Advance Access publication March 15, 2012

Political Analysis (2012) 20:292–315
doi:10.1093/pan/mps001

A Bayesian Poisson Vector Autoregression Model

Patrick T. Brandt and Todd Sandler
School of Economic, Political, and Policy Sciences, The University of Texas at Dallas, 800 W. Campbell Road, GR 31, Richardson, TX 75080
e-mail: pbrandt@utdallas.edu (corresponding author), tsandler@utdallas.edu

Edited by Vera Troeger

Multivariate count models are rare in political science despite the presence of many count time series. This article develops a new Bayesian Poisson vector autoregression model that can characterize endogenous dynamic counts with no restrictions on the contemporaneous correlations. Impulse responses, decomposition of the forecast errors, and dynamic multiplier methods for the effects of exogenous covariate shocks are illustrated for the model. Two full illustrations of the model, its interpretations, and results are presented. The first example is a dynamic model that reanalyzes the patterns and predictors of superpower rivalry events. The second example applies the model to analyze the dynamics of transnational terrorist targeting decisions between 1968 and 2008. The latter example's results have direct implications for contemporary policy about terrorists' targeting that are both novel and innovative in the study of terrorism.

1 Introduction

Event count models are workhorses in the analyses of many social phenomena. This is largely because there are many concepts where counts of events are easily collected (particularly over time). Despite counts of many phenomena, the bulk of the count models, employed in political science, are univariate. These univariate count models have been extended to allow for various dynamic count processes from moving averages, to autoregressions, to changepoints (Brandt et al. 2000; Brandt and Williams 2001; Park 2010).

As widely as event count methods, such as Poisson and negative binomial regressions, have been used in political science over the last two decades, analysts are still forced to confront the fact that the models in their toolkit are generally univariate and possibly dynamic. If one wishes to analyze dynamic multivariate counts—such as those we see in different types of transnational terrorism events—there are few if any options for capturing the properties of the data in a model. Data such as these are *dynamic multiple time series of counts*. Here, there are multiple outcomes of interest such as counts of attacks on different kinds of terrorist targets over time. What the literature is missing is a flexible, dynamic, multivariate model for the counts of these different events.

Here, we present a new general model for dynamic multivariate counts. The reason for this model is that there are often questions about the endogenous and possibly dynamic relationships among count series. For example:

1. Militarized interstate dispute (MID) or event data on conflicts (either monadic or dyadic actions) involving multiple belligerents over time may have dynamic feedback.
2. Studies of terrorism often are interested in the endogenous dynamic interactions among different types of attacks (bombing, kidnappings, etc.) since this is evidence of complementarity or substitution by terrorists.
3. Analyses of word or topic counts may want to capture dynamic responses across subject areas or speakers.

Authors' note: We are grateful to Janet Box-Steffensmeier, Walter Enders, John Freeman, Jeff Gill, Sara Mitchell, Xun Pang, Vera Troeger, and three anonymous reviewers for their comments. Supplementary materials for this article are available on the *Political Analysis* Web site.

© The Author 2012. Published by Oxford University Press on behalf of the Society for Political Methodology.
All rights reserved. For Permissions, please email: journals.permissions@oup.com

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press

Each of these examples—and many other applications where the outcome variables are several counts—calls for a model that is dynamic and endogenous. Modeling endogeneity of the counts is critical in these applications as well. The endogeneity can either be contemporaneous correlation among the counts (e.g., correlation at time t in the attacks on different kinds of targets) or dynamic (e.g., correlation over time whereby attacks on one target type predict attacks on another target type or count series). Without this model, these event counts will be part of a possibly misspecified count regression. This paper addresses the need for dynamic endogenous count models. We review the extant options and discuss their limitations. We then develop a new multivariate event count time-series model, the Bayesian Poisson vector autoregression (BaP-VAR), to characterize the dynamics of a vector of counts over time (e.g., terrorist targeting decisions that account for the interdependencies of the four target-type time series). Next, we outline how these models are specified, estimated, and interpreted. The interpretation methods include the derivation of an impulse response function tracing out the anticipated speed of reaction of the count series to a shock (based on Cholesky decomposition orderings). This allows one to model the endogenous interactions of the counts. We also show how to compute the decompositions of the forecast error variances (DFEVs) to glean additional dynamic insights. Finally, our examples also include an application with exogenous covariates, allowing us to present a method for dynamic multiplier analyses using the model.

2 Review of Possible Count Models

The goal here is to explore and develop a model for counts that is both multivariate and dynamic. There are three reasons why existing methods are unsatisfactory for addressing possibly dynamic multivariate count data. First, past research typically assumes that the count series are independent and thus ignore the endogeneity among multiple counts. Second, existing multivariate time-series models often make distributional assumptions that ignore the event count nature of the data. Third, the dynamic models of counts in the literature are univariate.

The literature on count models moves in several different directions, addressing these concerns about dynamics and endogeneity. King (1988) shows with Monte Carlo simulations and an empirical example that using an ordinary least squares or logged ordinary least squares produces biased and inefficient estimates compared with a Poisson regression specification. This means that linear regression-based estimators (or analogous time-series methods) will generally be inadequate for modeling counts since the estimates will be biased and inefficient. This is particularly the case if there are observed zero counts in the data.

Brandt et al. (2000) and Brandt and Williams (2001) extend the Poisson regression estimator to allow for dynamic processes. They present Monte Carlo results that demonstrate that when counts are serially correlated, the use of dynamic Poisson regression models produces less-biased and more efficient estimates than a standard Poisson regression estimator. Their Poisson exponentially weighted moving average (PEWMA) and Poisson autoregression (PAR(p)) models produce estimates of dynamic count processes that are less biased to those from (a) a log-linear regression, (b) a Poisson regression, or (c) a negative binomial regression. They also document that one cannot simply include lagged count terms in a Poisson regression to capture serial correlation effects (Brandt et al. 2000, 825, footnote 2).

Static multivariate count models have been proposed. King (1989a) offers a bivariate count model (the seemingly unrelated Poisson regression model, SUPREME) that builds on the logic of the Gaussian seemingly unrelated regression model. King's model estimates a multivariate count model and tests whether there is contemporaneous correlation among the counts; however, it only allows positive correlations among the count series. Karlis and Meligkotsidou (2005) present a different multivariate Poisson model that allows different *positive* correlations and covariates across multiple counts. However, they admit that their model is more constrained than those described below because it does not allow for negative correlations and cannot account for the common overdispersion seen in count data (Karlis and Meligkotsidou 2005, 256). So, these models offer little in terms of discovering potential endogenous effects because they only allow for positive correlations among the event series.

A more fruitful approach is that of Aitchison and Ho (1989), who develop a (conditionally) Poisson log-normal model that allows for a mixture of multiple Poisson count models with a multivariate log-normal distribution for the latent means and covariance of the process. There are two advantages of this

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press

process: (1) it has an unrestricted (i.e., positive or negative) covariance structure across the counts, and (2) it admits overdispersion. However, they document that the maximum likelihood estimator for this model is computationally expensive.

Chib and Winkelmann (2001) extend the Aitchison and Ho (1989) model using a Bayesian specification and propose a quick robust estimation method. Chib and Winkelmann (2001) establish the superiority of their multivariate count model over earlier, separate, or independent univariate count models for observed multivariate data (cf. Deb and Trivedi 1997; Munkin and Trivedi 1999). Thus, when the counts are multivariate jointly modeling, the counts lead to less-biased and more-efficient estimations across the system of count variables.

Combining the univariate dynamic Poisson models and multivariate count models is the goal of this paper. The approach of Brandt et al. (2000) and Brandt and Williams (2001) is not easily extended to a multivariate model. The reason why is seen in Ord, Fernandes, and Harvey (1993), who develop a multivariate count model where the predictive distribution is an exponentially weighted moving average of the latent means of the counts (an extension of the PEWMA of Brandt et al. 2000). The latent variable for the counts follows a special multivariate generalization of the univariate gamma latent variable used in the construction of the PAR(p) and PEWMA models. Ord, Fernandes, and Harvey's (1993) model, however, suffers from two assumptions or complications that make it a less than ideal candidate. First, the correlations among the counts are restricted to lie between $\left(\frac{-1}{m-1}, 1\right)$, where m is the dimension of the multiple counts. Second, the dynamic process is based on an exponentially weighted discounting that does not admit cross-equation effects like those necessary to identify complementary or substitution effects across the count time series. This complicates inference about dynamic endogenous effects among the count time series. This means that multivariate generalizations of the PAR(p) or PEWMA models are difficult to interpret and are unnecessarily complex.

Another recent effort to build a dynamic multivariate count model is that of Heinen and Rengifo (2007). They propose a multivariate autoregressive conditional double Poisson model where the contemporaneous correlations depend on copulas. In this model, the counts follow a double Poisson rather than standard Poisson process, meaning that the observed Poisson process is itself the result of another unobserved Poisson process. In many political science applications, there is no *a priori* reason to believe that the event counts will arise from such a model.[1]

Finally, some authors use Gaussian vector autoregression models for serially correlated multivariate counts. For example, Hanson and Schmidt (2011) employ a Gaussian VAR model to look at the relationship between the weekly number of Coalition casualties and Coalition in Iraq from late 2004 until 2006. However, multivariate time-series models, like the vector autoregression (and related) methods discussed in Brandt and Freeman (2006) and Brandt and Williams (2007), are based on a Gaussian error process. For dynamic event count data, such an assumption is not tenable and produces biased estimates (per the findings in King 1988 and Brandt and Williams 2001). So, although these methods will work for count series that are approximately normal, they may fail when applied to the dynamic events like those that we are interested in here.

Combining all these past results, we can summarize in Table 1 the existing knowledge about how to model counts that may be static, dynamic, multivariate, or univariate. In Table 1, the leftmost specification columns describe the assumptions or nature of the data (univariate or multivariate) and their dynamics (static or dynamic). The remaining columns are various estimators that have been proposed to address the count or dynamic aspects of the kind of data considered here. Cell entries document properties of the estimates for a given estimator (column) for different assumptions (rows). All these alternatives produce potentially biased or inefficient estimates based on *known results* in the existing literature.[2] Using any

[1]Such a process actually makes sense in the Heinen and Rengifo (2007) analysis. Their dependent variable is the number of New York Stock Exchange trades for five retail department stores. Here, one can think about discrete numbers of trades generated by the translation of dollar amounts into discrete shares (thus generating the number of trades). Such logic does not seem to apply to terrorist targeting, where we are thinking about how marginal utility and resource constraints affect the intensity of attacks or arrival rate of the number of new attacks.

[2]The cumulative conclusion of the existing literature, based on analytic and Monte Carlo studies, is that if counts are dynamic or multivariate, estimators that include these properties will be superior to those that do not include them. This means that a Monte Carlo analysis is not necessary since the dynamic multivariate nature of the data in our model is not going to be captured well

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press

Table 1 Model assumptions versus estimators

Specification		*Estimators and properties*						
Multi- v. univariate	*Dynamic v. Static*	*(L)OLS*	*Poisson*	*PAR(p)/PEWMA*	*SUPREME*	*Multivariate count*	*Linear VAR*	*BaP-VAR*
Univariate	Static	Biased	Unbiased	Unbiased	Unbiased	Unbiased	NA	NA
Multivariate	Static	Biased	Inefficient	Inefficient	Unbiased	Unbiased, efficient	Unbiased	Unbiased
Univariate	Dynamic	Biased	Biased	Unbiased			Biased	Unbiased
Multivariate	Dynamic	Biased	Biased	Inefficient	Biased	Biased	Biased	Unbiased

Note. Left columns give assumptions or properties of data. Estimator columns document bias or inefficiency results. See text for explanations.

of these extant models for multivariate counts with dynamics can lead possibly to biased estimates, thus providing the rationale for the general model proposed here.

The estimator proposed here, the BaP-VAR, reduces in many cases to the alternative univariate and static or dynamic models in Table 1. For instance, if the counts are uncorrelated, but dynamic, the model reduces to a series of PAR(p)-like models for each count time series. If the counts are correlated contemporaneously, *but not dynamically or over lags* (so that all the autoregressive coefficients are zero), the model reduces to the multivariate count model of Chib and Winkelmann (2001).

3 A BaP-VAR Model

Our basic multivariate count model extends that proposed by Chib and Winkelmann (2001). Suppose that we have a series of m grouped counts for time periods $t = 1, \ldots, T$, denoted by y_{tj}, for $j = 1, 2, \ldots, m$. The observed data are assumed to have marginal Poisson distributions of the form

$$y_{tj}|b_t, \beta_j \sim \text{Poisson}(\mu_{tj}), \tag{1}$$

$$\mu_{tj} = \exp(x_{tj}\beta_j + b_{tj}), \quad \forall\, j \leqslant m, \quad t \leqslant T, \tag{2}$$

where y_{tj} is the tth observation of the jth count variable, $b_t = (b_{t1}, \ldots b_{tm})$ is a set of observation and equation-specific latent effects (which capture the contemporaneous correlations among the series), μ_{tj} is the mean of the (conditionally) independent Poisson distributions for observation t in equation j, and x_{tj} and β_j are the regression components.

The covariance of the latent effects terms describes the contemporaneous relationships among the m count time series. Let the $m \times m$ matrix D be the covariance of these latent effects. The latent effects are distributed normally or

$$b_t|D \sim N_m(0, D), \quad \forall\, t. \tag{3}$$

Note that these effects are unrestricted, so that the covariance matrix D will be positive definite and admit positive or negative covariances among the counts.

Let $w_{tj} = \exp(b_{tj})$ (i.e., the natural log of the w terms are the b terms). Then the latent variables for the random effects in this model are log-normally distributed since

$$\exp(b_{tj}) \sim \exp(N(0, D)) \rightarrow \log(w_{tj}) = b_{tj} \sim \text{LN}(\mu^*, \Sigma), \tag{4}$$

where $\text{LN}(\mu^*, \Sigma)$ is a log-normal distribution with mean μ^* and covariance Σ. Chib and Winkelmann (2001, 429) outline the properties of the model; namely, it has an unrestricted, flexible latent correlation structure, and it can be used for longitudinal data.[3]

by the other estimators. Our two specific examples highlight dynamic properties, not captured in previous analyses that do not incorporate our innovation.

[3]The mapping from the latent mean and variance in equation (4) to the conditional mean of the response variables is in the appendix. Chib and Winkelmann (2001) propose a number of extensions to the basic model. These are different specifications of b_{tj}: using a multivariate student-t distribution for the latent effect or parameterizing them with covariates (cf. Karlis and Meligkotsidou 2005). Additionally, Congdon (2005) presents the Chib and Winkelmann (2001) model, replicates the bivariate Poisson model of Munkin and Trivedi (1999), and presents a Bayesian factor model for the latent effects as an extension.

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press

The likelihood function for one observation is found from the convolution of equations (1) and (3):

$$\Pr(y_t|\beta, D) = \int \prod_{j=1}^{m} f(y_{tj}|\beta_j, b_{tj})\phi_m(b_t, 0, D)\partial b_t. \tag{5}$$

Here, we have the likelihood for the m counts in observation t, conditional on the regression coefficients and the latent effects. The first density $f(\cdot)$ is the Poisson mass function from equation (1). The second component of the likelihood function is the m-dimensional multivariate normal density for the latent effects.[4]

A natural extension to allow a dynamic process for the conditional mean in equation (2) is to add autoregressive terms. Such an approach would allow for modeling time-series dynamics and contemporaneous correlations among the counts via the D matrix. This adds a data augmentation step to the model to account for the possible dynamic effects and builds on the logic of the conditional linear process of the PAR(p) model (Grunwald et al. 2000; Brandt and Williams 2001).

A conditionally linear approach autoregression can be implemented by writing the latent means μ_{tj} as a vector autoregression (VAR) process. This is derived from the (column) vectorized latent means for observation t for equations $j \leqslant m$, denoted by $\mu_t = (\mu_{t1}, \mu_{t2}, \ldots, \mu_{tm})'$. Then an observation-driven approach for the dynamics writes this latent mean as a function of the past observed counts and the predictors, or see Brandt and Williams (2001) and Congdon (2005):

$$\mu_t = \begin{pmatrix} \mu_{t1} \\ \mu_{t2} \\ \vdots \\ \mu_{tm} \end{pmatrix} = \begin{pmatrix} a_{11}y_{t-1,1} + a_{12}y_{t-1,2} + \cdots + a_{1m}y_{t-1,m} \\ a_{21}y_{t-1,1} + a_{22}y_{t-1,2} + \cdots + a_{2m}y_{t-1,m} \\ \vdots \\ a_{m1}y_{t-1,1} + a_{m2}y_{t-1,2} + \cdots + a_{mm}y_{t-1,m} \end{pmatrix} + \begin{pmatrix} \exp(x_{t1}\beta_1 + b_{1t}) \\ \exp(x_{t2}\beta_2 + b_{2t}) \\ \vdots \\ \exp(x_{tm}\beta_m + b_{mt}) \end{pmatrix} \tag{6a}$$

$$= A \begin{pmatrix} y_{t-1,1} \\ y_{t-1,2} \\ \vdots \\ y_{t-1,m} \end{pmatrix} + \exp(X_t\beta + b_t), \tag{6b}$$

where A is the $m \times m$ matrix of the VAR(1) coefficients, X_t is an $m \times k$ vectorized set of predictors, β is a $k \times 1$ matrix of regressors, and b_t is the $m \times 1$ vector of random effects at time t.[5] This is a multivariate analog of Brandt and Williams (2001), but using a different nonconjugate density for the random effects.

With respect to the Markov chain Monte Carlo (MCMC) implementation proposed in Chib and Winkelmann (2001), this generates no problems because we are only adjusting the latent mean μ_{tj} by the observable VAR(1) process. Although the conditional posterior of the VAR(1) parameters is nonstandard, this presents no difficulties via a (block) Metropolis–Hastings sampler.

Several facts should be noted about the specification of the BaP-VAR in equation (6b). First, when $A = 0$ (for all elements), the model reduces to that of equation (2), which is the same latent mean process from Chib and Winkelmann (2001) since then there is no dynamic process. So, if the data of interest have unknown dynamics, one can directly evaluate this fact using the estimates from the BaP-VAR(p) model to determine whether to use this model or the nondynamic version. Second, if the estimates of D, the latent contemporaneous correlations among the counts, are a diagonal matrix, but $A \neq 0$, then this model reduces to a series of PAR(p)-like estimations (cf. Brandt and Williams 2001).

3.1 *Properties and Interpretation*

The benefit of this approach is that we can use standard VAR results to derive the dynamic properties (i.e., impulse responses and DFEV) for the correlated count model. For the BaP-VAR(1) matrix A, the

[4]Note that this likelihood lacks a closed form for all but the simplest cases of D. The obvious case is when D is a diagonal matrix so that we have m independent Poisson processes with heterogeneity. This is why beyond the $m = 2$ case, most implementations will need some numerical approximation (see Aitchison and Ho 1989) or a Monte Carlo implementation. Given our need for inference, this is a case where Bayesian posterior simulation is preferred on both computational and substantive grounds.

[5]The extension to a VAR(p) process is obvious.

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press

regressors, and latent terms, the vectorized version latent means are

$$\mu_t = Ay_{t-1} + \exp(X_t\beta + b_t), \tag{7}$$

where $\exp(X_t\beta + b_t)$ is a vector of the elements x_{tj}, β_j, and b_{tj} as in equation (6b).

The impulse response function is derived as in Lütkepohl (2005, 51–53) to trace out the effect of a shock to one equation in the system. For a two-equation case, this would be

$$\begin{bmatrix}\mu_{1t}\\ \mu_{2t}\end{bmatrix} = \begin{bmatrix}a_{11}\, a_{12}\\ a_{21}\, a_{22}\end{bmatrix}\begin{bmatrix}y_{1,t-1}\\ y_{2,t-1}\end{bmatrix} + \begin{bmatrix}\exp(x_{1t}\beta_1)w_{1t}\\ \exp(x_{2t}\beta_2)w_{2t}\end{bmatrix}, \tag{8}$$

where the first subscript denotes the equation and $w_{jt} = \exp(b_{jt})$. The change in μ_{1t} for a surprise change in w_{1t} is recursively traced from $t = 0$ to $t = 2$ as follows:

$$\begin{bmatrix}\mu_{10}\\ \mu_{20}\end{bmatrix} = \begin{bmatrix}\exp(x_{10}\beta_1)\cdot 1\\ \exp(x_{20}\beta_2)\cdot 0\end{bmatrix} = \begin{bmatrix}\exp(x_{10}\beta_1)\\ 0\end{bmatrix} \tag{9a}$$

$$\begin{bmatrix}\mu_{11}\\ \mu_{21}\end{bmatrix} = \begin{bmatrix}a_{11}\, a_{12}\\ a_{21}\, a_{22}\end{bmatrix}\begin{bmatrix}y_{1,t-1}\\ y_{2,t-1}\end{bmatrix} \tag{9b}$$

$$= \begin{bmatrix}a_{11}\, a_{12}\\ a_{21}\, a_{22}\end{bmatrix}\begin{bmatrix}\exp(x_{10}\beta_1)\\ 0\end{bmatrix} = \begin{bmatrix}a_{11}\exp(x_{10}\beta_1)\\ a_{21}\exp(x_{10}\beta_1)\end{bmatrix} \tag{9c}$$

$$\begin{bmatrix}\mu_{12}\\ \mu_{22}\end{bmatrix} = A\mu_1 = A^2\mu_0. \tag{9d}$$

This is the same iterative derivation seen in standard VAR results, accounting for the initial scaling effect of the unit shock. Here, we assume that there is only a unit shock in the initial period that is traced out and that the initial effects prior to $t = 0$ are in equilibrium. This allows us to replace the values of the lagged y_t terms with their lagged predictors from the previous steps. The result is a prediction of the path of the one-step-ahead forecast errors or the impulse response. The other assumption here is that we are not tracing out the mean of the system but rather the dynamics around the mean for a unit shock. So, the $\exp(x_t\beta)$ terms only enter in the initial step to scale the shock. The generic vector version of this impulse response result is shown in an appendix. The derivation there also allows for one-standard-deviation shocks that may be correlated across the equations based on the reduced-form error covariance of the random effects, Σ. Note that when $A = 0$, the marginal effects of any covariates are going to be the same as those from the Chib and Winkelmann (2001) model and these dynamic multipliers are not relevant, and the impulse responses will trail off immediately at lag zero.

As in standard VAR models, the decomposition of the forecast variance for this model can be derived. Let C_i be the ith moving average coefficient matrix for the infinite expansion of the lag polynomial (see equation (A6c) of the appendix), $\exp(X_t\beta + b_t) = e_t$ be the fixed effect and the random error vector, and $\hat{y}_{t+s}$ be the prediction at period $t + s$. The vector moving average (VMA) representation of the process in period $t + s$ is then

$$y_{t+s} - \hat{y}_{t+s} = e_{t+s} + C_1e_{t+s-1} + C_2e_{t+s-2} + \cdots + C_{s-1}e_{t+1}. \tag{10}$$

This VMA representation of the BaP-VAR differs only from a standard Gaussian VAR in the inclusion of the $\exp(x_t\beta)$ terms in the process. If these are not time varying, then the derivation of the variance of the forecast errors, $y_{t+s} - \hat{y}_{t+s}$, follows the standard results for VAR processes (accounting for the $\exp(x_t\beta)$ terms). We can then use the standard results in Brandt and Williams (2007, 45–48) and Hamilton (1994) for the innovation accounting or DFEVs. These trace out the amount of variation in each series due to the changes in the other series in the systems of counts. This is computed using the variance of the counts in period $t + s$ or

$$V(y_{t+s} - \hat{y}_{t+s}) = E[(y_{t+s} - \hat{y}_{t+s})'(y_{t+s} - \hat{y}_{t+s})] \tag{11a}$$

$$= \Sigma + C_1\Sigma C_1' + C_2\Sigma C_2' + \cdots + C_{s-1}\Sigma C_{s-1}'. \tag{11b}$$

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press

The effects of the change in each of the count series explains how the covariance in equation (11b) changes over time (changes in s), which is captured by looking at different one-standard-deviation changes in the e_t elements across the variance of equation (11b). In the sequel, we compute these changes in the covariance of the forecast errors over time and use them to look at the relative proportion of the variance of the series that is explained by the changes in the other series. This decomposition provides insights about how each of the series can help predict or explain the other.

3.2 *Priors and Estimation*

A natural prior for the (β, D^{-1}) parameters is a normal-Wishart distribution with independent priors for the regression coefficients $\beta \sim N(\beta_0, B_0^{-1})$ and the latent log-normal variances $D^{-1} \sim \text{Wishart}(R_0, \nu_0)$. This parallels the analysis in Chib and Winkelmann (2001). The remaining parameters are those for the autoregressive matrices A.

Let $\pi(A_0, \Psi) = \phi(A_0|\Psi) \cdot f_W(\Psi^{-1}, T)$ be the normal-Wishart prior for the BaP-VAR autoregressive coefficients in $A_1, \ldots, A_p$, the covariance, Ψ, be the matrix of priors for the autoregressive coefficients, and T be the sample size or degrees of freedom. Note that in an m equation model with p lags, this implies setting prior means and covariances on $m^2p + m$ parameters. We use a multivariate normal prior for the autoregressive component and set the prior means and covariance following (Sims and Zha 1998; see also Brandt and Freeman 2006). This prior mean for the autoregressive coefficients is centered on a random walk model where $A_1 = I$ and $A_2 = A_3 = \cdots = A_p = 0$ for the VAR regression coefficients. The prior variance for the autoregressive coefficient for variable j in equation i for the ℓ lag is

$$\psi_{\ell,j,i} = \left(\frac{\lambda_0\lambda_1}{\sigma_j \ell^{\lambda_3}}\right), \tag{12}$$

where λ_0, λ_1, and λ_3 are hyperparameters in the Sims–Zha prior. The first of these hyperparameters controls the overall scale of the prior coefficients, the second scales the variance around the A_1 coefficients, and the third allows the variances of the coefficients at higher lags to shrink toward 0. The σ_j term is the prior standard deviation of the count variable in equation j from the sample estimate of D for equation j.[6] This prior centers beliefs about the dynamics on low-order frequencies and smoothes the coefficients of higher-order lags toward 0.

The posterior is then proportionate to

$$\underbrace{\phi(A|\Psi)\, f_{\text{W}}(\Psi^{-1}, T)\phi(\beta|\beta_0, B_0^{-1})\, f_{\text{W}}(D^{-1}|\nu_0, R_0)}_{\text{Prior}} \underbrace{\prod_{j=1}^{m} f(y_{ij}|A, \beta_j, b_{tj})\phi_m(b_t, 0, D)}_{\text{Likelihood}}, \tag{13}$$

where $\phi(\cdot)$ is a multivariate normal density and f_{W} is the Wishart density. In equation (13), the first term is the Sims–Zha prior for the autoregressive process, the second term is the prior for the regression parameters, the third term denotes the prior for the latent covariance, and the fourth term is the likelihood. The MCMC algorithm for simulating this posterior distribution is given in the appendix.

3.3 *Lag Length and Model Selection*

The previous model specification discussions assume that the lag length for the process (p) is known. Since one does not generally know the dynamic specification of the process, this needs to be tested as part of the estimation of the models. Given the large number of parameters in these models, the Bayesian model fit and selection criteria need to penalize overfitting and to compute easily as part of the posterior sampling process. Our candidate measure is the *Deviance Information Criteria* (DIC), detailed in Spiegelhalter et al. (2002) and Plummer (2008). The DIC measures the expected variation of the model plus a penalty for the model complexity in terms of the effective number of parameters. Like the Akaike Information Criteria

[6] As a rule of thumb, the prior standard deviation for equation j, σ_j should be smaller than the observed sample value. This is because the sample value will be less than the long-run belief about the standard deviation of the series for the jth equation.

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press

commonly used in comparing maximum likelihood models, smaller values are evidence of better fit. We will select the lag length of our BaP-VAR models based on the value of p that minimizes the DIC statistic.[7]

4 Example 1: Superpower Rivalry

The first example employing the BaP-VAR model is an analysis of the number of MIDs between the United States and the Soviet Union. King (1989b, 201–207) analyzes this rivalry using a multivariate count model, the SUPREME. There are three reasons why the BaP-VAR(p) model is recommended for this application:

1. There is significant well-known dynamics in the United States and Soviet rivalry and defense expenditures in the period of analysis (annual, 1951–1978) (e.g., Williams and McGinnis 1988, 1992; McGinnis and Williams 1989).
2. Aggregate MIDs are serially correlated (Brandt et al. 2000).
3. King's (1989b) model requires that the structure of the common drivers is a linear, equally weighted additive latent variable contribution to the number of MIDs for each superpower actor.

King's (1989b) model misses the potential dynamic contributions of the BaP-VAR(p) model. This model can estimate the more complex and potentially lagged relationships rather than just additive contemporaneous effects of the covariates on the outcome variables. King's analysis estimates the joint or tit-for-tat behavior between the United States and the Soviet Union with a separate equation. He assumes that there are three latent variables for United States to Soviet (US2SU) actions that are initiated domestically, Soviet to United States (SU2US) actions that are initiated domestically, and a third joint (JOINT) set of tit-for-tat behaviors that are the unobserved sum of the two series. This model estimates three sets of parameters for the effects of United States and Soviet military spending and the party of the president on each series.

This example shows how to specify, estimate, and interpret a simple BaP-VAR(p) model for these rivalry data (with the covariates and their interpretation). We then compare the findings of the BaP-VAR model to both SUPREME and Gaussian VAR regressions. Our models use only the two initial series and identify the tit-for-tat behavior via impulse responses and the contemporaneous effects of the two count series on one another, as in Brandt, Colaresi, and Freeman (2008).

4.1 *Data and Model Specification*

The dependent variables in King's (1989b) analysis are the number of United States and Soviet MIDs annually from 1951 to 1978, recorded as counts (see Azar 1982). Figure 1 shows the data series being modeled, their ACFs, and histograms. Given the small sample size ($N = 28$) and the degree of autocorrelation seen in the ACF plots, there is weak evidence for some serial correlation in these data.

There are three covariates (in each equation) in King's (1989b) analysis: military expenditures by the United States, military expenditures by the Soviets (both in constant 1970 $billions), and the party of the U.S. president (1 = Democrat, 0 = Republican). The BaP-VAR(p) models include each of these as exogenous covariates (without any *a priori* testing): Here, we just extend the earlier contribution to see what happens when dynamic effects are added to the model in a less restrictive correlation structure across the specification. King (1989b) averaged these covariates to find the joint effects that can be separately (weightedly) identified in the BaP-VAR model.

[7] Sims, Waggoner, and Zha (2008) document how hard it is to generate model fit statistics like log marginal data densities for Gaussian Bayesian VAR and Markov-switching Bayesian VAR models. Given these difficulties and the complexities of the BaP-VAR, our use of the DIC seems advisable.

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press

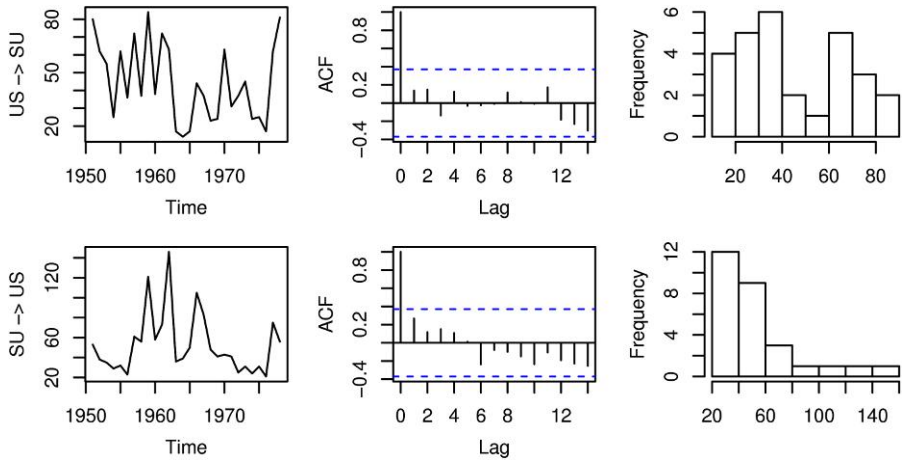


Fig. 1 Superpower rivalry data and diagnostic plots.

The prior distribution for the parameters in this example are as follows (per our recommendation from Sims and Zha 1998; Brandt and Freeman 2006):

$$A_1 = N\left[\begin{pmatrix}1&0\\0&1\end{pmatrix}, \begin{pmatrix}1&0\\0&1\end{pmatrix}\right], \tag{14}$$

$$A_i = N\left[\begin{pmatrix}0&0\\0&0\end{pmatrix}, \begin{pmatrix}\frac{1}{p}&0\\0&\frac{1}{p}\end{pmatrix}\right], \quad i = 2, \ldots, p \tag{15}$$

$$\beta = N(0, \sigma_\beta), \tag{16}$$

where β is the vector of the intercept and the exogenous variable prior coefficients. The prior sets these to have a mean of 0 and a variance of $\sigma_\beta = 3.33$ for the intercepts and 1.67 for the other exogenous covariates. Consistent with the application of the Sims–Zha prior in Brandt and Freeman (2006), this allows for sufficient variation to be explained by the dynamics of the process and the exogenous variables at the same time.[8]

4.2 *Results*

We fit the BaP-VAR(p) models for $p = 1, 2, 3$ lags to the superpower rivalry data. Three covariates—military spending by each rival and the party of the U.S. president—are also included in the BaP-VAR model. The burn-in sample is 100,000 draws with a final sample of 250,000 draws reported below. The best-fitting model is the one with a single lag of the counts of the MIDs for each belligerent, based on the DIC values. Table 2 reports the posterior estimates for the coefficients from the BaP-VAR(1) model, King's (1989b) SUPREME model, and a Gaussian VAR for the logged counts.

The estimates in Table 2 support several conclusions and comparisons across the models. First, there are dynamic effects of the series on each other. The estimated residual correlation across the two-equation BaP-VAR(1) model is 0.77. King's (1989b) model assumes that this is zero, which is a serious form of misspecification. It is also substantively important since this correlation is evidence of positive contemporaneous reciprocity between the superpowers: Each unanticipated MID leads to more MIDs.

Second, there is strong evidence for the autoregression in the BaP-VAR(1) since the autoregressive coefficients for the lag of US2SU actions affect other series positively. Note that the Gaussian VAR results show no autoregressive parameters that differ from zero, reflecting the downward bias documented in using both linear regression-based and nondynamic count models (King 1988; Brandt and Williams 2001).

[8]Given the small sample size, this choice of the prior is more heuristic than equation (12).

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press

Table 2 Estimates of BaP-VAR(1), SUPREME, and log VAR(1) models for superpower rivalry

	BaP-VAR(1)		*SUPREME*			*VAR(1)*	
	US2SU	*SU2US*	*US2SU*	*SU2US*	*Joint*	*US2SU*	*SU2US*
$US2SU_{t-1}$	0.46	0.33				0.42	−0.13
	(0.22, 0.66)	(0.05, 0.56)				(−0.04, 0.89)	(−0.26, 0.67)
$SU2US_{t-1}$	−0.10	0.08				−0.37	0.20
	(−0.23, 0.04)	(−0.10, 0.26)				(−0.96, 0.21)	(−0.26, 0.67)
Intercept	0.83	1.55	2.67	2.59	8.58	3.67	4.24
	(−1.90, 3.28)	(−1.02, 3.84)	(0.12)	(0.17)	(0.75)	(1.20, 6.13)	(2.27, 6.20)
U.S. Military $	0.02	0.03		0.01		−0.001	−0.005
	(−0.02, 0.06)	(−0.01, 0.07)		(0.002)		(−0.027, 0.025)	(−0.026, 0.016)
S.U. Military $	0.01	−0.01	0.010			−0.003	−0.010
	(−0.01, 0.04)	(−0.03, 0.01)	(0.002)			(−0.015, 0.010)	(−0.020, 0.001)
Average Military					−0.103		
					(0.016)		
U.S. President	0.18	0.78	0.438	0.794	−0.674	0.21	0.41
	(−0.74, 1.08)	(0.04, 1.62)	(0.055)	(0.062)	(0.202)	(−0.260, 0.676)	(0.040, 0.785)

Note. 95% credible intervals for the Bayesian model in parentheses; standard errors for the SUPREME model; 95% confidence intervals for VAR.

The Gaussian VAR estimates reflect the expected downwardly biased estimates when compared to the count-based estimators.

The presence of these autoregressive dynamics (and the absence of evidence that the SU2US lag coefficients differ from zero) indicates that the U.S. MIDs Granger caused Soviet MIDs. This pattern is seen clearly in an examination of the impulse response functions for the BaP-VAR model. Figure 2 presents these impulse responses for one-standard-deviation shocks to each equation. The equations are ordered with the US2SU shock hitting the system before the SU2US shock (consistent with the results in Table 2). The median responses and their 68% (approximately one standard deviation on each side) credible regions are plotted as well.

These responses show that most of the dynamic reactions of the series are due to innovations in the US2SU series. The effects of Soviet-initiated MIDs are small, and the effects of U.S. MIDs on the Soviets come with a lag. Thus, it is the change in U.S. MIDs that generated most of the reactions in this system. Further, King reports that based on his model, about 6.9 MIDs are due to the tit-for-tat behavior among the belligerents. This is nearly identical to the observed impulse responses in Fig. 2.[9]

Third, King's analysis assumes that the effects of the covariates are static rather than dynamic. Based on the estimates reported in Table 2, the effects of the military expenditures are very similar in both models. The effects of the party of the U.S. president are distinct: Having a Democrat in the White House leads to more MIDs directed from the Soviets to the United States. This estimate is nearly identical in the SUPREME model. The models differ, however, when looking at the partisan effects of the control of the White House on the other equation. In the BaP-VAR model, there is no evidence that Democratic control of the White House leads to a change in MIDs, in contrast to the results reported in King's (1989b). For the Gaussian VAR, there are no effects for the military spending covariates in either equation and the U.S. president party variable has half of the reported effect seen in the BaP-VAR results for the SU2US equation. Figure 3 reports the dynamic impact multipliers for one-unit changes in each of the covariates along with the 90% confidence intervals for each effect in the BaP-VAR. The plot demonstrates that each additional billion dollars in United States and Soviet military expenditures leads to an additional MID and that the effect decays over 6 years. Having a Democrat in the White House leads to one additional U.S. MID toward the Soviets and two or more additional MIDs from the Soviets toward the United States.

[9]The replication materials for the paper include the Gaussian VAR-based impulse responses. These Gaussian estimates are biased downward from those reported for the BaP-VAR.

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press

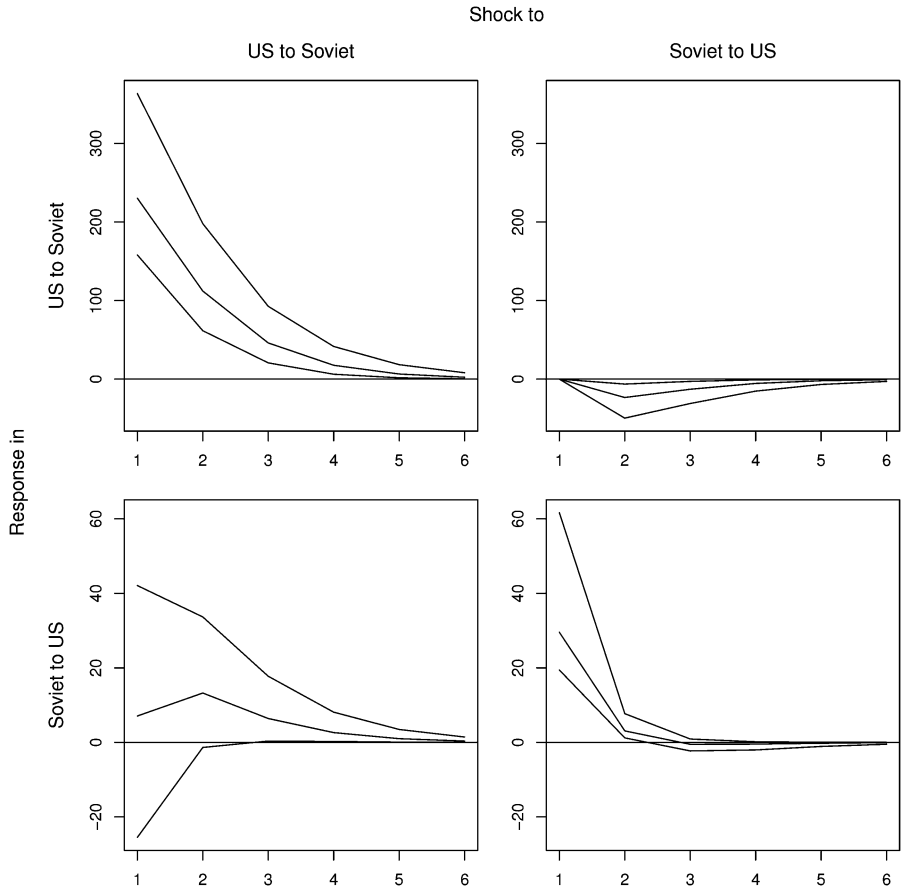


Fig. 2 Impulse responses for the superpower rivalry. 68% error bands are shown with the outer solid lines.

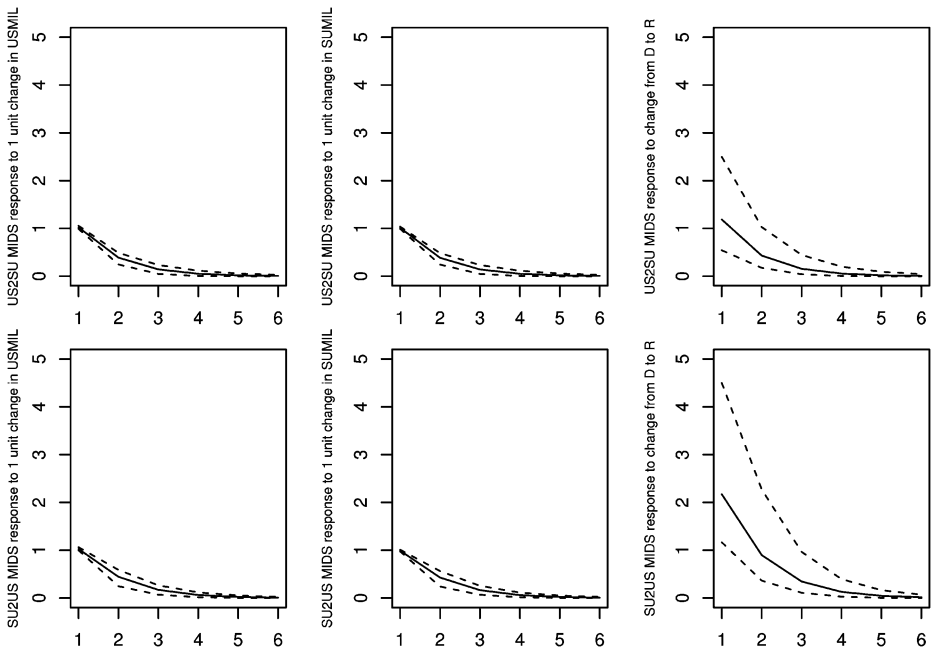


Fig. 3 Impact multipliers for the covariates in the superpower rivalry example.

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press

The military spending results are consistent with King's finding that it is not a deterrent to conflict. And, contrary to King, the partisan effects are in the direction of more MIDs and not tit-for-tat behavior. Here, this tit-for-tat behavior is seen in the dynamics of the impulse responses, not in the covariates.

5 Example 2: Transnational Terrorist Targeting

The second application of the BaP-VAR model is to transnational terrorism targeting decisions. Since 1968, transnational terrorism has posed a worldwide security concern, punctuated by the four skyjackings on 11 September 2001 (henceforth, 9/11) that killed almost 3000 people and had short-term global financial reverberations (Chen and Siems 2004). Governments' efforts to harden potential targets have induced terrorists to redistribute their attack efforts among targets, attack modes, and geographic locations based on their perceived marginal gain per dollar spent (Enders and Sandler 1993, 2006a). Some terrorist choices react opposite to one another, whereas other choices move in tandem. The former are substitutes, whereas the latter are complements.[10] Here, the BaP-VAR model is used to investigate transnational terrorist targeting choices for 1968–2008 and to discern the patterns of complementarity and substitution among four target classes—officials, military, business, and private parties.

5.1 *Transnational Terrorism Data and Samples*

In this analysis, terrorism is the premeditated use or threat to use violence by individuals or subnational groups in order to obtain a political or social objective through the intimidation of a large audience beyond that of the immediate victims. Here, we examine transnational terrorist incidents that through their perpetrators, victims, institutions, governments, targets, or venues involve two or more countries. Transnational terrorism remains a key security concern for the global community that draws large public expenditures. An understanding of the changing dynamics of transnational terrorist targeting can inform policy, especially when target interdependencies are taken into account.

Brandt and Sandler (2010) apply univariate Bayesian Poisson changepoint regression models to identify how terrorists adjusted their target choices over time in response to target hardening and the changing orientation of terrorist preferences. Their analysis endogenously locates regimes for the four separate target types; however, this earlier analysis cannot investigate the dynamic multivariate interdependency among the target series. The regimes, generally, fall into five time intervals. Regime 1 starts in January 1968 (denoted by 1968:1) and concludes in 1973:1 when metal detectors were first introduced in airports. Next, Regime 2 represents the era of technological and other barriers and runs from 1973:2 to 1979:12. Regime 3 is that of state sponsorship and encompasses 1980:1–1989:11. It ends just at the start of the fall of communism, which marked a decline but not the elimination of state sponsorship. In Regime 4, from 1989:12 to 2001:9, the fundamentalist terrorists (e.g., al-Qaida, Jemaah Islamiyah, and others) became the dominant transnational terrorist influence, eclipsing the left-wing terrorists (e.g., the Italian Red Brigades and many others) whose numbers dwindled with the collapse of communism and effective police actions (Alexander and Pluchinsky 1992; Hoffman 2006). Finally, Regime 5 encompasses the post-9/11 era, a period during which the fundamentalist terrorists were under attack (2001:10–2008:12). In response, transnational terrorists engaged in fewer attacks with greater carnage and sought softer targets (i.e., private parties; see Enders and Sandler 2005; Brandt and Sandler 2010).[11]

The most comprehensive transnational terrorist event data are *International Terrorism: Attributes of Terrorist Events* (ITERATE). Our analysis relies on ITERATE, which now runs from 1968 to 2008 and includes 13,179 incidents (Mickolus et al. 2009). For our purposes, ITERATE records the incident date

[10]For attack modes, kidnappings are viewed as substitutes for skyjackings; that is, metal detectors in airports resulted in fewer skyjackings and many more kidnappings (Enders and Sandler 1993, 2006b). By contrast, skyjackings and other hostage (nonkidnapping) events (i.e., barricade and hostage taking and nonaerial hijackings) are complements (Brandt and Sandler 2009). That is, hostage-taking events at known locations tend to rise and fall together.

[11]Neither of these papers find a regime change at 9/11. We split the sample here to confirm that the patterns are the same pre- and post-9/11. It is important to remember that these are simplified descriptions of each regime.

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press

Table 3 DIC statistics for each subsample and BaP-VAR(p) lag length specification

Sample	*p*	*DIC*
1968:1–1973:1	1	934
	2	941
	3	957
1973:2–1979:12	1	1488
	2	1500
	3	1504
1980:1:2–1989:11	1	2268
	2	2269
	3	2280
1989:12–2001:9	1	2248
	2	2271
	3	2287
2001:10–2008:12	1	1198
	2	1208
	3	1213

with a unique marker to distinguish multiple incidents on the same date.[12] We combine incident types into three general event types that include hostage incidents, bombings, and assassinations.[13] Next, we construct four monthly time series for our four target types—officials, military, business, and private parties—that include these three aggregate modes of attack. The ITERATE variable, "type of immediate victim," allows us to classify our target types. Officials include host country and foreign officials along with foreign diplomats, whereas military includes domestic and foreign military victims. Business targets consist of corporate officials and employees. Finally, prominent opinion leaders and private individuals or civilians (e.g., tourists, missionaries, and students) constitute private parties. We do not include terrorist or unknown victims in the analysis.

5.2 *BaP-VAR Analysis of Targeting*

For each of the multiple time series of counts, we fit five models—one for each subsample. We also investigated models with one to three lags.[14] Based on the DIC statistics in Table 3, the best-fitting models for each subsample are those with only a single lag.

As a first step in the analysis, we need to ascertain whether the Poisson processes for each of the count series are independent (as in Brandt and Sandler 2010). This assessment of dependency across the equations can be done by computing the variance–covariance and cross-equation correlations (see equation (A4) in the appendix). Table 4 reports the median estimated 4×4 posterior mode of the Σ matrices, cross-equation covariances, correlations, and their 90% credible intervals in parentheses for each subsample. The row and column letterings correspond to the equation (O = Officials, M = Military, B = Business, P = Private Parties). The upper (lower) triangle reports the variances and covariances (correlations). Nondiagonal entries in bold have credible intervals that do not cover zero.

There are three important patterns in Table 4. First, the covariances of the latent effects or errors in the multivariate count model are *positive*, thus indicating that surprise increases in attacks on one

[12]For incidents without a fully specified date, we resort to a modal date. If, for example, the month is known but not the date, then we assign the incident to the 15th of the month. When only the year is given, we assign the incident to June 15 of that year. Very few of the terrorist incidents have an incomplete date.

[13]We exclude threats and hoaxes because these incidents do not respond to added security or changing terrorist preferences, the subject of this study. Given that there were only a handful of the sundry events such as exotic pollution, nuclear-related attacks, and sabotage, these are culled. Moreover, arms smuggling is excluded because the planned use for these weapons is not known; the weapons may have been intended for nonterrorist incidents such as robberies or insurgencies.

[14]A total of 15 BaP-VAR models are fitted, using the prior described earlier, across the samples and lag lengths. The results reported here are based on an MCMC posterior sample that consisted of 100,000 draws (after a burn-in of 30,000 draws). A second parallel chain was run for each model to check the convergence of the posterior samples. A comparison of the two chains for each model shows that they pass standard diagnostic tests for convergence. These results are included in the online replication materials.

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press

Table 4 Error covariances and correlations for the four series across the five subsamples

Period covariance/correlation		*O*	*M*	*B*	*P*
$\sum_{1968:1-1973:1}$	O	59.7	**10.2**	**21.7**	**9.0**
		(33.0, 152)	(3.5, 49.3)	(5.4, 113)	(0.33, 32.3)
	M	**0.38**	11.2	4.1	−0.23
		(0.16, 0.57)	(3.2, 310)	(−0.93, 44.6)	(−3.1, 5.1)
	B	**0.22**	0.09	157	7.25
		(0.05, 0.40)	(−0.01, 0.29)	(30.6, 6411)	(−4.0, 50.1)
	P	**0.19**	0.00	0.09	37.3
		(0.01, 0.38)	(−0.08, 0.18)	(−0.03, 0.30)	(16.6, 214)
$\sum_{1973:2-1979:12}$	O	27.2	2.15	**14.9**	**13.7**
		(18.1, 46.7)	(−0.26, 5.8)	(6.1, 27.5)	(5.1, 27.5)
	M	0.16	5.22	2.01	2.50
		(−0.02, 0.37)	(2.60, 39.6)	(−0.78, 5.70)	(−0.21, 8.65)
	B	**0.43**	0.12	43.0	**12.5**
		(0.20, 0.62)	(−0.04, 0.33)	(28.3, 82.2)	(3.4, 26.2)
	P	**0.46**	0.18	**0.23**	31.4
		(0.20, 0.67)	(−0.01, 0.41)	(0.10, 0.54)	(20.1, 66.5)
$\sum_{1980:1:2-1989:11}$	O	36.9	3.52	**16.0**	**19.1**
		(25.6, 54.4)	(−0.45, 8.3)	(8.51, 25.0)	(9.98, 30.1)
	M	0.19	9.24	3.09	2.04
		(−0.03, 0.34)	(6.38, 17.3)	(−0.06, 7.04)	(−1.51, 6.07)
	B	**0.55**	0.21	22.3	**11.7**
		(0.34, 0.71)	(−0.01, 0.42)	(15.5, 35.7)	(5.24, 19.9)
	P	**0.53**	0.11	**0.41**	35.5
		(0.32, 0.69)	(−0.08, 0.30)	(0.20, 0.60)	(24.9, 54.4)
$\sum_{1989:12-2001:9}$	O	23.1	**7.40**	**8.51**	**17.9**
		(12.0, 93.9)	(3.7, 25.5)	(3.90, 28.6)	(11.1, 35.5)
	M	**0.32**	22.2	2.20	**12.5**
		(0.14, 0.49)	(7.46, 268)	(−0.50, 7.87)	(7.20, 30.6)
	B	**0.20**	0.05	70.7	**16.8**
		(0.06, 0.38)	(−0.01, 0.17)	(18.8, 1445)	(9.00, 44.0)
	P	**0.51**	**0.37**	**0.28**	50.5
		(0.35, 0.65)	(0.18, 0.53)	(0.11, 0.43)	(35.2, 83.4)
$\sum_{2001:10-2008:12}$	O	3.30	**1.02**	**2.08**	**2.52**
		(2.36, 5.65)	(0.46, 1.85)	(0.57, 4.19)	(1.12, 4.71)
	M	**0.38**	2.03	**3.12**	**2.20**
		(0.16, 0.59)	(1.32, 6.49)	(1.60, 12.0)	(1.04, 4.39)
	B	**0.21**	**0.43**	24.3	**7.94**
		(0.04, 0.44)	(0.20, 0.65)	(9.93, 291)	(4.24, 20.5)
	P	**0.37**	**0.41**	**0.42**	13.7
		(0.17, 0.56)	(0.19, 0.59)	(0.20, 0.60)	(9.35, 24.7)

Notes. Subsample error covariances are denoted by the dates in the leftmost column. Median covariance (correlation) estimates and 90% credible intervals are presented in parentheses above (below) the diagonal. Entries in bold have intervals that do not cover zero.

type of target generally lead to more attacks on other types of targets. Second, the patterns of these positive correlations vary over time in response to changes in terrorist tastes, their resources, and defensive measures. These changes are consistent with our expectations about transnational terrorists' targeting decisions over time. In the 1968–1973 time period, attacks on officials drive international terrorism attacks on the other targets. In the 1973–1979 subsample, it is attacks on business and private parties that are most highly correlated with attacks on officials. Private party attacks are also correlated with business attacks as terrorists responded to enhanced security of officials by broadening their targets. During the period of state sponsorship (1980–1989), attacks on officials and business are positively correlated with those on private parties, where the correlations are somewhat higher than in the previous period. In the 1989–2001 period of Islamic fundamentalism, attacks on officials are positively correlated with all other target types

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press

and attacks on business and military targets are positively correlated with those on private parties. Finally, in the post-9/11 period, attacks on all targets are correlated with each other, albeit at lower levels than seen in the earlier periods. Third, all this is evidence of ever greater complementarity in terrorist targeting decisions: Targets are being attacked in more collective ways. Yet while these attacks are more correlated as Islamic-based attacks generate more collective carnage, the overall variation in the number of attacks is lowest in the most recent period (cf. the diagonals of Table 4 across the subsamples). These patterns change over time and will be investigated in more detail below using impulse responses and DFEVs.[15]

To investigate the dynamics of these correlations and determine the multiplier effects, we compute the impulse responses for each subsample. The posterior parameter draws are used to construct 100,000 sets of impulse responses. We assume a Wold causal ordering for the contemporaneous effects, where one-standard-deviation shocks first enter the officials equation, then the military, then the business, and finally the private parties equation. This is theoretically and empirically justified by the reaction times of the target types to attacks: Within a given month, one would expect that officials, then military, and then the business targets have the capacity to react to terrorist attacks. Collective action concerns and capacity inhibit a defensive response by private parties. This contemporaneous causal ordering is implemented with a Cholesky decomposition of the error covariance matrices for each draw in each subsample. For each median impulse response, a 68% percentile credible interval (approximately one standard deviation around the mode) is presented over 6 months (Sims and Zha 1999; Brandt and Freeman 2006).[16]

The impulse responses for the subsamples are plotted in Figs. 4–7. Each row in these figures corresponds to one of the equations in the model. The columns list that one-standard-deviation shock being examined. First, we see from the figures that the effects of shocks on the number of attacks for each type of target are relatively short: The main effects are seen in the immediate 6 months after a shock. Second, the impulse responses show the following patterns: The size of the shocks to the officials, military, and business equations shrinks over time, while the size of shock to the private parties equation is relatively highest in the 1989–2001 period. Note in particular the changes in the scales of the figures; the dynamics are similar, while the scales and heterogeneity differ.

Figure 4 shows the impulse responses for the 1968:1–1973:1 subsample. The responses to surprise increases in the number of attacks on officials are short-term (about 1–2 months) increases in the number of attacks on the other target types (see the first column). A one-standard-deviation increase in attacks on officials leads to about two additional attacks on military targets. Innovations in attacks on military or business targets lead to no meaningful increases in attacks on other target types. This captures the insights in Table 4 where the statistically significant contemporaneous effects in this initial period were those between attacks on officials and the other targets.

During the 1973–1979 subsample when barriers are being deployed, Fig. 5 shows that the pattern of responses is different. A one-standard-deviation increase in the number of attacks on officials leads to no meaningful change in the number of attacks on military targets. There are increases in the number of attacks on business and private parties from these innovations. The number of attacks on private parties immediately reacts to more attacks on officials with about eight more attacks in this period. The response of a surprise increase in attacks on military targets is seen mainly in the roughly two additional attacks on private parties. In this subsample, there are no substantive responses by other targets to increased attacks on business or private parties. This is consistent with the focus by terrorists on state targets in the initial periods of transnational terrorism. It is also consistent with earlier barriers protecting business and private parties in select venues such as airports.

Figure 6 shows the responses of the equations during the state sponsorship of the 1980s. Most notable here is that the responses of the officials, military, and business equations are smaller than in the previous figures. Surprise attacks on officials lead to an additional two or three attacks on business and about

[15]In an additional analysis, reported in the supplemental replication materials, the subsamples are pooled and dummy variables are included to capture the shifts in the intensity of attacks across the subsamples. The results of this analysis (improperly) constrain the dynamics of the BaP-VAR to be the same across the subsamples. Nevertheless, the findings for the covariances are similar to those reported in Table 4. Further, we estimate Gaussian VAR regressions for each subsample and compute these same error covariances. Based on reasons stated earlier, these Gaussian results are biased downward relative to those reported here.

[16]The reason that we chose this method for computing the error bands is that it accounts for the main shape of the error bands and presents magnitudes of the responses that are identical to those from the other methods in Sims and Zha (1999). Future applications of the BaP-VAR(p) model should consider these alternative methods for reporting the uncertainty of the responses as well.

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press

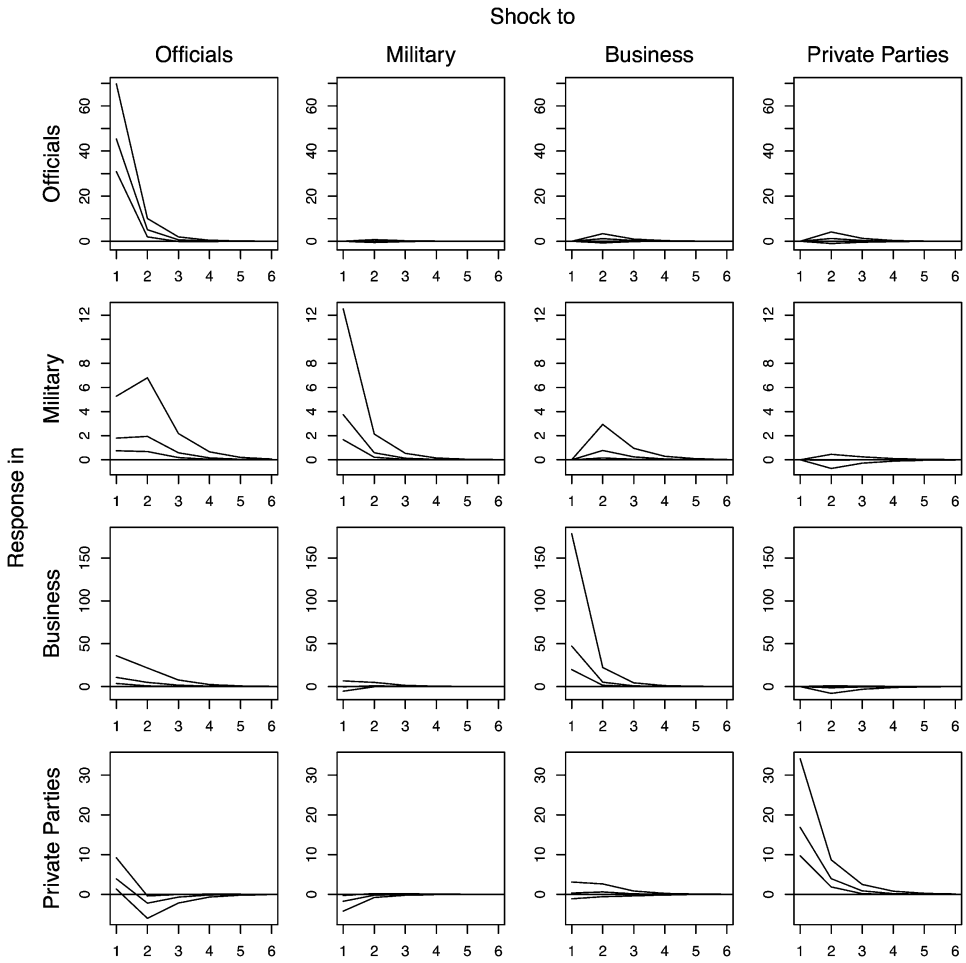


Fig. 4 Impulse responses for count VAR(1) model, 1968:1–1973:1 subsample.

five attacks on private party targets (see the first column of graphs in the figure). Innovations in attacks on military targets have little impact on the other target types. Shocks to business attacks lead to about two additional attacks on officials over 2 months. There is a small dynamic response in business attacks to attacks on private parties; about four additional attacks are seen in the second and third months after a new attack on private parties. State sponsorship marked an increase in carnage (Hoffman 2006) as sponsored groups (e.g., Abu Nidal Organization and PFLP) sought notoriety; consequently, "soft" business and private party targets became more favored.

The responses for the subsample covering 1989:12–2001:9 have a great contrast to the others. First, we see that the innovations in attacks on officials lead to many more attacks on private party targets (see the first column of Fig. 7). This larger effect is the result of smaller shocks (around 9 attacks) to the officials' innovations. Thus, there is a larger response in attacks on private parties relative to the smaller shock. The response of the private parties equation is now larger for all shocks relative to the earlier periods. Surprise shocks to the military and business variables each lead to roughly 10 new attacks on private parties over a 3-month period. This is direct evidence for the BaP-VAR over the static multivariate (Chib and Winkelmann 2001) specification since the effects of the shock are not immediate. So, while terrorist attacks on official, military, and business targets drop in this period, they are the stronger predictors of subsequent attacks on private parties. This pattern indicates a substitution-like shift to the softest of all target types—private parties—during the start of the fundamentalist era.

The impulse responses for the final period, 2001:10–2008:12, are nearly identical to those in Fig. 7 and are omitted from the presentation. The only substantive difference in this latter period is that the diminution of the responses in the officials, military, and business series continues, evidence that there is a

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press

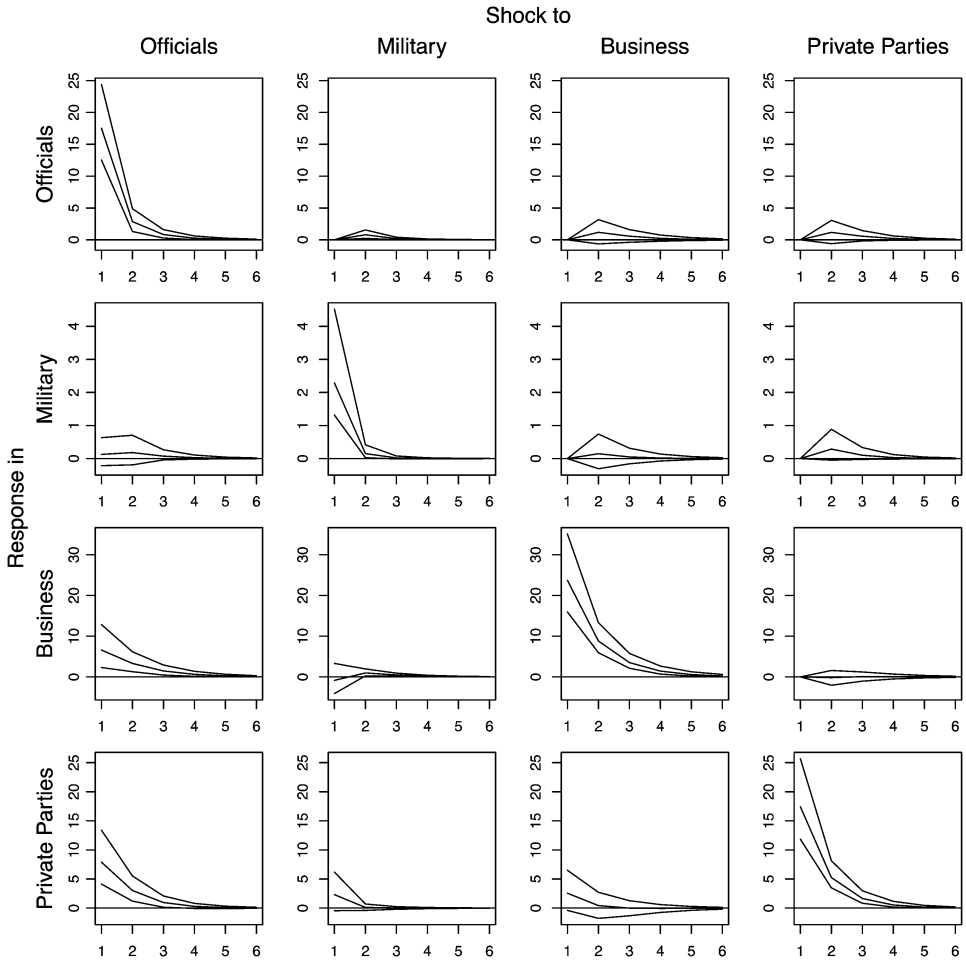


Fig. 5 Impulse responses for count VAR(1) model, 1973:2–1979:12 subsample.

shift away from attacks on these targets. The general pattern is that the responses to shocks are positive—like those seen in Fig. 7. This is evidence that the post-9/11 responses are complementary with attacks on official, military, and business targets feeding into more attacks on private parties and vice versa.

The impulse responses in Figs. 4–7 present the median dynamic responses to one-standard-deviation shocks to each equation for each subsample, but to make inferences about changes in the covariances over time, we compute the DFEV. These tell us how much of the *variation* (not just the multiplier) is explained by innovations in each series or how much of the forecast variation is due to changes in each variable. This is a dynamic analysis of interseries variance because we can trace out how much of the variation in each target count equation is due to changes in the other counts. Analogous to the impulse response analysis, the contemporaneous covariance of the DFEV is identified using a Cholesky decomposition and the same ordering of the equations. The DFEV report the percentage of the total variation in each equation explained by each innovation over 12 months. These DFEVs are computed using the posterior sample of coefficients.

Figure 8 shows the DFEV for our BaP-VAR(1) model for each of the subsamples. In this figure, the rows of graphs show the decomposition of the variance of the responses to a one-standard-deviation shock in each of the variables. The responses to shocks in the officials series are in black, the military series in red, the business series in blue, and the private parties series in green. The columns correspond to the subsamples defined earlier. At each time point in each panel, the total proportion of the forecast variation sums to 100%. Each graph then shows the time path of the proportion of the variation of the responses for each of the possible shocks.

If each target series is driven by its own dynamics, then the percentage of the variation of a given response is mainly due to its own innovations or dynamic response to shocks so that the color of the

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press

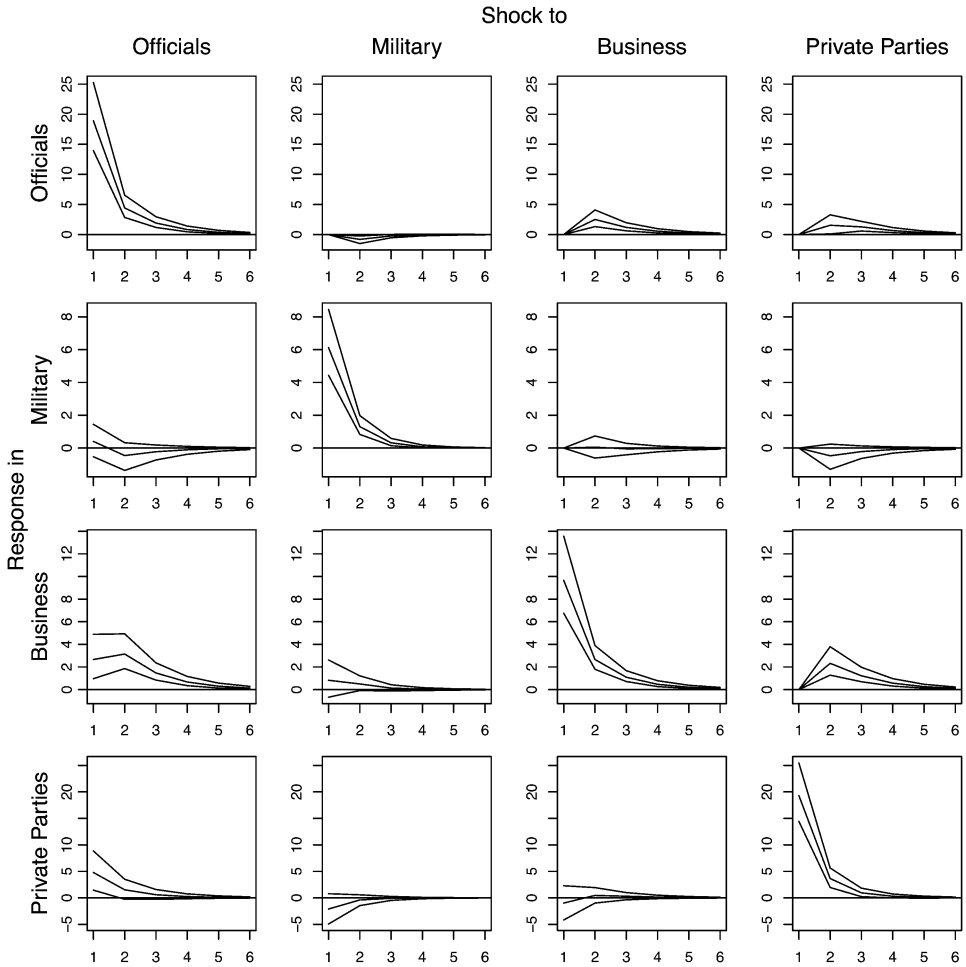


Fig. 6 Impulse responses for count VAR(1) model, 1980:1–1989:11 subsample.

dominant response should match across a row: black in the first, red in the second, blue in the third, and green in the fourth. Simple inspection reveals that this is not the case, meaning that there is some important and meaningful *variation of the responses over time* rather than in the modal responses to shocks (cf. Figs. 4–7). Thus, the initial review of Fig. 8 shows that while the earlier results primarily show complementarity in the *intensity* of attacks across the target types, the *variation* across target types in a given subsample demonstrates dynamic changes in the main factors that drive the predictability of future attacks on each target type.

Specifically, we see several patterns of changed variation across targets and subsamples in the DFEVs. In the first subsample (column), the birth of modern terrorism, the black, red, and blue lines are dominant in the first three panels. This means that innovations in attacks on officials, military, and business predict most of the variation in the observed attacks on official, military, and business targets. These three target types drive 80% of the variation in attacks on each other. By contrast, during this initial period, over 80% of the forecast variation in attacks on private parties is driven by changes in their own attacks.

In the period of technological barriers (second column), there is a different pattern: While most of the variation in attacks on officials is due to attacks on officials, around 95% of the variation in attacks on the military is now equally explained by attacks on military, business, and private parties. This is evidence of a dynamic shift in the drivers of attacks on the military in this period since we see a shift in how much variation we can explain in one target type from attacks on another. During this period, attacks on business and private parties are driven by their own dynamics since the forecast decompositions are driven by their own innovations.

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press

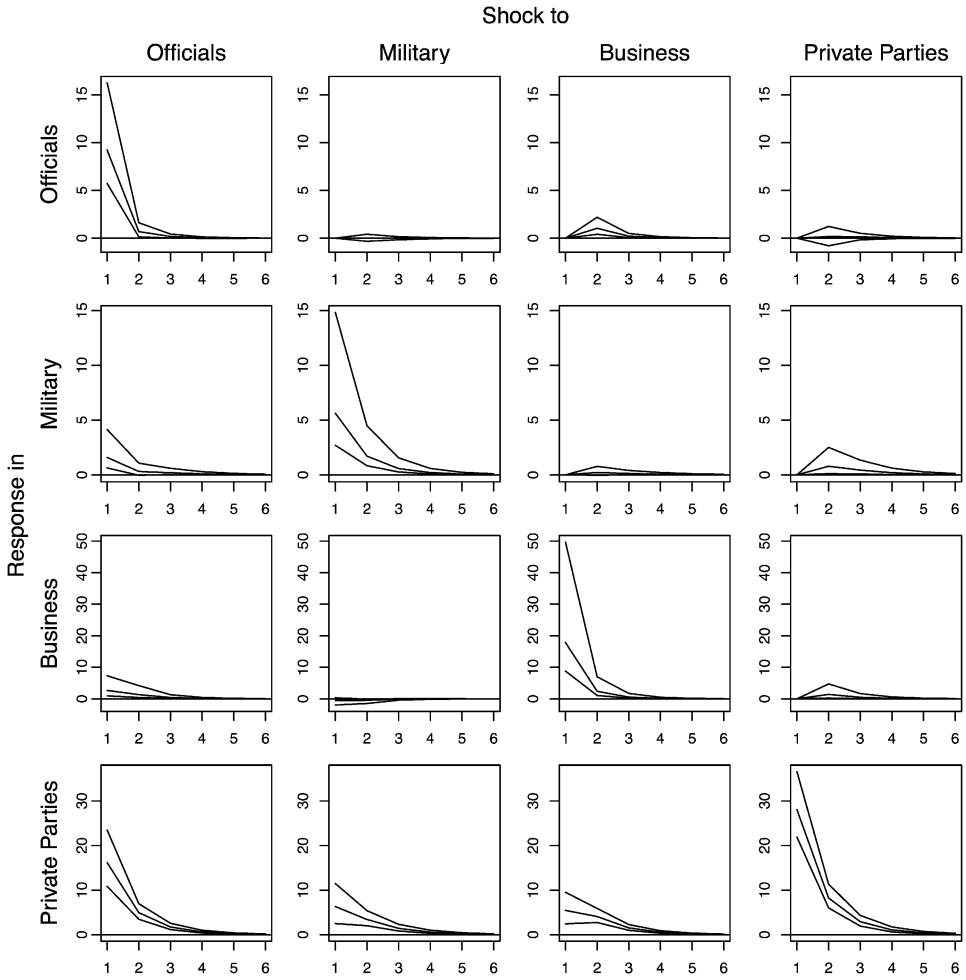


Fig. 7 Impulse responses for count VAR(1) model, 1989:12–2001:9 subsample.

During the state sponsorship regime (third column of Fig. 8), there is an alternative predictive dynamics among the target classes. The main source of variation in the military target attacks now becomes attacks on military targets themselves. In this period, only about 20% of the variation in attacks on business targets are the result of innovations in attacks on the other target types. In the last row of column three, over 90% or more of the attacks on private parties are explained by attacks on private parties in the 1980–1989 period.

In the 1989:12–2001:9 period of fundamentalist dominance (column 4 of Fig. 8), there is another shift in what drives attack variation on the four target classes. First, we see that the attacks on business and private parties are dominant drivers of the variation in the number of attacks. Attacks on private parties explain over 65% of the variance in attacks on officials and over 50% of the variance in attacks on military targets. Second, it is attacks on business and private party targets that basically explain the plurality of attacks in general. Unlike the pre-1990 period, there is generally little effect of attacks on officials for predicting other attack targets. These main patterns are repeated in the most recent period, 2001:9–2008:12, where the variation in attacks is mainly explained by innovations in attacks on private parties and business targets. Over 80% of the variation in attacks on official and military targets is explained by attacks on private parties and business. For the business and private party variation over time in this subsample, 99% of the variation is explained by attacks on these two targets.

The DFEVs provide clear evidence of the changing dynamics of targeting decisions by transnational terrorists over time. While the impulse responses showed that there is complementarity in the intensity of terrorist targeting decisions, the DFEVs complete the picture and show how the total variation in the mix of target decisions demonstrates change over time and across target types.

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press

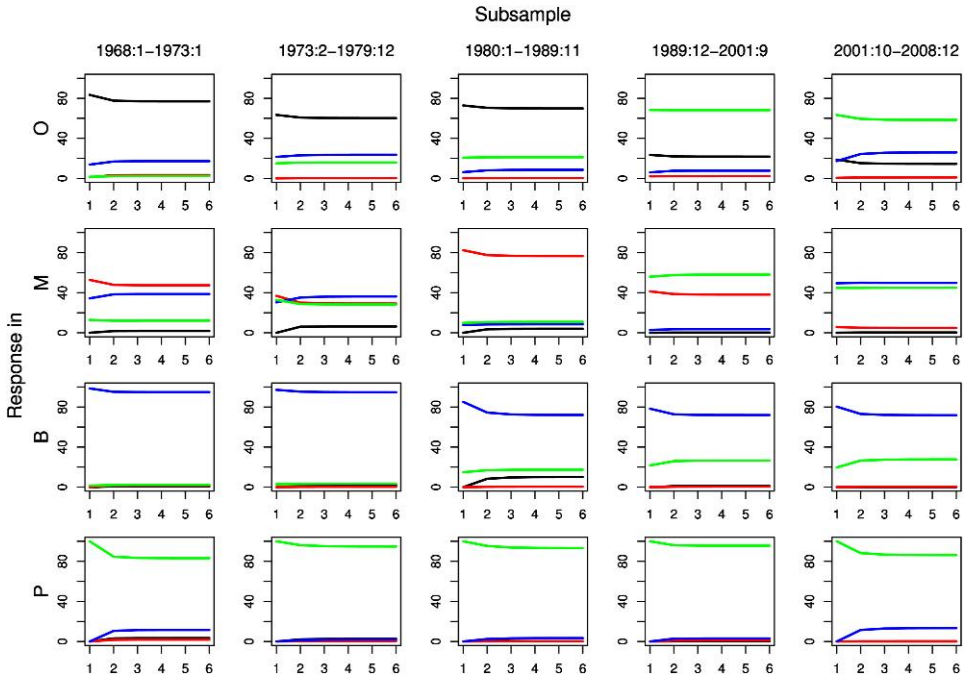


Fig. 8 Decompositions of the forecast error variances. Responses to shocks in officials (black), military (red), business (blue), and private parties (green) for each subsample.

6 Conclusion

This research develops a BaP-VAR model that permits the investigation of endogenous dynamic relationships among event count time series. The model is general and admits a dynamic specification via an observation-driven vector autoregression process. Unlike existing multivariate and count methods, the BaP-VAR allows for unrestricted endogenous correlation among the data being modeled. Our BaP-VAR estimator recovers endogenous relationships and dynamics that cannot be recovered by other count and time-series estimators.

The presentation of the BaP-VAR model also includes dynamic inference methods, such as impulse response functions, that trace out the effects of surprise shocks to the system of equations. In the BaP-VAR framework, DFEV measure the forecastability of one series based on another. Finally, the first example shows how to compute and present dynamic multipliers that measure the effects of changes in exogenous covariates on the outcomes of interest.

The superpower rivalry example shows the gains from using this new method over an extant one. The dynamic responses indicate that much of the rivalry was driven by the number of previously U.S.-initiated MIDs. Unlike King's (1989b), we are able to estimate jointly the effect of each country's military spending on the number of MIDs and determine the level of residual correlation among the MIDs. These results indicate that the earlier model is both dynamically misspecified and underestimates the contributions of these variables. Finally, the example shows that the dynamic interaction, not the exogenous variables, dominates the explanation of the pattern of MIDs.

In our terrorism targeting example, we uncover a clear pattern of dynamic interdependency; this pattern changes over the five distinct regimes and has not previously been identified. Moreover, there is a marked interdependency among the target time series in which complementarity prevails over the five regimes. Over time, innovation spillovers favor attacks on private parties, especially during the two regimes where fundamentalist terrorists are dominant. The regime with the least interdependency among targets is that of state sponsorship in the 1980s when terrorists appear more fixated on specific targets.

In contrast, the DFEVs display a marked dynamic pattern over time that is regime specific. During the first two pre-1980 regimes, attacks on officials and business explain a plurality of the variation in the other targets. The state sponsorship in the 1980s is a transition period with each target type being the

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press

dominant determinant of its own variation—this is not the case in the other four regimes. Moreover, attacks against private parties influence the variation in the other three target series. Thus, the state sponsorship regime demarcates the increased importance of private party attacks, consistent with the rising carnage per incident. By the post-9/11 regime, the dominant drivers of the variation in terrorist attacks are those of the softer targets—private parties and business.

Appendix A: BaP-VAR(p) Derivations and Properties

A.1 Properties of a Log-normal Distribution in the BaP-VAR

For a log-normal distribution, $w_{jt} \sim \text{LN}(\mu^*, \Sigma)$, where $w_{jt} = \exp(b_{jt})$ with $b_j \sim N(0, D)$ (as above), the mean and variance are $\mu^* = \exp(0.5\,\text{diag}(D))$ and $\Sigma = (\text{diag}(\mu^*))[\exp(D) - \mathbf{1}\mathbf{1}'](\text{diag}(\mu^*))$, where $\mathbf{1}$ is a vector of ones.

As in Chib and Winkelmann (2001), the latent mean of the count process for count j at time t is $\tilde{\lambda}_{tj} = \lambda_{tj}\mu^*_{tj}$, where $\lambda_{tj} = \exp(x_{tj}\beta_j)$ and $\mu^*_{tj} = w_{tj} = \exp(b_{tj})$. Here, the mean of y_{tj} differs from λ_{tj} by a scale factor (the random effects) for the basic model. For $\tilde{\lambda}_t = (\tilde{\lambda}_{t1}, \ldots, \tilde{\lambda}_{tm})$, define the diagonal matrix $\tilde{\Lambda}_t = \text{diag}(\tilde{\lambda}_t)$. Using iterated expectations, we have

$$E[y_{jt}|\beta, D] = \tilde{\lambda}_{jt} = \exp(x_{tj}\beta_j), \tag{A1}$$

$$V[y_{jt}|\beta, D] = \tilde{\Lambda}_j + \tilde{\Lambda}_j[\exp(D) - \mathbf{1}\mathbf{1}']\tilde{\Lambda}_j \tag{A2}$$

for a given estimate of β_j.

In the BaP-VAR mean equation (6b), the same iterated expectations apply. The conditional mean and variance for equation (7) are

$$E[y_t|\beta, A, D] = E(Ay_{t-1}) + \exp(x_{tj}\beta_j), \tag{A3}$$

$$V[y_t|\beta, A, D] = A\,\text{Cov}(y_{t-1})A' + \tilde{\Lambda}_j + \tilde{\Lambda}_j[\exp(D) - \mathbf{1}\mathbf{1}']\tilde{\Lambda}_j. \tag{A4}$$

The covariance including the lagged counts only adds the terms for the lagged counts, which are computed trivially with the posterior simulation outputs.

A.2 Derivation of the impulse response function (IRF) for the BaP-VAR

Impulse responses are found by computing

$$\frac{\nabla y_i(\ell + s)}{\nabla w_{sj}}, \tag{A5}$$

where ∇ is the discrete difference operation. This is the change in equation i at time $\ell + s$ for a shock in w_{sj} at time s. We can use the same argument as in the text to work out this computation for each equation-shock combination. This process can be generalized, however, if we recognize a slight abuse of notation and replace μ_t in equation (7) with its sample analog y_t. This makes sense since the expected value of y_t is μ_t in the model.

Let L^k be the lag operator that shifts y_t k periods or $L^k y_t = y_{t-k}$. Rewrite equation (7) as

$$y_t - Ay_{t-1} = \exp(x_t\beta)w_t, \tag{A6a}$$

$$(I - AL)y_t = \exp(x_t\beta)w_t, \tag{A6b}$$

$$y_t = (I - AL)^{-1}\exp(x_t\beta)w_t. \tag{A6c}$$

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press

Assuming that the roots of A are consistent with stationarity, we can invert the lag polynomial term, $(I - AL)$, and derive the VMA representation in equation (A6c). The inverse of the finite $(I - AL)$ polynomial will have an infinite number of terms that correspond to the VMA or impulse responses of the process.[17]

The error term in equation (A6c) is $w_t = \exp(b_t)$. This last equation is the impulse response for this count VAR model (e.g., Brandt and Williams 2007). In the present case, this is the impulse response for the latent variable of the multivariate count model. The sole remaining issue concerns how to handle the initial correlations of the shocks in the impulse response computation. We do this using a standard Cholesky decomposition of the estimated covariance matrix in equation (A4). Various methods can then be used to summarize the posterior impulse responses: See Brandt and Freeman (2006) for a discussion of the computation of error bands using eigendecompositions that account for the serial and cross-correlation of the responses (see Sims and Zha 1999).

A.3 MCMC Sampler for the BaP-VAR Model Parameters Posterior

The posterior sampler for the BaP-VAR model can be split into four basic blocks:[18]

Sample the latent effects, b: Draw from a multivariate log-normal target density with mean b and precision D.

Sample the regression effects, β: Conditional on b_i, sample β either for $j = 1, \ldots, m$ or as a single block using Metropolis–Hastings steps.

Sample the autoregressive process: Conditional on the latest effects and the regression parameters, sample the VAR process using a block Metropolis–Hastings step.

Sample the (inverse) covariance, D^{-1}: This last set of terms (the inverse of the precisions of the latent covariance matrix) is sampled from a Wishart, conditional on the first two steps.

Only the last of the densities in the algorithm has a standard representation (see Chib and Winkelmann 2001, 429–431). Hence, a custom Metropolis sampler must be implemented.[19] We can then implement the impulse responses and DFEV computations described above using the posterior sample of the parameters.

Funding

The U.S. Department of Homeland Security through the Center for Risk and Economic Analysis of Terrorism Events (CREATE) at the University of Southern California, Grant (007-ST-061-000001 to T.S., 2010-ST-061-RE0001 to T.S. and P.B.); and the National Science Foundation (0921051 to P.B.). However, any opinions, findings, conclusions, and recommendations are solely those of the authors and do not necessarily reflect the views of the Department of Homeland Security, CREATE, or the National Science Foundation.

[17] Let C_ℓ be the ℓth-period moving average coefficients for the impulse responses. For a VAR(1), these can be found using

$$\begin{aligned} C_1 &= A \\ C_2 &= AC_1 \\ &\vdots \\ C_\ell &= AC_{\ell-1}. \end{aligned}$$

[18] The details of a tuned MCMC sampler for the posterior can be found in Chib and Winkelmann (2001).

[19] One way to avoid the need to tune a Metropolis sampler for such an application is to program the algorithm in a general purpose Bayesian sampler like BUGS, JAGS, or UMACs. For the present application, we adapt the BUGS framework for the Chib and Winkelmann (2001) model proposed by Congdon (2005) and also employed in El-Basyouny and Sayed (2009), Congdon (2007), and Ma (2006). Here, we implement the model using JAGS and R for a general implementation.

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press

References

Aitchison, J., and C. H. Ho. 1989. The multivariate Poisson-log normal distribution. *Biometrika* 76:643–53.

Alexander, Y., and D. A. Pluchinsky. 1992. *Europe's Red Terrorists: The Fighting Communist Organizations*. London: Frank Cass.

Azar, S. 1982. *The Codebook of the Conflict and Peace Data Bank (COPDAB)*. College Park, MD: Center for International Development.

Brandt, P. T., M. P. Colaresi, and J. R. Freeman. 2008. The dynamics of reciprocity, accountability, and credibility. *Journal of Conflict Resolution* 52:343–74.

Brandt, P. T., and J. R. Freeman. 2006. Advances in Bayesian time-series modeling and the study of politics: Theory testing, forecasting, and policy analysis. *Political Analysis* 14:1–36.

Brandt, P. T., and T. Sandler. 2009. Hostage taking: Understanding terrorism event dynamics. *Journal of Policy Modeling* 31:758–78.

———. 2010. What do transnational terrorists target? Has it changed? Are we safer? *Journal of Conflict Resolution* 54:214–36.

Brandt, P. T., and J. T. Williams. 2001. A linear Poisson autoregressive model: The Poisson AR(p) model. *Political Analysis* 9:164–84.

———. 2007. *Multiple Time-series Models*. Thousand Oaks, CA: Sage.

Brandt, P. T., J. T. Williams, B. O. Fordham, and B. Pollins. 2000. Dynamic modeling for persistent event count time series. *American Journal of Political Science* 44:823–43.

Chen, A. H., and T. F. Siems. 2004. The effects of terrorism on global capital markets. *European Journal of Political Economy* 20:249–66.

Chib, S., and R. Winkelmann. 2001. Markov chain Monte Carlo analysis of correlated count data. *Journal of Business & Economic Statistics* 19:428–35.

Congdon, P. 2005. *Bayesian Models for Categorical Data*. Chichester, UK: John Wiley & Sons.

———. 2007. Bayesian modeling strategies for spatially varying regression coefficients: A multivariate perspective for multiple outcomes. *Computational Statistics & Data Analysis* 51:2586–601.

Deb, P., and P. K. Trivedi. 1997. Demand for medical care by the elderly: A finite mixture approach. *Journal of Applied Econometrics* 12:313–36.

El-Basyouny, K., and T. Sayed. 2009. Collision prediction models using multivariate Poisson-lognormal regression. *Accident Analysis and Prevention* 41:820–28.

Enders, W., and T. Sandler. 1993. The effectiveness of anti-terrorism policies: A vector-autoregression-intervention analysis. *American Political Science Review* 87:829–44.

———. 2005. After 9/11: Is it all different now? *Journal of Conflict Resolution* 49:259–77.

———. 2006a. Distribution of transnational terrorism among countries by income classes and geography after 9/11. *International Studies Quarterly* 50:367–93.

———. 2006b. *The Political Economy of Terrorism*. Cambridge: Cambridge University Press.

Grunwald, G. K., R. J. Hyndman, L. Tedesco, and R. L. Tweedie. 2000. Non-Gaussian conditional linear AR(1) models. *Australian and New Zealand Journal of Statistics* 42:479–95.

Hamilton, J. 1994. *Time-series Analysis*. Princeton, NJ: Princeton University Press.

Hanson, M. A., and M. B. Schmidt. 2011. The impact of coalition offensive operations on the Iraqi insurgency. *Applied Economics* 43:2251–65.

Heinen, A., and E. Rengifo. 2007. Multivariate autoregressive modeling of time-series count data using copulas. *Journal of Empirical Finance* 14:564–83.

Hoffman, B. 2006. *Inside Terrorism: Revised and Expanded.* New York: Columbia University Press.

Karlis, D., and L. Meligkotsidou. 2005. Multivariate Poisson regression with covariance structure. *Statistics and Computing* 15:255–65.

King, G. 1988. Statistical models for political science event counts: Bias in conventional procedures and evidence for the exponential Poisson regression model. *American Journal of Political Science* 32:838–63.

———. 1989a. A seemingly unrelated Poisson regression model. *Sociological Methods & Research* 17:235–55.

———. 1989b. *Unifying Political Methodology: The Likelihood Theory of Statistical Inference*. Princeton, NJ: Princeton University.

Lütkepohl, H. 2005. *New Introduction to Multiple Time-series Analysis*. Berlin: Springer.

Ma, J. 2006. *Bayesian Multivariate Poisson-Lognormal Regression for crash prediction on rural two-lane highways*. PhD diss., University of Texas, Austin, TX.

McGinnis, M. D., and J. T. Williams. 1989. Change and stability in superpower rivalry. *American Political Science Review* 83:1101–23.

Mickolus, E. F., T. Sandler, J. M. Murdock, and P. Flemming. 2009. *International Terrorism: Attributes of Terrorist Events, 1968–2008 (ITERATE)*. Dunn Loring, VA: Vinyard Software.

Munkin, M. K., and P. K. Trivedi. 1999. Simulated maximum likelihood estimation of multivariate mixed-Poisson regression models, with applications. *Econometrics Journal* 2:29–48.

Ord, K., C. Fernandes, and A. C. Harvey. 1993. Time-series models for multivariate series of count data. In *Developments in time-series analysis in honor of Maurice B. Priestley*, ed. T. S. Rao, 295–309. London: Chapman & Hall.

Park, J. H. 2010. Structural changes in the U.S. president's uses of force abroad. *American Journal of Political Science* 54:766–82.

Plummer, M. 2008. Penalized loss functions for Bayesian model comparison. *Biostatistics* 9:523.

Sims, C. A., D. F. Waggoner, and T. A. Zha. 2008. Methods for inference in large multiple-equation Markov-switching models. *Journal of Econometrics* 146:255–74.

Sims, C. A., and T. A. Zha. 1998. Bayesian methods for dynamic multivariate models. *International Economic Review* 39:949–68.

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press

———. 1999. Error bands for impulse responses. *Econometrica* 67:1113–56.
Spiegelhalter, D., N. Best, B. Carlin, and A. Van Der Linde. 2002. Bayesian measures of model complexity and fit (with discussion). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 64:583–639.
Williams, J. T., and M. D. McGinnis. 1988. Sophisticated reaction in the U.S.-Soviet arms race: Evidence of rational expectations. *American Journal of Political Science* 32:968–95.
———. 1992. The dimension of superpower rivalry: A dynamic factor analysis. *Journal of Conflict Resolution* 36:86–118.

https://doi.org/10.1093/pan/mps001 Published online by Cambridge University Press