

**TRANSNATIONAL TERRORISM 1968-2000:
THRESHOLDS, PERSISTENCE, AND FORECASTS**

by

Walter Enders*
Department of Economics, Finance, and Legal Studies
University of Alabama
Tuscaloosa, AL 35487
wenders@cba.ua.edu
205-348-8972
205-348-0590 (fax)

and

Todd Sandler†
School of International Relations
University of Southern California
Von Kleinsmid Center 330
Los Angeles, CA 90089-0043
tsandler@usc.edu
213-740-9695
213-742-0281 (fax)

Revised: June 2004

* Walter Enders is the Bidgood Chair of Economics and Finance in the School of Business at the University of Alabama.

† Todd Sandler is the Robert R. and Katheryn A. Dockson Professor of International Relations and Economics at the University of Southern California.

The research for this study was funded, in large part, by a grant from the Development Research Group of the World Bank. All opinions expressed are solely those of the authors.

TRANSNATIONAL TERRORISM 1968-2000:
THRESHOLDS, PERSISTENCE, AND FORECASTS

Abstract

This paper applies a threshold autoregression (TAR) model to a casualties time series to show that the autoregressive nature of such events depends on the level of terrorism at the time of a shock. Following a shock, persistence of heightened attacks characterizes low-terrorism regimes, but not high-terrorism regimes. Similar findings are associated with incidents with deaths, bombings with deaths, and hostage taking. In contrast, the assassinations series indicates some persistence even in the high-terrorism state, while the threats/hoaxes series displays persistence in only the high-terrorism state. For all series studied, the TAR model outperforms a standard autoregressive representation. A forecasting method is engineered based on the TAR estimates, and nicely tracks resource-using events.

JEL codes: D74, C32, H56

TRANSNATIONAL TERRORISM 1968-2000: THRESHOLDS, PERSISTENCE, AND FORECASTS

1. Introduction

The death, destruction, and economic impacts resulting from the four hijackings on September 11, 2001 (henceforth, 9/11) made policymakers and the public acutely aware of the threat posed by transnational terrorism to liberal democracies worldwide.¹ The sheer magnitude of 9/11 led to repercussions that have touched most countries through economic spillovers, collateral damage, or security expenditures. As the public learned about the global reach of the al-Qaida network, the true threat of transnational terrorism became better understood. To assist in this understanding, social scientists must apply theoretical and empirical tools to analyze the actions of terrorists and identify the most appropriate response of governments. In particular, economic methods and models can enlighten policymakers on transnational terrorism (Sandler and Enders 2004). Consumer-choice models can be tailored to display terrorists' choices when constrained by their resources and the policies of the authorities (Landes 1978; Sandler, Tschirhart, and Cauley 1983). Empirically, time-series techniques can derive policy implications and forecasts (Brophy-Baermann and Conybeare 1994; Enders and Sandler 1993, 1995).

Terrorism is the premeditated use or threat of use of extranormal violence or brutality by subnational groups or individuals to obtain a political objective through intimidation or fear directed at a large audience. Terrorist acts are purposely brutal to create an atmosphere of fear, while publicizing the terrorists' cause. As the public becomes numb to their acts of violence, terrorists ratchet up the carnage to maintain media interest. An essential ingredient in defining terrorism is the presence of a political objective, which can involve getting a country out of an international organization, the formation of an Islamic state, the United States ending its support of Israel, or other goals.

A crucial factor of modern-day terrorism is the transnational implications of many terrorist attacks. When a terrorist incident in one country involves victims, perpetrators, or audiences in two or more countries, terrorism takes on a *transnational* character. A terrorist act may be transnational owing to its impact, its planning and execution, its perpetrators (if known), or its targets and resulting damage. A skyjacking of a plane with passengers from multiple countries is a transnational terrorist event, as is a skyjacking of a domestic flight originating in one country but terminating in another country. The skyjackings on 9/11 were transnational terrorist acts, as were the near-simultaneous bombings of the US embassies in Nairobi, Kenya, and Dar es Salaam, Tanzania, on 7 August 1998. We focus on transnational terrorism because of its importance in the post-9/11 world and data availability.

Although transnational terrorist attacks are down in the post-Cold War period, the lethality of these attacks has increased with each incident almost 17 percentage points more likely to result in death or injuries (Enders and Sandler 2000). The *casualties* time series, involving one or more deaths or injuries, displays predictable factors unlike the noncasualties series, which is largely random following detrending. In a subsequent study, Enders and Sandler (2002) use a threshold autoregression (TAR) model applied to just the *death* series to identify nonlinearities following shocks. When terrorism is in a low-intensity regime with attacks below a data-determined threshold, shocks that raise the level of events are sustainable. In contrast, shocks during a high-intensity regime of terrorism (i.e., attacks above a data-determined threshold) are episodic with series returning rapidly to the mean number of events.

The primary purpose of the current study is to apply TAR analysis to additional transnational terrorist time series including important subcomponents of the overall all-incident series. In particular, we analyze both the *death* and *casualties* time series for nonlinearities as well as the component series of *bombings* with one or more deaths. Additionally, we apply a

TAR model to series involving *assassinations*, *hostage taking*, and *threats and hoaxes*. A wide variety of nonlinearities are uncovered for component series – e.g., shocks to the *assassinations* series display some persistence in both the low- and high-intensity regime; shocks to the *threats and hoaxes* series result in persistence in only the high-intensity regime. An important finding is that the response to an external shock depends on current activity levels, which may vary among different kinds of terrorist events so policy prescriptions must be tailored to the kind of event. A secondary purpose is to engineer a forecasting technique based on the TAR representation of the component time series.

2. Terrorists' Choice-Theoretic Model

We view terrorists as rational actors who maximize their well-being by allocating their resources among alternative goals while accounting for underlying constraints.² Rationality is judged by the terrorists' predictable response to changes in their constraints and not by their particular goals or the means used to achieve them. Terrorist reactions to shocks in high- and low-terrorism states are dependent on available resources. For instance, the ability to shift sufficient resources between activities and intertemporally to sustain higher levels in a low-intensity period is typically greater than in an already high-intensity regime.

At the group level, terrorists must allocate resources at three different layers: (i) between terrorist (T) and nonterrorist (N) activities, (ii) among alternative terrorist modes of attack, and (iii) among different time periods. For conceptual simplicity, we assume that each of these decisions is independent. When deciding between terrorist and nonterrorist activities (e.g., propaganda, seeking election), the terrorist group – treated as a unitary entity – must determine N and T so as to maximize expected utility $[E(U)]$:³

$$E(U) = \pi U(W^S) + (1 - \pi)U(W^F), \quad (1)$$

where $U(W)$ is the standard von Neumann-Morgenstern utility function, π is the probability of success, and $1 - \pi$ is the probability of failure. Moreover, W^S and W^F are the net wealth equivalent measures over the two states – success (S) and failures (F). These measures can be expressed as:

$$W^S = w + g^N(R^N) + g^T(R^T, e), \quad (2)$$

$$W^F = w + g^N(R^N) - f(R^T, \tilde{e}). \quad (3)$$

In Equations 2 and 3, w denotes the group's current assets, net of current earnings or losses. Terrorists receive w and the net gain, g^N , from nonterrorist incidents regardless of the outcome of the terrorist activities. These latter gains increase as the extent of nonterrorist activities, as measure by R^N , increases. The certainty of g^N is not as important as these gains being more certain than those from terrorist activities, g^T . The latter depends on the resources expended on terrorism, R^T , and environmental factors, e , which involve actions by the authorities to limit the gains from terrorism (e.g., freezing assets, hardening targets, and intelligence).

In Equation 3, the net wealth for a terrorist failure includes the group's current assets plus its gains from legitimate activities *minus* the monetary equivalence of the fine or penalty, f , from failure. When failure is associated with no arrests, f represents the lost resources. In more dire circumstances where terrorists are apprehended, the penalty refers not only to wasted resources (including fallen comrades), but also the opportunity cost of imprisonment and any monetary penalties. Fines are likely positively related to the level of terrorist activity, R^T . In $f(\bullet)$, \tilde{e} indicates an environmental factor that influences the level of losses during failed campaigns.

A secondary choice involves the allocation of R^T among alternative modes of attacks. Suppose that a group only uses two kinds of operations – hostage taking (h) and bombings (b). Further suppose that the per-unit costs of each kind of operation is $C(h, e_h)$ and $B(b, e_b)$ for hostage taking and bombings, respectively, where unit cost depends on the level of operations

and actions by the authorities – e_h and e_b . The terrorists now must choose h and b to maximize

$$U(h, b),$$

subject to:

$$R^T = C(h, e_h)h + B(b, e_b)b. \quad (5)$$

If the bordered Hessian determinant, $|H|$, is positive,⁴ then a maximum exists. Under standard assumptions, comparative statics show that hostage taking and bombings decrease when the authorities manage to limit R^T . Moreover, actions by the authorities to increase the unit cost of, say, hostage taking cause terrorists to switch some operations to the now relatively cheaper bombing events. Sufficient economies of scale in specific modes of operations may create nonconvexities where the terrorists specialize in a single type of event.⁵ With economies of scale, efforts by the authorities to raise the unit cost of a mode of operation may cause a group to shift entirely from one mode to another.

A third choice of terrorists involves an intertemporal allocation of resources. Analogous to other investors, terrorists can invest resources to earn a rate of return, r , per period. When terrorists want to augment operations, they can cash in some of their invested resources. Suppose that terrorists have a two-period horizon and must decide terrorist activities today (T_0) and tomorrow (T_1) based on resources today (R_0) and tomorrow (R_1). The intertemporal budget constraint is:

$$T_1 = R_1 + (1+r)(R_0 - T_0), \quad (6)$$

where tomorrow's terrorism equals tomorrow's resource endowment plus (minus) the earnings on savings (the payments on borrowings) from the initial period. Terrorists maximize an intertemporal utility function, $U(T_0, T_1)$, subject to Equation 6 and, in so doing, decide terrorist activities over time. Thus, terrorists can react to shocks by augmenting operations not only from

curbing nonterrorist activities, but also through an intertemporal substitution of resources. If high-terrorism regimes are supported by an intertemporal substitution, then shocks during such a regime can typically elevate resource-using terrorist activities for only a short time, unlike low-terrorism regimes where terrorists can draw on accumulated resources. This prediction may not characterize non-resource-using threats and hoaxes.

Thus far, we analyze the choice-theoretic decision of a single terrorist group. When groups are linked in a network, the same model applies to the network by attributing the utility and resource constraints to be those of the network. If terrorists are tied together implicitly through similar hatreds and grievances, then multiple terrorist groups may act as a unified whole and respond identically to shocks, so that our choice-theoretic representation would also characterize this conglomerate of groups. By copying one another's innovations, terrorists worldwide take on the appearance of a single group, thereby giving rise to distinct peaks and troughs in transnational terrorist events (Enders and Sandler 1999; Faria 2003).

3. Data and Terrorist Time Series

Data on transnational terrorist incidents are drawn from *International Terrorism Attributes of Terrorist Events* (ITERATE), which records the incident date, type of event, casualties (i.e., deaths or injuries), if any, and other variables. By splicing together the earlier ITERATE data sets, ITERATE 5's "common" file contains 40 or so key variables common to all terrorist events from 1968 to 2000 (Mickolus et al. 2002). Thus, earlier ITERATE data for various subperiods no longer need to be consulted. Coding consistency for ITERATE and its updates has been maintained by applying the same criteria for defining transnational terrorist events and associated variables.⁶ ITERATE draws its information from the world newsprint and electronic sources.

In total, we extract six quarterly time series from ITERATE for the purposes of the paper. The most inconclusive of these series is the *casualties* series, in which one or more persons were injured or killed in a transnational terrorist incident. The *death* series is a subset of the *casualties* series and includes only incidents where one or more persons – terrorist or victim – died. We focus on these time series, because earlier work indicates that they are more predictable than the inclusive series of all incidents (Enders and Sandler 2000). We use quarterly, rather than monthly, data to avoid periods with zero observations. Because bombing is the favorite tactic of terrorists, accounting for almost half of all events, a time series involving bombings with one or more deaths (henceforth called the *bomb* series) is also extracted. *Bomb* combines seven types of events – explosive bombings, letter bombings, incendiary bombings, missile attacks, car bombings, suicide car bombings, and mortar and grenade attacks – involving at least one death.

The remaining three series are components of the *all-incidents* series. The *assassinations* series consists of politically based murders or attempted murders. Although most assassinations end in one or more deaths, some are unsuccessful so that the *assassinations* series is not a subset of the *death* or *casualties* series. The *hostage* series includes kidnappings, transnational skyjackings, nonaerial hijackings, and barricade and hostage-taking events drawn from ITERATE. These events need not end with a casualty. Finally, the *threats/hoaxes* series combines two component events: threats or promises of future actions and hoaxes or falsely claimed past actions. By their nature, threats and hoaxes involve no casualties and uses little or no resources.

4. Estimates of Casualties and Deaths

There are several reasons to believe that the dynamic patterns of transnational terrorists display nonlinear behavior. In relatively tranquil times, terrorists can replenish and stockpile

resources, recruit new members, raise funds, and plan for future attacks. Terrorism can remain low until an event occurs that switches the system into a high-terrorism regime. Because each terrorist attack utilizes scarce resources, increased terrorist activities in response to shocks during high-terrorism states are not anticipated to exhibit a high degree of persistence as resource stocks become depleted quickly. Since periods with little terrorism involve few resource costs, an elevated campaign in response to shocks during a low-terrorism regime can be highly persistent as terrorists use stockpiled resources or normal resource inflows from activities (e.g., extortion) or contributions.

Moreover, Faria (2003) shows that the attack and counterattack strategies played by the authorities and terrorists can lead to cyclical patterns. As in a game of cat-and-mouse, high-terrorism regimes induce political leaders to undertake counter-terrorism policies. Once the overall amount of terrorism declines, the political will to sustain a vigilant counter-terrorism policy dissipates and a new cycle begins. Enders and Sandler (1995, 2000) use spectral analysis to show that the most logistically complex terrorist events have the longest cycles; however, the length of a cycle need not be deterministic. As such, an alternative estimation strategy is needed to account for the possibility of high- versus low-terrorism states.

A realistic way to capture the nature of terrorist campaigns is to use a regime-switching model, which allows for several states of the world (or regimes) such that the severity of terrorism differs across states. To formally model this type of behavior, we utilize the two-regime version of the TAR model developed by Tong (1983, 1990):

$$y_t = I_t \left[\alpha_0 + \sum_{i=1}^p \alpha_i y_{t-i} \right] + (1 - I_t) \left[\beta_0 + \sum_{i=1}^p \beta_i y_{t-i} \right] + \varepsilon_t, \quad (7)$$

$$I_t = \begin{cases} 1 & \text{if } y_{t-d} \geq \tau \\ 0 & \text{if } y_{t-d} < \tau, \end{cases} \quad (8)$$

where y_t is the series of interest, α_i and β_i are coefficients to be estimated, τ is the value of the threshold, p is the order of the model, d is the delay parameter, and I_t is the Heaviside indicator function.⁷

The nature of the TAR model is that there are two states of the world. In the high-terrorism state, y_{t-d} exceeds the value of the threshold τ , so that $I_t = 1$ and y_t follows the autoregressive process, $\alpha_0 + \sum \alpha_i y_{t-i}$. Similarly, in the low-terrorism state, $y_{t-d} < \tau$, so that $I_t = 0$, and y_t follows the autoregressive process, $\beta_0 + \sum \beta_i y_{t-i}$. There are two attractors or potential “equilibrium” values: in the high-terrorism state, the system is drawn toward $\alpha_0/(1 - \sum \alpha_i)$, and, in the low-terrorism state, the system is drawn toward $\beta_0/(1 - \sum \beta_i)$. Note that the TAR model is equivalent to an AR(p) model when all $\alpha_i = \beta_i$.

To avoid the problem of overfitting, we use the same type of nonlinear model for each of the various terrorist series. It seems natural to apply a TAR model in the form of Equation 7 and 8 to compare the persistence of high- versus low-terrorism regimes. First, a TAR model permits us to estimate the value of the threshold without imposing an a priori line of demarcation between the regimes. Second, the key feature of the TAR model is that a sufficiently large ε_t shock can cause the system to switch between regimes. The dates at which the series crosses the threshold are not specified beforehand by the researcher; instead, the data determine whether the series is in the high- or low-terrorism state. Third, because the variance of ε_t is regime independent, the different degrees of persistence can be compared using characteristic roots. When the characteristic root for a regime is zero, there is an immediate jump to the attractor for that regime. If, however, a regime has a unit characteristic root, the number of incidents will meander indefinitely until a shock occurs that causes a regime change. In intermediate cases, the speed of adjustment within a regime is inversely related to the largest root for that regime.

Chan (1993) shows that a grid search over all potential values of the thresholds yields a

super-consistent estimate of the unknown threshold parameter τ . To use the method, order the observations from smallest to largest such that:

$$y^1 < y^2 < y^3 \dots < y^T. \quad (9)$$

Each value of y^j is then allowed to serve as an estimate of the threshold τ . For each of these values, the Heaviside indicator is set using Equation 8 and an equation in the form of Equation 7 is then estimated. The regression equation with the smallest residual sum of squares contains the consistent estimate of the threshold. We follow the conventional practice of excluding the highest and lowest 15% of the $\{y^j\}$ values to ensure an adequate number of observations on each side of the threshold. Similarly, the delay factor, d , is selected as the arranged autoregression providing the best fit. As discussed in Hansen (1997), since τ is estimated consistently, the estimates of α_i and β_i converge to a t -distribution as sample size increases.

Incidents with Casualties

We begin by estimating the number of incidents with casualties (cas) as the linear AR(3) autoregressive process:

$$cas_t = 5.91 + 0.261cas_{t-1} + 0.310cas_{t-2} + 0.209cas_{t-3} + \varepsilon_t; \quad AIC = 1205.72, \quad (10)$$

(2.83) (2.98) (3.59) (2.40)

where cas_t represents the number of incidents with one or more casualties in period t , ε_t is the error term, and the coefficients' t -statistics are in parentheses.⁸ The model appears adequate in that it satisfies the standard diagnostic tests. All t -statistics are significant at conventional levels and the point estimates of the autoregressive coefficients imply stationarity. The results of a Dickey-Fuller test allow us to reject the null hypothesis of a unit-root at the 5% significance level. Moreover, the Ljung-Box Q -statistics indicate that the residuals are serially uncorrelated. For example, Q -statistics using the first 4, 8, and 12 lags of the residual autocorrelations have

prob-values of 0.98, 0.52, and 0.72, respectively.⁹

Correlation coefficients are measures of linear association and may not detect misspecifications due to nonlinearities in the data. We begin a search for the most appropriate TAR representation by estimating a model in the form of Equations 7 and 8 using a value of $p = 3$. The resulting model was clearly over-parameterized since a number of coefficients were insignificant at the 20% level. Once we eliminated these coefficients (although for the usual reasons, we retained the intercepts), we obtained:

$$\begin{aligned} cas_t = & [17.87 + 0.189cas_{t-1} + 0.237cas_{t-2}] I_t \\ & (3.19) \quad (1.83) \quad (1.83) \\ & + [3.92 + 0.423cas_{t-1} + 0.398cas_{t-3}] (1 - I_t) + \varepsilon_t; \end{aligned} \quad (11)$$

(1.48) (2.97) (3.12)

AIC = 1205.07, $\tau = 25$, and $d = 2$.

Diagnostic checking indicates that the model is appropriate. For example, the first twelve autocorrelations of the residuals are less than 0.14 in absolute value and the *prob*-values for the Ljung-Box $Q(4)$, $Q(8)$, and $Q(12)$ statistics are 0.98, 0.57, and 0.81, respectively. Even though the TAR model contains seven parameters (i.e., six coefficients plus τ), the AIC selects it over the linear model. Note that the coefficients for cas_{t-1} and cas_{t-2} have *prob*-values near 0.07; paring down the model by eliminating either of these terms would substantially reduce the AIC and decrease the estimated persistence of the high-terrorism state.

The threshold model yields very different implications about the behavior of the *cas* series than the linear model. The latter indicates that the *cas* sequence converges to its long-run mean of 25.4 incidents per quarter. Since the linear specification makes no distinction between high- and low-terrorism states, the degree of autoregressive decay is always constant. Regardless of whether the number of incidents is above or below the mean, the degree of persistence is quite large; the largest characteristic root of Equation 10 is 0.88.¹⁰

In contrast, the “skeleton” (i.e., the model without the error term) of the TAR model implies that there are two equilibrium states. In the high-terrorism state (i.e., when the number of incidents is 25 or more), $I_t = 1$, and Equation 11 becomes $cas_t = 17.87 + 0.189cas_{t-1} + 0.237cas_{t-2}$, for which the number of incidents gravitates toward the attractor 31.1 [= $17.87 \div (1 - 0.189 - 0.237)$]. As measured by the largest characteristic root, the speed of adjustment is 0.59: when terrorism is high, approximately 60% of each incident is expected to persist into the next period. In the low-terrorism state, $I_t = 0$ and Equation 11 becomes $cas_t = 3.92 + 0.423cas_{t-1} + 0.398cas_{t-2}$ with an attractor of 21.9. The largest characteristic root is 0.91, indicating very persistent behavior following a shock with little tendency to return to the attractor.

The different speeds of adjustment imply that the short-run forecasts from the two models will differ. As analyzed in Koop, Pesaran, and Potter (1996), the forecasts from a nonlinear model are state-dependent. In terms of Equation 11, a positive shock when $y_{t-2} > 25$ will be less persistent than the same shock when y_{t-2} is far below the threshold. For a model with 3 lags, we select a particular history for y_t , y_{t-1} , and y_{t-2} . For example, in the last three quarters of 1985 – a high-terrorism regime – the number of casualty incidents was 33, 50, and 40, respectively. Hence, to forecast the subsequent number of *cas* incidents from the perspective of 1985:4, we let $y_{t-2} = y_{1985:2} = 33$, $y_{t-1} = y_{1985:3} = 50$, and $y_t = y_{1985:4} = 40$. It is straightforward to verify that the linear model forecasts the values of $y_{1986:1} = 38.75$, $y_{1986:2} = 38.87$, $y_{1986:3} = 36.43$, $y_{1986:4} = 35.57$, and so on.

Since there is a possibility of regime switching, the multistep-ahead forecasts from the TAR model are more difficult to calculate. To employ Koop, Pesaran, and Potter’s methodology, we select 25 randomly drawn realizations of the residuals of Equation 11. Because the residuals may not have a normal distribution, the residuals are selected using standard “bootstrapping” procedures. In particular, the residuals are drawn with replacement using a uniform distribution.

Call these residuals $\varepsilon_{t+1}^*, \varepsilon_{t+2}^*, \dots, \varepsilon_{t+25}^*$. We then generate y_{t+1}^* through y_{t+25}^* by substituting these “bootstrapped” residuals into Equation 11 and setting I_t appropriately for high- or low-terrorist states. In essence, y_{t+1}^* is one possible realization of the *cas* series for 1986:1, y_{t+2}^* is one possible realization of the *cas* series for 1986:2, and so on. For this particular history, we repeat the process 1000 times. Under very weak conditions, the Law of Large Numbers (see Koop, Pesaran, and Potter 1996) guarantees that the sample average of the 1000 values of y_{t+1}^* converges to the conditional mean of y_{t+1} denoted by $E_t y_{t+1}$. Similarly, the Law of Large Numbers guarantees that the sample means of the various y_{t+i}^* converge to the true conditional i -step ahead forecasts, *i.e.*,

$$\lim_{N \rightarrow \infty} \left[\sum_{k=1}^N y_{t+i}^*(k) / N \right] = E_t y_{t+i}. \quad (12)$$

The essential point is that the sample averages of y_{t+1}^* through y_{t+25}^* yield the 1-step through 25-step ahead conditional forecasts of the *cas* series from the perspective of 1985:4. Intuitively, because the number of casualty incidents exceeds the threshold, the value of *cas* should quickly decline from 40 toward the attractor of 31.1. In contrast to the forecasts from the linear model, the TAR model forecasts the values $y_{1986:1} = 36.97$, $y_{1986:2} = 34.43$, $y_{1986:3} = 33.71$, and $y_{1986:4} = 32.45$.

The entire set of conditional forecasts is shown by the solid line in the top panel of Figure 1. Although the expected number of *cas* incidents does decline toward 31.1, there are two reasons why the long-term forecast continues to decline. Since incidents below the threshold are (on average) more persistent than those above, the system’s mean will be below the attractor. Moreover, the forecasts allow for the possibility of a regime-switch into the low-terrorist state. As shown in Panel *a*, the long-run forecast is about 28.5 *cas* incidents per quarter. It is interesting

to note that this long-run forecast is greater than the sample mean of 25.4. When the number of *cas* incidents is high, there is a rapid decline to the threshold, as terrorist networks cannot maintain high-level, resource-using offensives. A comparison of the forecasts with the actual number of casualty incidents (the dashed line in the figure) is instructive. The close fit is remarkable given that the forecasts are *not* 1 step-ahead forecasts, and, instead, traces out the 1-step through 25-step ahead forecasts.

In contrast, the number of terrorist incidents in the last three quarters of 1998 were quite low; $y_{1998:2} = 5$, $y_{1998:3} = 15$, and $y_{1998:4} = 6$. As shown in Panel *b* of Figure 1, reversion back toward the attractor of 21.9 is quite slow in the low-terrorism state. In fact, conditional on the history of 1998:4, the forecasts remain below 21.9 until the third quarter of 2001. The forecasts seem to track the actual number of incidents occurring through the end of our data set reasonably well and ultimately converge to those for Panel *a*.

Incidents with Deaths

Incidents with deaths are a major component of the *cas* series comprising 62% of such incidents. The *death* series is more complete than the series in Enders and Sandler (2002), because it runs from 1968:1 through 2000:4. The sample mean is 15.7 incidents per quarter. We first estimate the *death* series as the linear AR(2) autoregressive process:

$$death_t = 4.54 + 0.406death_{t-1} + 0.314death_{t-2} + \varepsilon_t; \quad AIC = 1100.47. \quad (13)$$

(3.58) (4.76) (3.69)

All *t*-statistics are significant at conventional levels and the point estimates of the autoregressive coefficients imply stationarity. The Ljung-Box *Q*-statistics also indicate that the residuals are serially uncorrelated. For example, the *Q*-statistics using the first 4, 8, and 12 lags of the residual autocorrelations have *prob*-values of 0.75, 0.24, and 0.10, respectively.¹¹

Regardless of whether the number of incidents is above or below the mean, the degree of

persistence is quite large; the largest characteristic root of Equation 13 is 0.80. Hence, the linear model indicates relatively slow convergence to the mean.

We start our search for the most appropriate TAR model by estimating the *death* series in the form of Equations 7 and 8 with $p = 2$. Eliminating the coefficients (except the intercept) with prob-values in excess of 10%, we obtain the following TAR (1) model:

$$death_t = [\underset{(19.76)}{20.90}] I_t + [\underset{(1.24)}{1.86} + \underset{(4.49)}{0.582} death_{t-1} + \underset{(3.88)}{0.379} death_{t-2}] (1 - I_t) + \varepsilon_t; \quad (14)$$

$$AIC = 1095.39, \tau = 22, \text{ and } d = 1.$$

The first eight autocorrelations of the residuals are less than 0.14 in absolute value and the *prob*-values for the Ljung-Box $Q(4)$, $Q(8)$, and $Q(12)$ statistics are 0.460, 0.375, and 0.063, respectively. Moreover, the AIC selects the TAR model over the linear model.

The difference between the threshold and linear models for the *death* series is quite pronounced. In the high-terrorism regime (i.e., when the number of incidents is 22 or more), $I_t = 1$, and Equation 14 becomes $y_t = 20.90$, owing to a zero characteristic root. Whenever the number of incidents exceeds 22, there will be an immediate decline to 20.90 incidents in the subsequent quarter, so that there is no persistence for incidents to exceed the threshold. A comparison of Equations 11 and 14 indicates that the *death* series has less persistence in the high-terrorism state than the *casualties* series.¹² Because a typical *death* incident is more difficult to execute than one with casualties (each incident of the *death* series is necessarily included in the *cas* series), the reduced persistence for *death* in the high-terrorism state is reasonable.

In the low-terrorism state, Equation 14 becomes $y_t = 1.86 + 0.582y_{t-1} + 0.379y_{t-2}$, with a large characteristic root of 0.97, consistent with near random-walk behavior. When the number of incidents is below the threshold value of 22, there is little tendency to return to a long-run mean value. In the low-terrorism state, the *death* series shows even less tendency to return to the

mean than the *cas* series. This finding poses a real concern during the pre- and post- 9/11 era of low terrorism, because the enhanced lethality of terrorist events of the post-Cold War era, found by Enders and Sandler (2000), is anticipated to persist.

The two forecast functions for the *death* series are shown Figure 2. In the last two quarters of 1985, death incidents numbered 33 and 28, respectively, indicative of a heightened-terrorism state. In order to forecast the subsequent number of *death* incidents from the perspective of 1985:4, we let $y_{t-1} = y_{1985:3} = 33$ and $y_t = y_{1985:4} = 28$ and apply the earlier methodology. In the mid-1980s high-terrorism state, the expected quarterly number of *death* incidents immediately falls from 28 to approximately 21; however, the long-term forecasts continue to decline. Because incidents below the threshold are far more persistent than those above it, the long-run forecast of 17.6 for the number of *death* incidents per quarter will be below the attractor. Panel *a* of Figure 2 shows that these 1-step through 25-step ahead forecasts track the actual number of *death* incidents quite well.

In 1998:3 and 1998:4 the quarterly number of incidents with deaths were 14 and 5, respectively, for the low-terrorism state. The forecast function for this history is shown in Panel *b* of Figure 2 where the number of such incidents for 1999:1 is expected to be about 10. Thus, from the perspective of 1998:4, the forecasted number of *death* incidents is anticipated to gradually rise until the long-run forecast of 17.6 is reached.

5. TAR Estimates of Other Series

In this section, we estimate how the key subcomponents of the *cas* and *death* series behave over time. Moreover, other incident types may be politically and/or economically important even when they do not always entail deaths or casualties. A credible threat may entail costly prevention. Although an assassination may fail, the attempt may provoke the same

political intimidation and utilize as many resources as a successful one.

A problem is that some of these series are somewhat thin. For example, the number of bombings with deaths has a mean of 7.1 incidents per quarter and there are two zero values near the beginning of the sample period. Because we use count data and the number of incidents cannot be negative, the assumption of normality is violated. One alternative is to perform the estimation assuming that the series is generated from a Poisson or negative binomial distribution, which forces the number of incidents to be non-negative. Experimentation with these specifications yield results that are in accord with those stated in Cameron and Trivedi (1998, p. 89): “Nevertheless, OLS estimates in practice give results quantitatively similar to those for the Poisson and other estimators using the exponential mean.” Since the models estimated using the Poisson are more difficult to interpret, we report the results for the AR and TAR models assuming normality.

Bombings with Deaths

In order to determine how an important subcomponent of the *death* series behaves, we examine the number of bombings with deaths (the *bomb* series). The linear model is:

$$bomb_t = 3.45 + 0.300bomb_{t-1} + 0.226bomb_{t-2} + \varepsilon_t; \quad AIC = 1010.95. \quad (15)$$

(4.50) (3.50) (2.63)

The residual autocorrelations are less than 0.17 in absolute value. Moreover, the Ljung-Box $Q(4)$, $Q(8)$, and $Q(12)$ statistics have *prob*-values in excess of 0.15, so there is little evidence of serially correlated residuals.

After we pared down the insignificant coefficients, we obtain the following TAR model:

$$bomb_t = [9.60] I_t + [2.22 + 0.439bomb_{t-1} + 0.322bomb_{t-2}] (1 - I_t) + \varepsilon_t; \quad (16)$$

(12.62) (2.27) (2.79) (3.38)

$AIC = 1009.80$, $\tau = 11$, and $d = 1$.

As measured by the AIC, the TAR specification has a better fit than the linear model. In the

high-terrorism state ($bomb_{t-1} \geq 11$), the quarterly number of incidents returns immediately to 9.60 after a shock; in the low-terrorism state, the quarterly number of incidents returns very gradually to the attractor after a shock. Given that the *bomb* series comprises 45% of the *death* series, the similarity of the TAR estimates is reasonable with the threshold of 11 for *bomb* being half that of *death*. Analogously, the high-terrorism intercept of the *bomb* series is just over 45% of that of the *death* series. In the low-terrorism state, the largest characteristic root for *bomb* is 0.83. As such, there tends to be more persistence in the low-terrorism state for *deaths* than for *bomb* following shocks.

Assassinations

Like the *bomb* series, the *assassination* series (*as*) is thin. The long-run mean number of assassinations was 8.1 incidents per quarter; there were no incidents in 1998:3 and the consecutive quarters of 1998:4, 1999:1 and 1999:2 each had a single incident. Standard lag length tests suggested an AR(3) specification. After eliminating an insignificant AR(2) coefficient, we obtain:

$$as_t = 1.99 + 0.461as_{t-1} + 0.303as_{t-3} + \varepsilon_t; \quad AIC = 1005.66. \quad (17)$$

(2.82) (5.94) (3.93)

The Ljung-Box $Q(4)$, $Q(8)$, and $Q(12)$ statistics have *prob*-values of 0.43, 0.09, and 0.13, respectively. All residual autocorrelations are less than 0.12 in absolute value except for $\rho_8 = 0.22$; however, estimates with eight lags do not yield a better fitting model. As such, Equation 17 seems to be a reasonable linear specification. The largest characteristic root of 0.83 suggests a slow rate of convergence to the long-run mean.

The TAR model with the best fit is:

$$as_t = [7.23 + 0.372as_{t-1}] I_t + [0.564 + 0.486as_{t-1} + 0.459as_{t-3}] (1 - I_t) + \varepsilon_t; \quad (18)$$

(5.68) (3.92) (0.622) (3.18) (3.10)

$$\text{AIC} = 1003.79, \tau = 9, \text{ and } d = 2.$$

The attractor in the high-terrorism state ($as_{t-2} \geq 9$) is approximately 11.5 incidents per quarter. Insofar as the sole characteristic root for the high-terrorism state is 0.372, there is relatively rapid convergence to the attractor. The assassination persistence, though small, in the high-terrorism state may stem from two sources: (i) modest resource requirements and (ii) scale economies in planning. The latter means that shocks during an active regime may trigger a wave of murders before the series converges. In the low-terrorism state following 9/11, there is a near unit-root implying almost no decay.

In the last three quarters of 1985, assassinations numbered 15, 18, and 20, respectively. The 1-step through 25-step ahead quarterly forecasts for the *assassinations* series are shown in Panel *a* of Figure 3 for this high-terrorism era. The forecasts quickly decay to the long-run value of approximately 10 incidents per quarter. In contrast, the last three quarters of 1998 saw 3, 2, and 1 assassinations, respectively. In Panel *b* of Figure 3, there appears to be no reversion of the forecasts back to a long-run value. Clearly, the low-assassination regime displays persistence as the long-run forecasts remain below the attractor of 10.25 past 2004.

Hostage Taking Incidents

Hostage taking incidents (hos_t) are well-estimated by the linear AR(2) model:

$$hos_t = 6.67 + 0.244hos_{t-1} + 0.241hos_{t-2} + \varepsilon_t; \text{ AIC} = 1094.33. \quad (19)$$

(5.01) (2.82) (2.77)

All of the residual autocorrelations are less than 0.15 in absolute value and the Ljung-Box $Q(4)$, $Q(8)$, and $Q(12)$ statistics for serial correlation have *prob*-values of 0.99, 0.98, and 0.64, respectively. The linear specification implies convergence to the long-run mean of 12.6 incidents

per quarter. The speed of convergence is estimated to be 0.63; thus, this linear specification for *hos* has a faster speed of adjustment than the linear specifications for the *cas* and *death* series.

When the *hos* series is estimated as a TAR process, we obtain:

$$hos_t = [15.59] I_t + [5.20 + 0.265hos_{t-1} + 0.328hos_{t-2}] (1 - I_t) + \varepsilon_t; \quad (20)$$

(18.52) (2.87) (1.45) (3.20)

AIC = 1091.64, $\tau = 15$, and $d = 1$.

When we purge the hos_{t-1} term from Equation 20, the AIC increases and there appears to be a significant autocorrelation coefficient at lag 1. As measured by the AIC, the TAR(1) model fits the data better than the linear specification. In the high-terrorism state ($hos_{t-1} \geq 15$), the skeleton predicts an immediate jump back to 15.59 incidents per quarter following a shock, so that elevated activity in the high-terrorism regime is unsustainable for more than a single quarter. In the low-terrorism state ($hos_{t-1} < 15$), the approach to the attractor is gradual; the largest characteristic root in the low-terrorism state is 0.72. Following the shock, the number of hostage-taking incidents can remain below the attractor for comparatively long periods of time. The *hostage* series behaves similarly in terms of its TAR representation to the *death* series, even though many hostage incidents do not result in a death. Over the sample period, the correlation coefficient between the *hostage* and *death* series is 0.536.

Threats and Hoaxes

A particularly interesting result is that the non-resource-using *threats/hoaxes* series (th_t) has precisely the opposite pattern of persistence than the other series. For *threats/hoaxes*, there is a good deal persistence in the high-terrorism state. The linear model is:

$$th_t = 4.36 + 0.290th_{t-3} + 0.314th_{t-4} + \varepsilon_t; \quad \text{AIC} = 1191.51. \quad (21)$$

(3.20) (3.57) (3.86)

All residual autocorrelations are less than 0.18 in absolute value and the Ljung-Box $Q(4)$, $Q(8)$,

and $Q(12)$ statistics have *prob*-values of 0.66, 0.51, and 0.38, respectively. Since the largest characteristic root of Equation 21 is 0.87, there is slow convergence to the long-run mean of 10.8 incidents per quarter.

When we estimate the model as a TAR and pare down the coefficients, we obtain:

$$th_t = [\underset{(3.54)}{4.38} + \underset{(2.84)}{0.404}th_{t-3} + \underset{(3.20)}{0.368}th_{t-4}] I_t + [\underset{(4.72)}{7.86}] (1 - I_t) + \varepsilon_t; \quad (22)$$

AIC = 1177.85, $\tau = 12$, and $d = 2$.

In the low-terrorism state (*i.e.*, $th_{t-2} < 12$), the point estimate of the skeleton indicates that the quarterly number of incidents will immediately jump to 7.86 following a shock. In contrast, the attractor for the high-terrorism regime is 19.2 incidents per quarter. Moreover, the speed of adjustment for the high-terrorism regime is very low, since the largest characteristic root is 0.93. This flip-flop in regime behavior for *threats/hoaxes*, compared with the other series studied, is surely due to the non-resource-using nature of these events. Shocks during high-terrorism regimes may be sustainable if heightened threats reinforce one another (*i.e.*, they are complementary) to create the desired state of anxiety for society. As such, there may be little diminishing returns to further threats and hoaxes.

The history for 1998:4 is such that $th_{1998:1} = 4$, $th_{1998:2} = 2$, $th_{1999:1} = 14$, and $th_{1999:2} = 6$. Insofar as $th_{1998:4}$ is less than the threshold value, there is a predicted jump to approximately 7.86 incidents per quarter. There are several plausible explanations why the *threats/hoaxes* series has drifted towards zero in recent years. The increase in religious fundamentalism has resulted in a greater willingness of terrorists to engage in behavior resulting in casualties and deaths (Enders and Sandler 2000; Hoffman 1998). The fact that the media now focuses only on the most severe events may have decreased the benefits terrorists derive from a threat or a hoax. Regardless of the best explanation, the world is clearly in a persistent regime in which threats and hoaxes are low.

The striking differences in the measures of persistence for each of the six series are summarized in Table 1. The sample means, thresholds and delay parameters are reported in columns 2 through 4 and the estimated attractors for the high and low terrorism regimes are shown in columns 5 and 6. Respectively, columns 7 through 9 report the largest characteristic root for the linear model, the low-terrorism regime, and the high-terrorism regime.

6. Concluding Remarks

For all six time series examined, a TAR model that allows for different autoregressive behavior during high- and low-terrorism regimes fits the data better than some linear AR model, which gives an average representation. For forecasting purposes, the TAR model leads to reasonably good forecasts for the *casualties*, *death*, and *assassinations* series. Because these forecasts vary significantly in terms of the persistence of incidents to shocks depending on the underlying terrorism regime, forecasts for such events should rely on the appropriate TAR representation. In the case of *threats/hoaxes*, the forecasts (not displayed), based on the TAR model, did not fit the actual series owing to a high degree of volatility that defies prediction. Forecasts appear to improve with the level of resources required of the underlying event. The development of these forecasting procedures can better allow the authorities to gauge terrorist reactions to shocks stemming from policies to crack down on terrorism or political events (e.g., a Middle East peace agreement).

The breakdown of the *all-incident* series into component series indicates a number of insights. First, the rate of convergence following shocks differs greatly between the aggregate series and its components, so that the behavior of an aggregate series is unlikely to characterize its component series. Second, the level of persistence following shocks tends to differ among component series themselves, so that generalizations must be resisted. Third, component series need not be in the same regime simultaneously – i.e., assassinations may be in a high-level

regime, while hostage taking is in a low-level regime. These can differ owing to substitutions induced by anti-terrorism policies: i.e., policy aimed at making one type of attack mode harder makes another attack mode *relatively* easier (cheaper).

The events of 9/11 and the subsequent “war on terror” came during a time when the number of incidents involving casualties was in a low-terrorism regime. Nevertheless, each terrorist event had a greater number of casualties on average than in past decades. Our analysis indicates that shocks following 9/11 would be met with heightened terrorist campaigns that can be sustained for some time. The low-terrorism regime era of religious terrorism means that shocks will be met with some persistence of heightened campaigns. Thus, the terrorist alert level which has been elevated since being instituted by the Bush administration is unlikely to fall in the near future owing to terrorists’ abilities to sustain campaigns during low-terrorism regimes.

Footnotes

1. The special threats and dilemmas of liberal democracies when confronting modern-day terrorism are addressed in Hoffman (1998) and Wilkinson (1986).

2. Support of the rational-actor representation of terrorists is provided by Sandler, Tschirhart, and Cauley (1983) and Landes (1978). Also, see cites within these articles.

3. The choice-theoretic model derives from Landes (1978) and Sandler, Tschirhart, and Caluley (1983). The introduction of an intertemporal choice in Equations 6 and 7 is novel for the terrorism literature and allows for high- and low-terrorism regimes.

4. $|H| = -(C_h)^2 U_{bb} + \lambda B_{bb} (C_h)^2 + 2C_h B_b U_{hb} - (B_b)^2 U_{hh} + \lambda C_{hh} (B_b)^2 > 0$, where subscripts indicate partial derivatives and λ is the Lagrangean multiplier associated with the resource constraint.

5. If the resource constraint is more convex-to-the-origin than the indifference curve, then a corner solution is anticipated.

6. For a fuller description of how ITERATE delineates transnational terrorist incidents, see the discussion in Mickolus, Sandler, and Murdock (1989).

7. It would be possible to estimate a multi-regime model; however, as the number of regimes increases, the number of observations in each regime decreases. Given our concern about thin series, we choose a two-regime model.

8. In the literature, there are several different ways to report the AIC; here we use: $AIC = T \ln(ssr) + 2n$ where T = number of usable observations, ssr = sum of squared residuals and n = number of parameters estimated (including the threshold). Note that the AIC is not a cardinal measure of fit and any monotonic transformation of this measure of AIC will select the same model.

9. Let ρ_i denote the i -th residual autocorrelation coefficient. Although the Q -statistics

allow us to reject the null hypothesis of serially correlated residuals, we are somewhat concerned that $\rho_7 = 0.21$. As such, we also estimate models using longer lag lengths. Estimating a model using y_{t-7} does not yield any substantial changes to the results reported below.

10. The characteristic roots are the values of r that satisfy: $r^3 - 0.261r^2 - 0.310r - 0.209$.

11. Although the Q -statistics allow us to reject the null hypothesis of serially correlated residuals, we are somewhat concerned that $\rho_7 = 0.21$. As such, we also estimate models using longer lag lengths with no substantial influence on the results reported in the text.

12. Given that we estimate each series as a univariate process, the *cas* series can be in the high (low) state, while the *death* series is in the (low) high state. To our knowledge there is no well-developed procedure to estimate TAR models in a multi-equation system.

References

- Brophy-Baermann, Bryan, and John A. C. Conybeare. 1994. Retaliating against terrorism: Rational expectations and the optimality of rules versus discretion. *American Journal of Political Science* 38:196–210.
- Cameron, A. Colin, and Pravin Trivedi. 1998. *Regression analysis of count data*. Cambridge: Cambridge University Press.
- Chan, K. S. 1993. Consistency and limiting distribution of the least squares estimator of a threshold autoregressive model. *The annals of statistics* 21:520–33.
- Enders, Walter, and Todd Sandler. 1993. The effectiveness of anti-terrorism policies: Vector-autoregression-intervention analysis. *American Political Science Review* 87:829–44.
- Enders, Walter, and Todd Sandler. 1995. Terrorism: Theory and applications. In *Handbook of Defense Economics, Vol. 1*, edited by Keith Hartley and Todd Sandler. Amsterdam: North-Holland, pp. 213–49.
- Enders, Walter, and Todd Sandler. 1999. Transnational terrorism in the post-cold war era. *International Studies Quarterly* 43:145–67.
- Enders, Walter, and Todd Sandler. 2000. Is transnational terrorism becoming more threatening. *Journal of Conflict Resolution* 44:307–32.
- Enders, Walter, and Todd Sandler. 2002. Patterns of transnational terrorism, 1970-99: Alternative time series estimates. *International Studies Quarterly* 46:145–65.
- Faria, Joao. 2003. Terror cycles. *Studies in Nonlinear Dynamics and Econometrics* 7, Article No. 3, <<http://www.bepress.com/snede>>.
- Hansen, Bruce. 1997. Inference in TAR models. *Studies in Nonlinear Dynamics and Econometrics* 2:1–14.
- Hoffman, Bruce. 1998. *Inside terrorism*. New York: Columbia University Press.

- Koop, Gary, M. Hashem Pesaran, and Simon Potter. 1996. Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics* 74:119–47.
- Landes, William M. 1978. An economic study of US aircraft hijackings, 1961-76. *Journal of Law and Economics* 21:1–31.
- Mickolus, Edward F., Todd Sandler, and Jean M. Murdock. 1989. *International terrorism in the 1980s: A chronology of events*, 2 vols. Ames, IA: Iowa State University Press.
- Mickolus, Edward F., Todd Sandler, Jean M. Murdock, and Peter Flemming (2002), *International terrorism: Attributes of terrorist events, 1968-2000* (ITERATE 5). Dunn Loring, VA: Vinyard Software.
- Sandler, Todd, and Walter Enders. 2004. An economic perspective on transnational terrorism. *European Journal of Political Economy* 20:forthcoming.
- Sandler, Todd, John T. Tschirhart, and Jon Cauley. 1983. A theoretical analysis of transnational terrorism. *American Political Science Review* 77:36–54.
- Tong, Howell. 1983. *Threshold models in non-linear time series analysis*. New York: Springer-Verlag.
- Tong, Howell. 1990. *Non-linear time series: A dynamical approach*. Oxford: Oxford University Press.
- Wilkinson, Paul. 1986. *Terrorism and the liberal state*. Revised edition. London: Macmillan.

Figure 1: Nonlinear Forecasts of Casualty Incidents

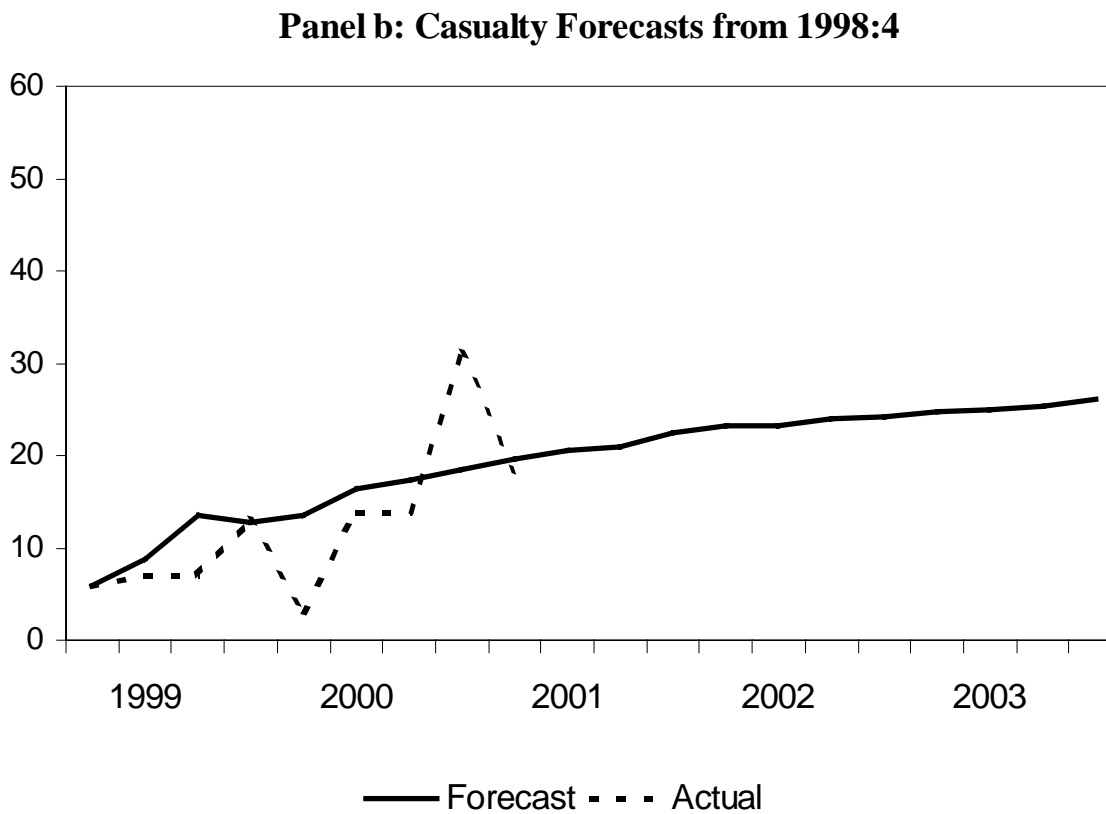



Figure 2: Nonlinear Forecasts of Death Incidents

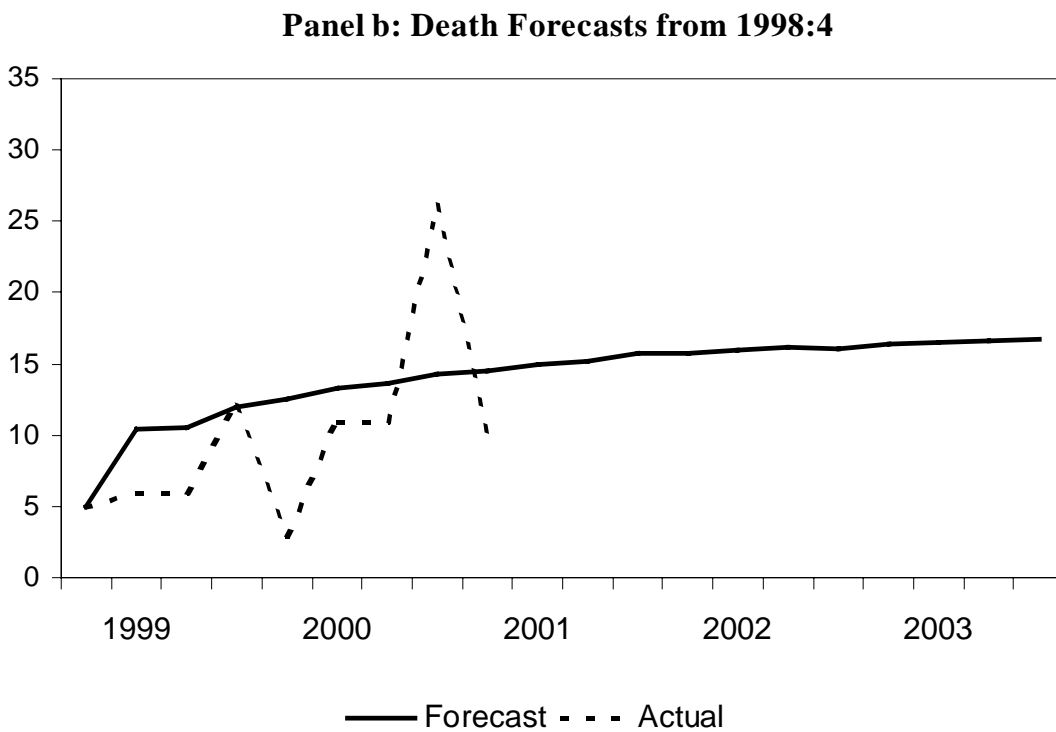





Figure 3: Nonlinear Forecasts of Assassination Incidents

Panel a: Assassination Forecasts from 1985:4



Panel b: Assassination Forecasts from 1998:4



— Forecast - - - Actual

Figure 4: Nonlinear Forecasts of Threats and Hoaxes

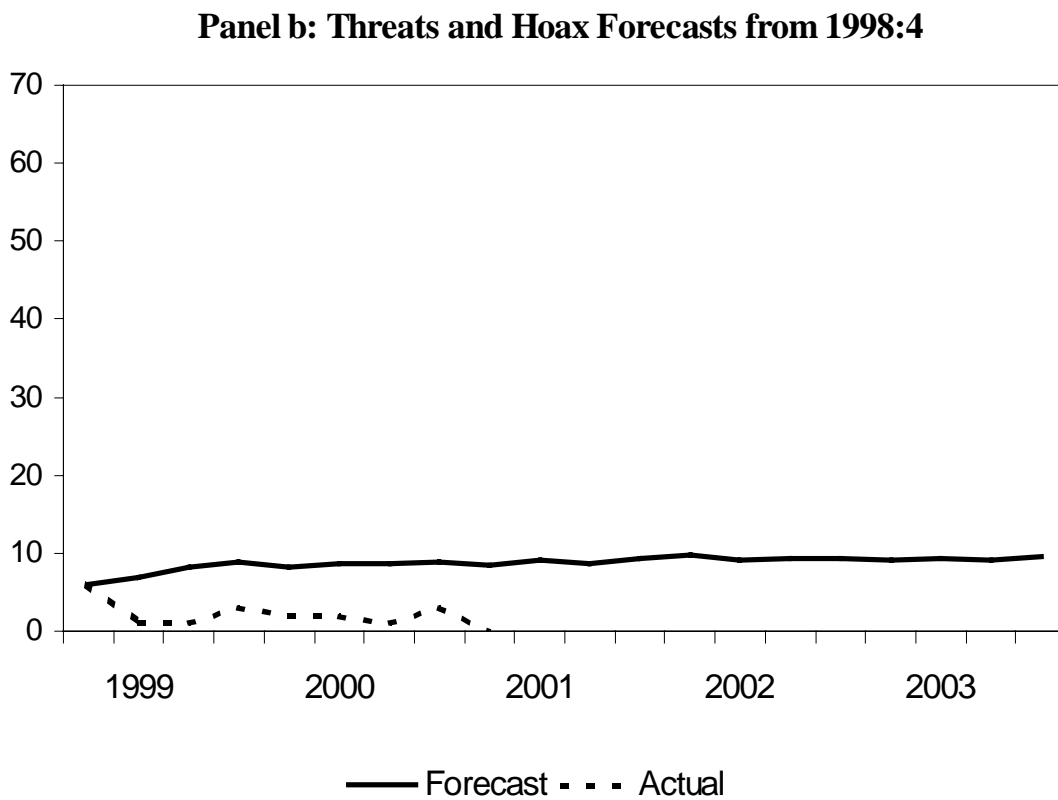
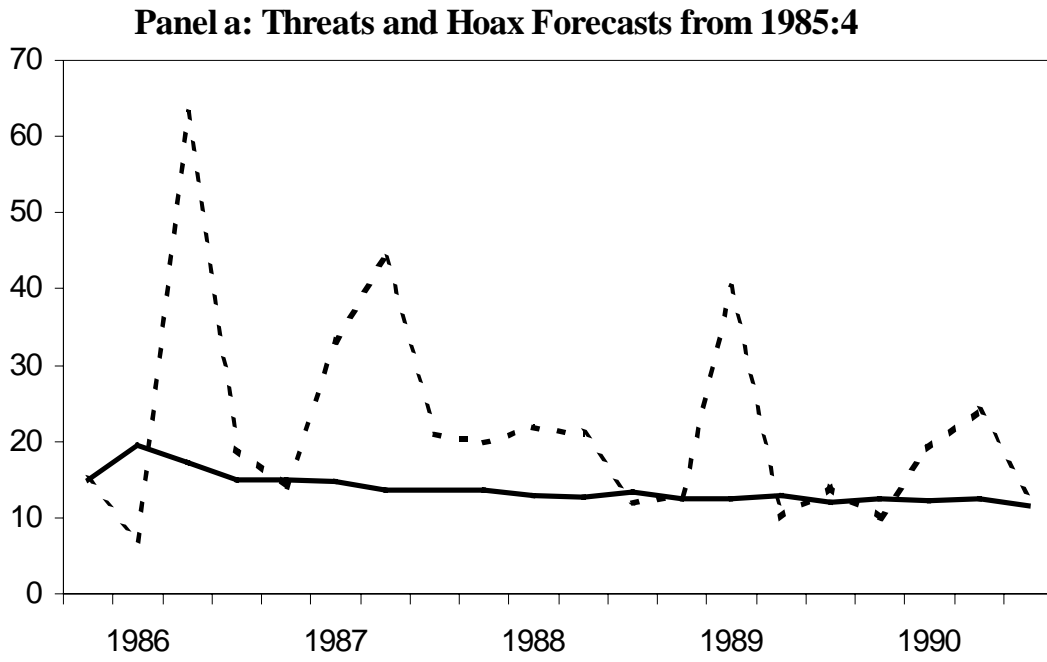


Table 1: Statistical Properties of the Incident Series

Series	Mean	τ	d	Attractor		Largest Characteristic Root		
				Low	High	Linear	Low	High
<i>cas</i>	25.4	25	2	21.9	31.1	0.88	0.91	0.59
<i>death</i>	15.7	22	1	NA ^a	20.9	0.80	0.97	0.00
<i>bomb</i>	7.1	11	1	9.29	9.60	0.65	0.83	0.00
<i>as</i>	8.1	9	2	NA ^a	11.5	0.83	0.97	0.37
<i>hos</i>	12.6	15	1	12.78	15.59	0.63	0.72	0.00
<i>th</i>	10.8	12	2	7.86	19.2	0.87	0.00	0.93

^a NA refers to the fact that the estimated attractors are not applicable since the series behaves as a unit-root process in the low-terrorism regime.