{X}_{t-l}\,|S\;])}{var(\mathbb{E}[\tilde{X}_{t-l}|S_t])} \qquad (5)$$

where $\tilde{X}_{t-l}$ is the residual of X_{t-l} on all other explanatory variables,

$$w_k^l = \frac{Pr(S_t = k)var(\tilde{X}_{t-l}\,|S_t = k)}{var(\tilde{X}_{t-l})} \qquad (6)$$

The proof follows from standard linear projection properties:

Proof. Let $l \in [0, p]$. Let $\tilde{X}_{t-l}$ be the of X_{t-l} regressed on all other regressors in Equation (3). Then $\beta_{x,l}$ can be decomposed as:

$$\begin{aligned}
\beta_{x,l} &= \frac{cov(Y_t, \tilde{X}_{t-l})}{var(\tilde{X}_{t-l})} && (7)\\
&= \frac{\mathbb{E}[cov(Y_t, \tilde{X}_{t-l}|S_t)] + cov(\mathbb{E}[Y_t|S_t]), \mathbb{E}[\tilde{X}_{t-l}|S_t])}{var(\tilde{X}_{t-l})} && (8)\\
&= \frac{\mathbb{E}[cov(Y_t, \tilde{X}_{t-l}|S_t)\frac{var(\tilde{X}_{t-l}|S_t)}{var(\tilde{X}_{t-l}|S_t)}] + cov(\mathbb{E}[Y_t|S_t]), \mathbb{E}[\tilde{X}_{t-l}|S_t])\frac{var(\mathbb{E}[\tilde{X}_{t-l}|S_t])}{var(\mathbb{E}[\tilde{X}_{t-l}|S_t])}}{var(\tilde{X}_{t-l})} && (9)\\
&= \frac{\mathbb{E}[\beta_{x,l}^{S_t} var(\tilde{X}_{t-l}|S_t)] + \ddot{\beta}_{x,l} var(\mathbb{E}[\tilde{X}_{t-l}|S_t])}{var(\tilde{X}_{t-l})} && (10)\\
&= \frac{\sum_{k=1}^{L+1} Pr(S_t = k) var(\tilde{X}_{t-l}|S_t = k)\beta_{x,l}^{k} + \ddot{\beta}_{x,l} var(\mathbb{E}[\tilde{X}_{t-l}|S_t])}{var(\tilde{X}_{t-l})} && (11)\\
&= \sum_{k=1}^{L+1} w_k^l \beta_{x,l}^{k} + \left(1 - \sum_{k=1}^{L+1} w_k^l\right)\ddot{\beta}_{x,l} && (12)
\end{aligned}$$

The second line applies the law of total covariance. Notice that $\tilde{X}_{t-l}$ must now be calculated separately for each subset of St, or else the residuals will also be biased. The third line multiplies by 1 to each term in the numerator. The fourth line rearranges while the fifth line expands the first term in the numerator. The final line rewrites the terms as weights. It can be shown

$$\sum_{k=1}^{L+1} Pr(S_t = k) var(\tilde{X}_{t-l}|S_t = k) + var(\mathbb{E}[\tilde{X}_{t-l}|S_t]) = var(\tilde{X}_{t-l})$$

meaning that the weights are nonnegative and sum to 1.

Equation (4) showcases two potential issues with estimating reaction functions with a VAR analysis. The first is model specification. If the reaction function includes terms or higher orders not included in the VAR, then the model will suffer from misspecification bias.

Assuming the VAR includes at least all terms in the reaction function, the analysis still may not recover the reaction parameters because of the aggregation across response times and the pooled estimate. Furthermore, the bias cannot be signed. By definition, $\beta_{x,l}^{k} = 0$ when $k \neq l$, meaning that the first term is attenuated toward 0 whenever $w_k^l \neq 1$. The second term captures the variation between means of reaction time. As Simpson (1951) showed in a cross-sectional setting, this term can cause sign reversals and erroneous inference.

Intuitively, analyzing action-level data at an arbitrary time-interval level creates a piecewise data aggregation process based on the response time. Estimating coefficients at the time-interval level leads to averaging over the different data aggregation processes, creating the illusion of long response times. Our results can be concisely summarized in terms of regimes (Hamilton 1989): a coefficient estimated over multiple regimes averages i) the relationship between the outcome and explanatory variable for each regime and ii) the relationship across regimes.

2.3.1 Illustration and Intuition

Imagine that there are two gas stations, X and Y, participating in a Bertrand competition. For simplicity, assume the fundamentals of the profit function remain constant. The objective is to study how the stations change their prices in response to the other price.[5] Let Y_a (X_a) be the *change* in price station Y (X) implements during action a. Also assume that $Y_a = \alpha_0 + \alpha_1 X_{a-1} + \epsilon_a$, where Y_a is the price at action a for gas station Y.

The econometrician records the price of both gas stations four times a day. Unbeknownst to them, the gas station managers who set the price are sometimes distracted with other tasks and don't always notice exactly when prices change. For simplicity, assume that the managers always check at the end of the day, meaning that at most it takes them three quarters of a day to respond to a price change, and the amount of time it takes them to respond is idiosyncratic.

Relating back to the econometric setting, actors X and Y respond after no lull periods, or after one, two, or three lull periods. The number of lull periods after which each side responds is randomly assigned with equal probability. Then for some t such that $S_t = l$, $Y_t = \alpha_0 + \alpha_1 X_{t-l} + \epsilon_t$ is the data generating process for $l \in \{0, 1, 2, 3, 4\}$. The data aggregation process partitions the observations into five groups: the first four refer to which lagged time-interval Y responds. The final group captures the time-interval observations in which Y does not move because they are busy with other tasks (e.g., $Y_a = 0$) or it is X's turn to perform an action.

Suppose the econometrician estimates the reaction function in quarter days (i.e., $\{Y_t, X_t\}$) following Equation (3) with $p = 4$. By definition, $\beta^k_{x,l} = \alpha_1$ for $k = l$ and 0 (including the coefficients on the Y lags). Because every regressor in the action-level data is included in the time-interval level data, the VAR coefficient simplifies to

$$\beta_{x,l} = w^l_l \alpha_1 + \left(1 - \sum_{k=1}^{L+1} w^l_k\right) \ddot{\beta}_{x,l}$$

Similarly, $\beta_{y,l}$ may not equal zero because of the additional cross-response time term, $\ddot{\beta}_{y,l}$.

Figure 2 plots the true reaction function from the action-level data, with slope $\beta^k_{x,l}$, versus the VAR estimates using the time-interval level data, $\beta_{y,l}$ The X axis shows

[5] Maskin and Tirole (1988) model and Noel (2007) estimates such a phenomena in the context of Edgeworth cycles.

X_{t-l} residualized on all other regressors while the y axis is Y_t. The blue dashed line is the true reaction function while the black is the estimate from the VAR. The pluses are time periods in which Y reacted to an X action $t - l$ periods back, the squares when Y reacts to other lag lengths, and the large triangle is when X reacts to Y and the circles are the lull periods.

Figure 2. Comparison of action-level response and VAR coefficient. The blue dashed line is the true action-level response while the black line is the estimate from the time-level VAR estimates.

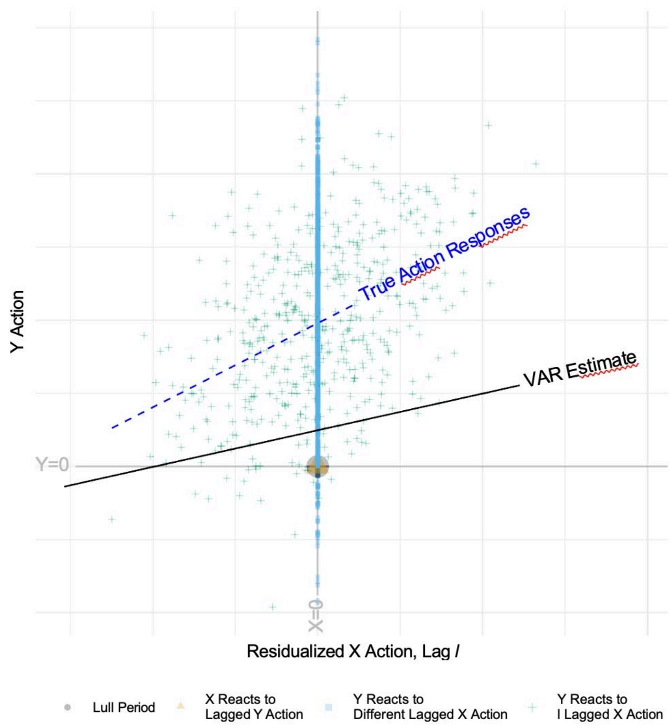

There are three main takeaways from this example. First, $\beta_{y,l}$ will be an attenuated estimate of the true reaction curve if $\beta^k_{x,l} = \alpha_1$. If $\beta^k_{x,l} \neq \alpha_1$., the bias of the estimate will depend on the cross-reaction covariances, which may be positive or negative. Second, side X and Y's varying response times create a facade of many previous prices affecting current prices. Regressing over all the data shows on average each side responds to multiple lags of price levels, when the data generating process shows only the previous price. Even lags in which the actor never responds to (e.g., a lagged Y action for side Y) may be biased because of the cross-group coefficient, $\ddot{\beta}$.Third, the difference in response times means that $\beta_{i,l} \neq 0$ or $i \in \{x, y\}$ and all values of l. This biproduct of the regression estimates aggregating the data to a time-interval can create the illusion of a cyclical behavior.

In conclusion, a VAR analysis may suggest drawn-out long cycles, even when both sides are following a Markov one sequential strategy, if the data is analyzed at a time-interval level rather than at the action level.

3. An Alternative Structural Approach

We instead propose a structural approach to estimating responses. This structural approach requires establishing a rule to determine which periods demonstrate waiting to act and which involve inaction. We discuss considerations in choosing the rule and potential biases introduced when the rule diverges from the oracle rule.

3.1 The Approach

This approach involves removing the waiting periods and estimating the reaction curve from Equation (1). This requires assuming a rule to identify the state of the form $Pr(\text{rule identifies waiting period} \mid \text{previous actions}) = 1$ and 0 otherwise. The *waiting* time intervals are then removed from the dataset. The remaining data follows Equation (1) and estimation of the reaction curve is performed on it.

A VAR specification assumes that the rule is $Pr(S_t = 1) = 0$, meaning that every observation is considered an action. There is no strategic waiting between actions, implying the constant lag structure is a reasonable approximation to the data generating process. Alternatively, assuming that $Pr(S_t = 1 \mid x_t = y_t = 0) = 1$ assumes that every time interval with null actions are assumed to be inactions and can be removed. In this case, both sides are assumed to not be mixing between action and inaction.

There are two considerations when crafting a rule: reaction time constraints and players signaling a response. First, a side may not have the necessary resources for an immediate action. This was the case with the United States reallocating resources to the Middle East after Iranian proxies killed three service members. Conversely, in prolonged conflicts like Israel-Gaza, both sides have all resources necessary to retaliate at the ready—Gazans have quickly deployable mortars while Israelis have drones continuously overhead (Haushofer, Biletzki, and Kanwisher 2010). The second constraint is signaling a response. If the action takes too long, the other side may interpret this as independent of their previous actions. This would suggest a reaction curve independent of the other player, defeating the intended purpose.

In Section 4.3, we investigate how *off* a rule can be and still produce reasonable point estimates and inference. Below, we provide an analytic illustration of the pitfalls of imperfect rules.

3.2 Illustration and Intuition of Misspecification

A misspecified rule can lead to removing inactions or keeping in waiting periods. Section 2.3 showcased the problems associated with keeping all the waiting periods. The other extreme case is removing all potential waiting periods (i.e., $Pr(S_t = 1 | x_t = y_t = 0) = 1$). The bias in the slope from this rule will be attenuated if the coefficient is less than 1 in absolute value.

Suppose $Y_a = \alpha_0 + \alpha_1 X_{a-1} + \epsilon_a$ and $X_a = \alpha_0 + \alpha_1 Y_{a-1} + \epsilon_a$. The researcher observes two series $\{y_t, x_t\}$. Following $Pr(S_t = 1|x_t = y_t = 0) = 1$, they remove all time intervals in which $y_t = x_t = 0$ creating $\{\acute{y}_t | \acute{x}_t\}$ and then estimate $\mathbb{E}[\acute{y}_t | \acute{x}_t] = \gamma_0 + \gamma_1 \acute{x}_t$.

Removing all of the actions such that $y_a = 0$ biases γ_0 term upward, while removing all the actions in which $x_a = 0$ may bias $\gamma 1$ up or down depending on the magnitude of α_1. Following the same steps as in Section 2.3, $\gamma 1$ can be decomposed into states $\mathcal{S}_t$ such that

$$\gamma_1 = \sum_k w_k \gamma_1^k + \left(1 - \sum_k w_k\right) \ddot{\gamma}_1 \qquad (13)$$

where k denotes the number of x actions removed due to the rule. If $k = 1$, then one x action was removed in the rule so $\gamma_1^1 = \frac{cov(y_a, x_{a-3} | \mathcal{S}_t = 1)}{var(x_{a-3} | \mathcal{S}_t = 1)} = \frac{cov(y_a, x_{a-3})}{var(x_{a-3})}$. Notice that $\gamma_1 = 0$ because $\mathbb{E}[\acute{x}_t | \mathcal{S}_t] = \mathbb{E}[\acute{x}_t]$ in this case.[6]

Finally, the sequential nature of the analysis at the action level implies that $\gamma_1^k = \alpha_1^{3k}$. Together, γ_1 can be written as:

$$\gamma_1 = \sum_k w_k \alpha_1^{3k}$$

Where $\sum_k w_k = 1$. Following the same logic from earlier, $w_k = \Pr(\mathcal{S}_t = k) \frac{var\,(x_{a-k} | \mathcal{S}_t = k)}{var(x_{a-k})} = \Pr\,(\mathcal{S}_t = k)$. In other words, γ_1 is weighted based on the probability of each state. If $w_0 = 1$, the rule perfectly removed the waiting periods without removing any of the inactions.

If $|\alpha_1| < 1$, then γ_1 will be attenuated following this rule. Conversely, if $|\alpha_1| > 1$, then γ_1 will be biased upward. The findings from this rule are informative if $|\gamma_1| < 1$ and statistically significant or $|\gamma_1| > 1$ and is a precisely estimated null effect.

---

[6] $\mathbb{E}[\acute{x}_t | \mathcal{S}_t = k] = \mathbb{E}[\alpha_0 + \alpha_1 y_{a-k} | \mathcal{S}_t]$ by definition. However, $\mathbb{E}[y_{a-k} | \mathcal{S}_t] = \mathbb{E}[y_{a-k}]$ because $\mathcal{S}_t$ only affects the indexing in time.

4. Simulations

We develop a simulation study to compare the VAR estimation to the structural approach. Section 4.1 introduces the simulation setup, Section 4.2 showcases how a VAR analysis may lead to erroneous conclusions in a strategic framework with waiting, and Section 4.3 discusses the pros and cons of our structural approach.

4.1 Simulation Setup

We first generate the data at the action level, then insert waiting periods. We define $Y_a = max(0, \alpha_0 + \alpha_1 X_{a-1} + \epsilon_a)$ and $X_a = max(0, \alpha_0 + \alpha_1 Y_{a-1} + \epsilon_a)$, where $\epsilon \sim \mathcal{N}(0, 1)$. Player Y makes a move on odd actions while player X makes a move on even actions. Turns in which a player is not moving is recorded as 0. We set $(\alpha_0, \alpha_1) = (1, .5)$ during the simulations.

After generating 500 actions for each side, we add zero, one, or two waiting periods between actions. The wait time is drawn from a uniform distribution. Section A.1 provides example data. We perform this Monte Carlo analysis 1000 times.

For the VAR, we first use the Bayesian Information Criterion (BIC) to determine optimal lags, then calculate the significance of each lag term. We present the point estimates from one side as well as the orthogonal impulse responses.

For the structural approach, we start with the oracle rule (i.e., $Pr(S_t = L + 1) = 1$) and then decrease the probability by 0.1 until we have a rule that keeps all the waiting periods and removes all the inactions (i.e., $Pr(S_t = L + 1) = 0$). In practice, we assign an inclusion probability based on whether the side is performing an inaction or waiting. In all iterations, we assume the reaction curve specification is known. Specifically, we estimate $\mathbb{E}[\acute{y}_t | \acute{x}_t] = \alpha_0 + \alpha_1 \acute{x}_t$ as the reaction curve.

Our ultimate goal is to understand if these methods properly recover the parameters of the reaction curve. For the VAR analysis, we focus on the i) optimal lag length, ii) coefficient point estimates and statistical significance, and iii) impulse response functions. For the structural approach, we report the point estimate for the intercept and slope and whether the estimates are statistically different from the true reaction curves parameters.

4.2 VAR Findings

We first investigate how a VAR performs in this setting using Equation 3 by estimating the optimal lag length based on the BIC, the coefficient estimates, and orthogonal impulse responses.

Figure 3 plots the histogram of optimal lag lengths over all the simulations. The median optimal lag length is four. This aligns with the maximum number of waiting periods being three in the simulation.

Figure 3. Optimal VAR lag length

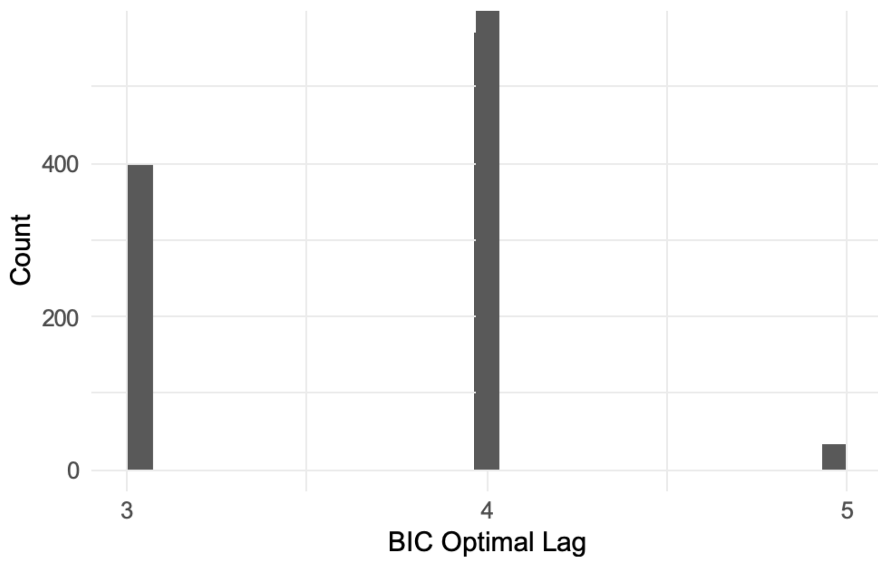

Taken at face value, this exercise suggests that an action is dependent on the past four actions. This is somewhat true—actions followed by three waiting periods are dependent on the fourth lag. However, some actions depend on only one lag while others depend on two or three. The BIC approach to optimal lag length is correctly identifying the largest number of periods one side waits to respond to another.

Figure 4 plots the average point estimate and 95 percent of the Monte Carlo distribution for each lagged coefficient. The coefficients are statistically significant and large across the simulations. While Y's actions are defined to depend only on X's previous action, the VAR analysis suggests that the five previous actions of X and Y are associated with the current action. In addition, none of the point estimates equal $\alpha 1$. Together, this suggests that a VAR analysis will not accurately capture the underlying coefficient of the reaction function.

Figure 4. Average VAR coefficient over Monte Carlo simulations with 95 percent simulated confidence intervals. Outcome is player Y's actions.

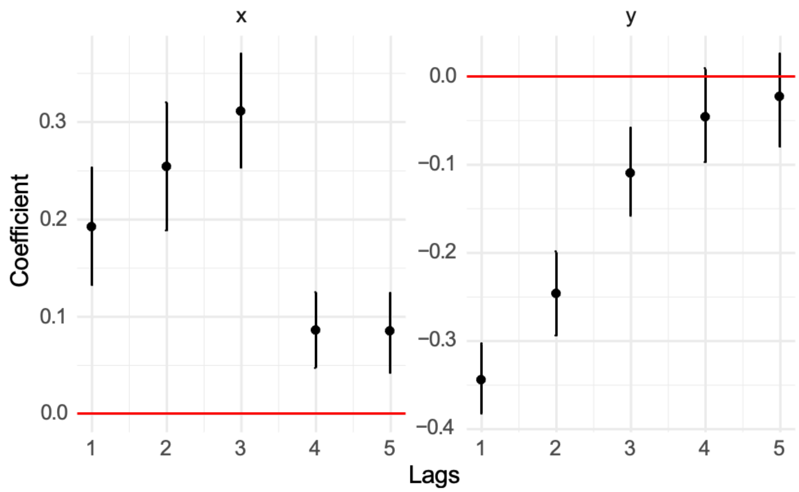

Finally, the impulse response function also suggests longer effects of an action than the reaction curve. Figure 5 plots the average and 95 percent confidence intervals of Y's impulse response to a shock on X. There is consistent evidence of persistent effects from a shock to X. However, the reaction curve was designed to depend only on the previous action. The persistent and prolonged effects are due to the heterogeneity in wait times. In practice, this means that a researcher may overstate the persistence of one side's action simply due to data aggregation.

Figure 5. Orthogonal impulse response function of Y on a shock to X with 95 percent simulated confidence intervals.

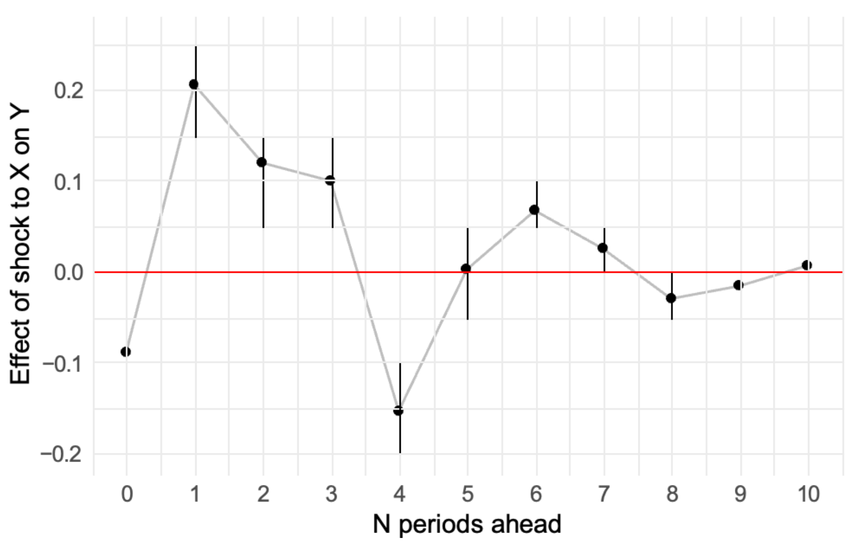

4.3 Structural Approach

Next, we present the effectiveness of our structural approach. Our proposal requires a researcher employ a rule identifying if periods of zero action are inaction or waiting periods. Table 1 varies the correctness of this rule. The first entry is 1, which means the researcher perfectly removes the waiting periods without removing the inaction periods. As the rule quality decreases, the probability of removing a waiting period decreases while the probability of inadvertently removing an action increases. The last row is 0, meaning the researcher removes all inaction periods while keeping the waiting periods. When the rule is 50 percent, it means the researcher’s rule is equivalent to a coin toss.

Table 1. 1000 Monte Carolo Simulations of Rules-Based Approach

	Monte Carlo parameter estimate				Monte Carlo *t*-statistic vs. true parameter value			
Pr(Rule removes waiting period)	Mean	Median	2.5 percentile	97.5 percentile	Mean	Median	2.5 percentile	97.5 percentile
$\alpha_0 = 1$								
1 (Oracle)	1.05	1.05	0.89	1.23	0.57	0.55	-1.28	2.50
0.975	1.15	1.16	0.99	1.34	1.74	1.74	-0.11	3.70
0.95	1.24	1.24	1.07	1.43	2.82	2.80	0.78	5.04
0.925	1.31	1.32	1.13	1.50	3.70	3.73	1.51	5.92
0.9	1.38	1.38	1.21	1.57	4.58	4.54	2.49	6.78
0.8	1.56	1.56	1.40	1.72	7.38	7.39	5.33	9.58
0.7	1.69	1.69	1.54	1.85	9.76	9.76	7.49	11.99
0.6	1.78	1.78	1.62	1.94	11.58	11.60	9.37	13.96
0.5 (Random Guess)	1.85	1.85	1.70	1.99	13.22	13.22	10.97	15.44
0.4	1.90	1.90	1.74	2.05	14.57	14.56	12.18	16.87
0.3	1.94	1.94	1.79	2.08	15.79	15.76	13.35	18.15
0.2	1.97	1.97	1.83	2.09	16.76	16.70	14.43	19.29
0.1	2.00	2.00	1.88	2.14	17.80	17.73	15.48	20.24
0 (Remove Inaction Only)	2.03	2.03	1.90	2.15	18.62	18.59	16.28	21.08
$\alpha_1 = .5$								
1 (Oracle)	0.48	0.49	0.41	0.56	-0.39	-0.38	-2.26	1.61
0.975	0.44	0.45	0.37	0.52	-1.41	-1.40	-3.38	0.50
0.95	0.41	0.41	0.33	0.49	-2.35	-2.34	-4.46	-0.20
0.925	0.39	0.38	0.30	0.47	-3.02	-3.05	-5.18	-0.82
0.9	0.36	0.36	0.29	0.44	-3.71	-3.72	-5.65	-1.71
0.8	0.29	0.29	0.22	0.36	-5.70	-5.67	-7.74	-3.76
0.7	0.24	0.24	0.17	0.32	-7.19	-7.23	-9.34	-5.00
0.6	0.21	0.21	0.14	0.28	-8.02	-8.05	-9.89	-6.01
0.5 (Random Guess)	0.19	0.19	0.11	0.26	-8.65	-8.64	-10.57	-6.74
0.4	0.17	0.17	0.10	0.25	-9.00	-9.00	-10.97	-6.97
0.3	0.16	0.16	0.08	0.23	-9.16	-9.15	-11.13	-7.10
0.2	0.15	0.15	0.07	0.22	-9.17	-9.17	-11.12	-7.24
0.1	0.13	0.13	0.05	0.22	-9.24	-9.22	-11.26	-7.21
0 (Remove Inaction Only)	0.13	0.13	0.05	0.21	-9.09	-9.08	-11.24	-7.07

[a] Monte Carlo simulations are repeated 1000 times. OLS is used to estimate the parameters and the latent coefficient values are presented. The intercept is set to 1 and the slope to 0.5. Column 1 shows the probability of correctly identifying a waiting period. If the rule was x percent effective, then a waiting row is removed with probability x and a row in which y moved (but caused no damage) is removed with probability 1-x. The oracle rule is displayed first. The next three columns present the average and Monte Carlo distribution for the intercept, followed by the Monte Carlo test statistics of whether the estimate is statistically different from the true parameter value. The last six rows are analogous for the slope.

A misspecified rule can lead to large distortions in magnitude and inference. When the researcher imposes a rule that is right 90 percent of the time, the point estimate for the slope is biased downward and the intercept is biased upward. The bias becomes bigger as the rule becomes worse and inference becomes unreliable. Therefore, slope coefficients from a misspecified rule should be observed as overly conservative.

The bias is driven by keeping the waiting periods in and inadvertently removing actions. Having erroneous rows of zero increases the intercept simultaneously attenuating the slope. Intuitively, this is driven by a decrease in the covariance between X and Y. In periods where Y is waiting, the value is always 0. This erroneous set of data points dampen the covariance used to estimate the slope, also mechanically forcing the intercept to be larger (in absolute terms) in order to minimize the sum of squares error. Conversely, a misspecified rule also removes a player's actions, further adding noise to the reaction curve estimation process.

There is a little slack for rule misspecification. Our simulations suggest that a nearly perfect rule introduces minimal bias and inference is not misleading. Anything else will cause slope coefficients to be attenuated and intercepts to be inflated.

5. Conclusion

Estimating a player's reaction curve during strategic interactions is important for both academics and policymakers. We show that if the underlying nature of a conflict is not taken into account, then standard econometric practices can lead to erroneous results. In our examples, we assume two players engage in a sequential game with varying response times. Even in an overly simplistic setting where players only respond to the opponent's last action, traditional VAR analysis can misspecify the length of response to a shock, the optimal lag length, and the coefficient estimates. To address these issues, researchers should consider organizing their data at the level of the players' actions rather than standard time intervals.

Work Cited

Ait-Sahalia, Yacine, and Per A. Mykland. 2003. "The Effects of Random and Discrete Sampling When Estimating Continuous-Time Diffusions." *Econometrica* 71(2): 483–549. https://doi.org/10.1111/1468-0262.t01-1-00416.

Angrist, Joshua D. 1998. "Estimating the Labor Market Impact of Voluntary Military Service Using Social Security Data on Military Applicants." *Econometrica* 66 (2):249. https://doi.org/10.2307/2998558.

Asali, Muhammad, Aamer S. Abu-Qarn, and Michael Beenstock. 2017. "The Cycle of Violence in the Second Intifada: Causality in Nonlinear Vector Autoregressive Models." *Journal of Applied Econometrics* 32 (6): 1197–1205. https://doi.org/10.1002/jae.2563.

Berman, Eli, Prabin Khadka, Danny Klinenberg, and Esteban F. Klor. 2023. "Deterrence Through Reaction Curves: An Empirical Analysis of the Israel-Gaza Conflict." *Working Paper.*

Brewer, K. R. W. 1973. "Some Consequences of Temporal Aggregation and Systematic Sampling for ARMA and ARMAX Models." *Journal of Econometrics* 1 (2): 133–54. https://doi.org/10.1016/0304-4076(73)90015-8.

Brown, Zach Y., and Alexander MacKay. 2023. "Competition in Pricing Algorithms." *American Economic Journal: Microeconomics* 15 (2): 109–56. https://doi.org/10.1257/mic.20210158.

Engle, Robert F. 2000. "The Econometrics of Ultra-High-Frequency Data." *Econometrica* 68 (1): 1–22. https://doi.org/10.1111/1468-0262.00091.

Engle, Robert F., and Jeffrey R. Russell. 1998. "Autoregressive Conditional Duration: A New Model for Irregularly Spaced Transaction Data." *Econometrica* 66 (5): 1127. https://doi.org/10.2307/2999632.

Freeman, John R. 1989. "Systematic Sampling, Temporal Aggregation, and the Study of Political Relationships." *Political Analysis* 1 (January): 61–98. https://doi.org/10.1093/pan/1.1.61.

Geweke, John. 1978. "Temporal Aggregation in the Multiple Regression Model." *Econometrica* 46 (3): 643. https://doi.org/10.2307/1914238.

Golan, David, and Jonathan D. Rosenblatt. 2011. "Revisiting the Statistical Analysis of the Israeli–Palestinian Conflict." *Proceedings of the National Academy of Sciences* 108 (15). https://doi.org/10.1073/pnas.1016378108.

Hamilton, James D. 1989. "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle." *Econometrica* 57 (2): 357–84. https://doi.org/10.2307/1912559.

Haushofer, Johannes, Anat Biletzki, and Nancy Kanwisher. 2010. "Both Sides Retaliate in the Israeli–Palestinian Conflict." *Proceedings of the National Academy of Sciences* 107 (42): 17927–32. https://doi.org/10.1073/pnas.1012115107.

Jaeger, David A, and M. Daniele Paserman. 2008. "The Cycle of Violence? An Empirical Analysis of Fatalities in the Palestinian-Israeli Conflict." *American Economic Review* 98 (4): 1591–1604. https://doi.org/10.1257/aer.98.4.1591.

Marcellino, Massimiliano. 1999. "Some Consequences of Temporal Aggregation in Empirical Analysis." *Journal of Business & Economic Statistics* 17 (1): 129–36. https://doi.org/10.1080/07350015.1999.10524802.

Maskin, Eric, and Jean Tirole. 1988. "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles." *Econometrica* 56(3): 571. https://doi.org/10.2307/1911701.

Noel, Michael D. 2007. "Edgeworth Price Cycles: Evidence from the Toronto Retail Gasoline Market." *The Journal of Industrial Economics* 55 (1): 69–92. http://www.jstor.org/stable/4622374.

Simpson, E. H. 1951. "The Interpretation of Interaction in Contingency Tables." *Journal of the Royal Statistical Society Series B: Statistical Methodology* 13 (2): 238–41. https://doi.org/10.1111/j.2517-6161.1951.tb00088.x.

Tiao, G. C., and W. S. Wei. 1976. "Effect of Temporal Aggregation on the Dynamic Relationship of Two Time Series Variables." *Biometrika* 63 (3): 513–23. https://doi.org/10.1093/biomet/63.3.513.

Wei, William W. S. 1978. "The Effect of Temporal Aggregation on Parameter Estimation in Distributed Lag Model." *Journal of Econometrics* 8 (2): 237–46. https://doi.org/10.1016/0304-4076(78)90032-5.

Working, Holbrook. 1960. "Note on the Correlation of First Differences of Averages in a Random Chain." *Econometrica* 28 (4): 916. https://doi.org/10.2307/1907574.

Zellner, Arnold, and Claude Montmarquette. 1971. "A Study of Some Aspects of Temporal Aggregation Problems in Econometric Analyses." *The Review of Economics and Statistics* 53 (4): 335–42. https://doi.org/10.2307/1928734.

A. Appendix

A.1 Example Action-Level Data

The data is originally generated at the action level. Table A1 shows the data after it is converted to time-interval notation. The labels show if player X performed an action, player Y performed an action, or it was a lull period.

Table A1. Example Excerpt of Data at Time-Interval Level.

Time	Y Action	X Action	Player Move	Data Aggregation for Y	Data Aggregation for X
1	0.00	0.0	Waiting	0	0
2	0.00	0.0	Waiting	0	0
3	0.37	0.0	Y	$g(x_{t-3})$	0
4	0.00	2.5	X	0	$g(y_{t-1})$
5	0.00	0.0	Waiting	0	0
6	0.00	0.0	Waiting	0	0
7	3.85	0.0	Y	$g(x_{t-4})$	0
8	0.00	0.0	Waiting	0	0
9	0.00	1.4	X	0	$g(y_{t-2})$
10	0.00	0.0	Y	$g(x_{t-1})$	0

A side may choose to perform no action. While this is recorded as a 0, it comes from a different data generating process than the zeros from the waiting rows.