A Conditional Probability Analysis of Pattern-Based Models Applied to Event Data in the Israeli-Palestinian Conflict

by

Philip A. Schrodt, University of Kansas

Valerie M. Hudson, Brigham Young University

Version 1.0

27 August 2006

Paper prepared for delivery at the Annual Meeting of the American Political Science Association, Philadelphia, 31 August-3 September 2006.

Authors can be contacted via email at the following addresses: `valerie_hudson@byu.edu` and `schrodt@ku.edu`.

This research was supported in part by grants from the U.S. National Science Foundation (SES-0455158),and the Brigham Young University College of Family, Home, and Social Sciences.

A link to the analytical web site for this project, which includes the graphic tools for analyzing event patterns can be found at the NKSS project web site `http://www.nkss.org`. The data set discussed in this paper, as well as a pdf version of the paper with color graphics, can be downloaded from the KEDS project web site: `http://www.ku.edu/~keds/`.

Abstract

Existing formal models of political behavior have followed the lead of the natural sciences and generally focused on methods that use continuous-variable mathematics. Stephen Wolfram has recently produced an extended critique of that approach in the natural sciences, and suggested that a great deal of natural behavior can be accounted for using rules that involve discrete patterns. Over the past three years we have developed software to display the presence of patterns in event data. This paper extends our earlier work, which focused on the presence of specific patterns over time, to look at the stochastic characteristics of these pattern. We are specifically interested in the relationship between patterns as reflected in their conditional probabilities: does the probability of a pattern increase or decrease the probability of other patterns? Using data from the Israeli-Palestinian conflict for the period April 1979 to October 2005, and using some of the common patterns that we identified in our earlier research, we find that these conditional relationship do hold, and are almost always positive. As we expected, and consistent with our earlier studies, the strength of the relationships varies over time and these variations are usually strongly correlated with the broader qualitative characteristics of the conflict, for example major conflict phases such as the first and second *intifadas* and the Oslo negotiation process.

Introduction and Background

In 2002, a methodological gauntlet was thrown down by Stephen Wolfram in his work, *A New Kind of Science*. Though his book was not written from or for a social science perspective, several of his assertions are pertinent to that endeavor. Wolfram asserts that most modern scientific methods used in the physical and biological sciences are but idiosyncratic and limited derivations from something much more basic, more fundamental, and more powerful. In place of the continuous-variable mathematical structures that underlie classical mechanics and statistics, Wolfram's approach focuses on the discrete transformation of patterns. Simple pattern-based models can, through iteration, produce surprisingly complex behavior in physical and biological systems. Biochemists, for example, search for patterns in amino acids as elements for understanding the functions of a strand of DNA, and then the patterns of those strands combine to produce the patterns formed by larger strands, then by chromosomes, then by the entire genome. Though the patterns themselves are simple, they can ultimately produce highly complex organisms, including human beings themselves.

Conveniently for social scientists, humans do not only originate from patterns, but human psychology is intensely linked to the ability to perceive patterns and to find meaning in patterns (Newell and Simon 1972, Abelson 1973, Simon 1982, Anderson 1983, Kohonen 1984, Holland et al 1986, Margolis 1987, Khong 1992, Reber 1993, *Political Psychology* 2003). Indeed, it is not far off the mark to suggest the ultimate basis of all human epistemology is discrete pattern identification. As Wolfram puts it, "observers will tend to be computationally equivalent to the systems they observe," (Wolfram, 2002, 737) an observation we will expound upon shortly.

Hudson, Schrodt and Whitmer (2004) was an initial descriptive validation of the potential of this approach. Since no one had looked for patterns in this fashion before, we first needed to demonstrate that we could find them, and that the patterns had some plausible correspondence to our underlying qualitative understanding of the situation we were analyzing. In that research, we developed a web-based tool for exploring pattern-based rules; this can be found at `http://kennedyosx.byu.edu/`. That site includes data from the Kansas Event Data System (KEDS) project, and provides a number of well-documented facilities for recoding the data, specifying rules, and visualizing event data as discrete patterns rather than scaled aggregations. In particular, the inputs titled “patterns” and “display” allow a researcher to have the capability to perform discrete pattern transformations on the graphic output. One can also experiment with possible rules, then display whether those patterns account for any of the behavior in the set.

In our initial probe of the approach, we specified some very simple rules and then ascertained how well they accounted for the behavior in the Israel-Palestine dyad. These rules were chosen from a combination of the general theoretical literature and a qualitative assessments of what some experts in the field assert are the rules these actors do use (e.g. Bickerton and Klausner 1998, Gauss 1998, Gerner 1994, Goldstein et al 2001, Tessler 1994).

Wolfram himself provides encouragement that the rules need not be many, and neither do they need be complex. For example, he states, “Simple and definite underlying rules can produce behavior so complex that it seems free of obvious rules” (Wolfram, 2002, 752) and then goes on to elaborate that in his years of experience analyzing complex systems,

> But when in general does complexity occur? [I]f the rules for a particular system are sufficiently simple, then the system will only ever exhibit purely repetitive behavior. If the rules are slightly more complicated, then nesting will also often appear. But to get complexity in the overall behavior of a system one needs to go beyond some threshold in the complexity of its underlying rules. The remarkable discovery that we have made, however, is that this threshold is typically extremely low. [I]t ultimately takes only very simple rules to produce behavior of great complexity. . . . Instead, once the threshold for complex behavior has been reached, what one usually finds is that adding complexity to the

> underlying rules does not lead to any perceptible increase at all in the overall complexity of the behavior that is produced. (Wolfram, 2002, 105-6)

Indeed, Wolfram found that the most complex behavior could be obtained with the use of approximately three rules. We feel that there is reason to believe that the set of rules being employed by the Israelis and Palestinians in enacting what they feel to be meaningful behavior toward one another is also not very large, nor very complex. Signaling between organized human collectives, especially those in conflict, almost mandates that only a small set of simple rules be used in order to maximize the chances that the other group will understand the meaning intended by the action.

Furthermore, because international politics is a complex, ill-structured problem solving environment, heuristics—simple rules used to partially solve complex problems—are of particular importance. Purkitt observes:

> To cope with limited cognitive capabilities, individuals selectively process information and use a limited number of heuristics or mental rules of thumb as cognitive aids in their effort to manage information. This apparently universal reliance on intuitive heuristics to solve all types of problems seems to be due to the need to compensate for the limitations of short-term memory and information processing capabilities. By using intuitive mental heuristics, people can develop a problem attack plan which permits them to develop a subjectively acceptable problem solution. (Purkitt 1991,43)

For example, rational choice and balance of power theories are both heuristics in the sense that they are relatively simple; they come with a complex set of side-conditions; and they are intended as general rules to guide decision-making, without providing a complete specification of actions to be taken. To the extent that an heuristic is shared by the decision-makers in a political system—for example balance of power in 19th century European diplomacy or the Chicken game in 20th century nuclear deterrence—it reduces uncertainty and becomes self-validating.

In our exploratory exercise (Hudson, Schrodt and Whitmer 2004), we endeavored to come up with a small set of fairly simple rules that could be justified on the basis of scholarship concerning Israeli and Palestinian actions. The first rule we used was the classic "tit-for-tat" (TFT) approach immortalized by Rapoport and, more recently, Axelrod (1984). Country experts have asserted that the Israelis and Palestinians consciously use this rule; and it has long been known that reciprocity is one of the strongest patterns in event data (for example Dixon 1986, Ward and Rajmaira 1992, Goldstein and Freeman 1992, Goldstein and Pevehouse 1997). The second rule we used was one that we have labeled the “olive branch”: one side responds to a period of conflict with cooperation rather than reciprocating the conflict. The olive-branch rule is the standard gambit for breaking out of the mutually-destructive DD/DD/.../DD sequence in the classical prisoners’ dilemma game. Finally, we looked at four more complex “meta-patterns” that involved patterns-of-patterns—that is, complex patterns that were built out of the occurrence of simpler pattern. These meta-patterns were designed to tap into escalation and de-escalation behavior that was more complex than the simple “olive branch.”

This analysis produced a rich set of results. For example, we found three general results on TFT, the simplest of our rules. First, the TFT behaviors are generally, but not totally, symmetric in the dyad—generally when one side is engaging in TFT, whether cooperative or conflictual, the other side is doing so as well. There is no reason that this must be the case, but the fact that we observe it suggests that the two antagonists are implementing a classical TFT solution to the prisoners’ dilemma game. Unsurprisingly, give our qualitative understanding of the conflict, they are far more likely to be playing DD than CC.

Second, most of spikes in the conflictual TFT correspond to periods of substantial violence such as the first and second *intifadas* and Israel’s 1982 invasion of Lebanon. The outbreak and decline of the first *intifada* from December 1987 to August 1990 shows the same exponential-decay shape that is seen in Goldstein-scaled data for the period (Schrodt and Gerner 1994). Similarly, the negotiations following the Oslo agreement in September 1992 and prior to the outbreak of the second *intifada* in September 2000 are evident.

The most surprising aspect of the TFT analysis was the juxtaposition of TFT conflict *and* cooperation during the post-Oslo period. We cross-checked this against the qualitative record and found that pattern to be a good illustration of the utility of objective events patterns over vaguely remembered narratives: While the Oslo period (1994-2000) saw nowhere near the levels of violence seen in the second *intifada,* there were periods of substantial conflict, such as the four suicide bombings in Tel Aviv and Jerusalem and subsequent Israeli reactions to these in the spring of 1996, shortly after Israel’s military withdrawal from Palestinian urban areas. Conversely, negotiations have continued at both the official and unofficial levels (e.g. the December 2003 Geneva Accords between Israeli and Palestinian citizen elites) during the second *intifada.*

This initial analysis was extended in Hudson, Schrodt and Whitmer (2005), which used the same set of tools (with minor improvements to make the site more user friendly) but “drilled deeper” into the event data. First, we switched from the nation-state level of aggregation found in most existing event data studies to a sub-state actor focus that used the new CAMEO sub-state actor coding scheme (Gerner et al 2005). This more detailed approach is readily accommodated by the analytical site, and as we expected, we found significant sub-state actor differences in many of the event streams—for example between the Israeli government and Israeli settlers, between the mainstream Palestinian Fatah and (post-Oslo) Palestinian Authority and Islamic militant groups, and between Palestinians in the West Bank and those in Gaza. Second, we analyzed changes in the pattern frequencies over time, and in particular compared these to changes in the Israeli government (prior to the death of Arafat, Palestinian leadership was essentially unchanged). Again, many of these government changes corresponded to discernible changes in pattern frequencies.

One of our concerns when we embarked on the analysis was whether we would posit plausible patterns and find nothing in the data. Our experience has, instead, been the opposite— the problem is not that we are finding too little, but we are still finding too much. When one combines the remarkably rich set of patterns that can be constructed using the quite simple methods aggregation methods available in the pattern-specification language with the ability to rapidly construct colorful, web-based displays at a very fine time interval, it is difficult to figure where to go next with the analysis. On the other hand, with a few exceptions, we are finding very credible “patterns in the patterns”—these do not occur at random, but instead their rise and fall generally tracks changes in the political situation which we know about from qualitative narratives.

In this paper, we are continuing to look at patterns, but we will be backing off from the high level of detail that we pursued in Hudson, Schrodt and Whitmer (2005) in three ways. First, instead of looking at the occurrence of patterns at a very fine level of temporal detail—the two earlier papers used a temporal resolution of two days—we will look at the aggregate behavior in 120-day blocks of time. This still allows us to look at changes in the patterns over time, which clearly are one of the most important features of the data, but allow these to be viewed in a single chunk. Second, we will be using conventional probability theory—specifically conditional probability—as the means of doing this aggregation. Notice that the use of probability, an interval-level measure, does not imply that we are shifting away from the discrete, categorical approach: we are still looking at patterns; we have merely shifted to considering the probabilities of those patterns.[1] Finally, in this initial analysis of the

[1] From the perspective of social science research, one of the frustrating aspects of Wolfram’s work is the absence of any stochastic element: Wolfram’s models are entirely deterministic. Wolfram argues, with quite a bit of evidence, that *many* patterns that appear to be random can in fact be generated by simple deterministic processes. However, it does not follow from this that *all* patterns, particularly those resulting from systems that are clearly highly complex (human social interactions presumably qualify), can be so reduced. The patterns found in event data contain a further stochastic component in that not all “events” are reported in the international media, some of those events are incorrectly coded, and the coding system itself is an artifact. Consequently event data must be

data, we shift back to the nation-state level: this gives us higher event counts and consequently a larger number of instances where the conditional probability can be computed.

Israel-Palestine Event Data

In this paper, we will continue to explore the possibilities and pitfalls of Wolfram's approach to understanding human behavior in the arena of international politics. As before, we are using event data on the Israel-Palestine conflict. This dyad involves actors whose behavior, while certainly affected by the initiatives of third parties, is highly interactive and has been the focus of sufficient media attention that we can be confident that the event data are a reasonably accurate description of the actual behavior in the system.

News reports on the interactions between Israel and Palestine were coded into the CAMEO scheme (Gerner et al 2002, 2005), an event data coding scheme that is similar to the classic WEIS (McClelland 1976) system but combines WEIS categories that were either ambiguous or difficult to distinguish using automated methods, and provides greater detail (which we will not be using in this exercise) in its sub-categories. A complete list of the verb phrases used to code CAMEO can be found at `http://kennedyosx.byu.edu/data/mideast.cameocodes_verbs.html`; a PDF copy of the manual is available at `http://www.ku.edu/~keds/data.dir/cameo.html`.

The data were coding using TABARI (`http://web.ku.edu/keds/software.dir/tabari.html`), a computer program from the Kansas Event Data System (KEDS) project that creates event data from machine-readable text.[2] The events were coded from Reuters News Service lead sentences obtained from the NEXIS data service for the period April 1979 through May 1997, the Reuters Business Briefing service for June 1997 through December 1998, and *Agence France Presse* (AFP) from January 1999 to November 2005. The data were run through a “one-a-day” filter to eliminate duplicate reports of the same event by allowing only one instance of any source-event-target combination in a day. The coding software, coding dictionaries and data are available at the KEDS web site, `http://www.ku.edu/~keds`.

The pattern definitions use the following general classes of events based on the CAMEO two-digit cue categories:

<u>verbal cooperation</u>: 01 to 05

<u>material cooperation</u>: 06 to 09

<u>verbal conflict</u>: 10 to 14

<u>material conflict</u>: 15 to 20

This reduces the number of distinct event categories to a manageable amount and is also likely to reduce the effects of coding error somewhat, since only broad categories of events are being considered.

The two actors we are looking at are the Palestinians (PSE and PAL[3]) and Israel (ISR). While CAMEO provides for extensive sub-state actor codes—and the default display differentiates these—in this analysis we are only looking at the most general level of “Palestinians” and

analyzed in a framework that allows for randomness, which Wolfram’s framework, if applied literally (which we have not done), does not provide.

[2] Discussions of machine coding can be found in Gerner et al. (1994), Schrodt and Gerner (1994), Huxtable and Pevehouse (1996), and Bond et al. (1997, 2003), and King and Lowe 2004.

[3] In the CAMEO coding scheme, “PSE”—the ISO designator for the Palestinian Authority—refers to Palestinians in the post-Oslo environment, whereas “PAL” refers to the pre-Oslo situation when there was not mutual recognition between a Palestinians governing entity and Israel. This is not an entirely satisfactory coding solution and may be altered in the future; we are combining the two sets of codes so it has no effect on this analysis.

"Israelis." This complicates the interpretation of some of the results, particularly in the post-2000 period of the al-Aqsa *intifada* because of the activity of Palestinian militant groups that are explicitly opposed to the policies of the Palestinian authority. To a much more limited degree, it also affects Israel, notably in the actions of Israeli settlers. We have used this aggregation in order to increase the number of cases where we can actually estimate the probabilities (and increase the number of events that are used to estimate those probabilities). We may look at some specific groups in the future: the pattern display can be very easily adapted to do this.

Patterns

The analysis in this paper builds on the patterns that we developed earlier in Hudson, Schrodt and Whitmer (2004, 2005); these are found in the default event display `http://kennedyosx.byu.edu/`. The patterns are built up from the individual events using a small but powerful set of basic functions that do basic operations such as counting events, lagging, summing event counts over a period of time, and evaluating logical conditions. These operations are described on the web page `http://kennedyosx.byu.edu/Event_Patterns_Guide.html`, which provides complete documentation for the system. The web-based tool itself was developed by Ray Whitmer of JHAX Ltd and is available for general use on our web server.

Once a set of patterns has been created, a graphical display is generated that shows when the various patterns occur over time. Figures 1a and 1b show examples of this display configured for aggregate material conflict (that is, PSE/PAL and ISR, rather than the substate actor limitations used in the default display) for two periods, 28 September 1995 to 28 December 1995, which corresponded to the initial implementation of the Israeli withdrawal from Palestinian urban areas as part of the implementation of the Oslo agreements, and 11 July 2000 to 10 January 2001, which corresponds to the outbreak of the al-Aqsa *intifada* in September 2000.

Each line of the display corresponds to a two-day period, so moving down the display corresponds to moving forward in time. The symbols correspond to the occurrence of various patterns: for example on the left side of the display, green open squares are Palestinian verbal cooperation directed towards Israel, solid blue triangles are Israeli material conflict directed to Palestinians; in the middle columns red and yellow squares correspond to conflict and cooperative tit-for-tat, and on the right side the various solid colored boxes correspond to complex "meta-patterns" that are composites of the simpler patterns. All of these patterns are discussed in detail in our earlier paper. A wide variety of graphical characteristics, including the placement, shape, color, transparency, and the direction of placement of multiple instances of a symbol can be modified within the display.

Following Wolfram's paradigm, our intention in developing the display was three-fold. First, we wanted to demonstrate that event data—which are a nominal time series—could be displayed as nominal data rather than, as is typical, artificially aggregated into interval-level series using a scaled such as that developed by Goldstein (1992). By preserving the distinctions made in the nominal data one can see, for example, that while both the first and second *intifadas* involve substantial levels of violence, in the second, al-Aqsa *intifada*, this is accompanied by very substantial verbal cooperation—a legacy of the Oslo process—which was absent in the first *intifada*.

Second, closely following Wolfram, we wanted to demonstrate that patterns do exist in the event data, and these occur non-randomly: this was discussed in detail in Hudson, Schrodt and Whitmer (2004). Finally, we wanted to show that changes in those patterns corresponded to plausible changes in the underlying political process, for example changes in the Israeli government, in the negotiation process, and in outbreaks of violence such as Israel's 1982 invasion of Lebanon and the two *intifadas*. This was done in Hudson, Schrodt and Whitmer 2005.

In this paper we take a different approach that links the nominal data to an interval value. We continue to work with patterns, but rather than deal with these visually, we look at them probabilistically. One of the options in the event display allows the generation of a tab-delimited text file that shows how frequently the pattern occurs: in other words, instead of generating symbols, it generates a series of numbers. This file is processed by a simple C program to compute marginal and conditional probabilities, which are then plotted over time.

We will examine the following patterns (all of these are instantiated in the default display)

Conflict events: CAMEO material or verbal conflict events as defined above

Cooperation events: CAMEO material or verbal cooperation events as defined above

XY Conflict tit-for-tat (TFT): in the period [t-14, t-8], there is at least one 8-day period when conflict behavior from Y to X exceeds the conflict threshold level, and in the period [t-8, t], conflict behavior from X to Y exceeds the conflict threshold level.

XY Cooperative tit-for-tat (TFT): in the period [t-14, t-8], there is at least one 8-day period when cooperation behavior from Y to X exceeds the cooperation threshold level, and in the period [t-8, t], cooperation behavior from X to Y exceeds the cooperation threshold level.

XY olive branch: in the period [t-14, t-8], there is at least one 8-day period when conflict behavior from Y to X exceeds the conflict threshold level, and in the period [t-8, t], cooperation behavior from X to Y exceeds the cooperation threshold level.

Because the data set reports substantially more conflict events from Israel to Palestinians than vice versa—a feature that could either be due to activity on the ground or the fact that most Israeli actions are due to institutionalized agents that are easier for the international media to report systematically—the conflict threshold was set to 2 for Palestinians and 4 for Israel. The cooperation threshold was set to 1 for both sides—cooperative behavior isn't real common in this dyad. The lag times for the TFT are a bit arbitrary and simply reflect our sense of the period of time that typically elapses between an action and response in the conflict. All of these parameters can be easily changed in the pattern display; we've done some experimentation and the results we are seeing do not seem to be highly sensitive to these choices. We looked at verbal and material events separately, so each of these behaviors generates two graphs.

Conditional Probability Analysis

The analysis in this paper will look at the *conditional* probabilities of pairs of patterns, Pr(Y|X). Specifically, we are interested in the *difference* between Pr(Y|X) and Pr(Y)[4]: a positive value of this difference indicates that Y is more likely to occur in the presence of X; a value near zero implies that the two behaviors are unrelated; and a negative value indicates that Y is less like to occur in the presence of X. This approach follows, but considerably extends, the analysis in Mintz and Schrodt (1988).

There are two advantages to this approach. First, it provides a level of aggregation, in the familiar paradigm of probability theory, that we did not have in our earlier work that directly dealt with the display. Of course, as with any aggregation, we lose detail here as well, so these two approaches should be viewed as complementary rather than one substituting for the other.

Second—and particularly critical for this data set—the probability calculations at least partially compensate for two major changes in the raw frequency of events that occur around 1999. First, at that point the source texts shift from Reuters to AFP, which generally generates a higher number of events. Second, the al-Aqsa *intifada* breaks out shortly thereafter, which further increases the event counts. This difference is very clear in the display, which is far denser in the final 20% of the data. As we shall see, it has far less effect on the conditional probability calculations, which deal with *relative* frequencies.

[4] In order to save typing—and reduce the frequency of *MS-Word* crashing—we have not included time subscripts on X and Y except in formal equations. These are implicit: Y and X always refers to the frequency of a pattern at a particular point in time.

The conditional probabilities of various patterns were computed from the textual output of the display tool using a short 400-line C program.[5] In general, we are using the conventional definition of a conditional probability

$$\Pr(Y|X) = \frac{\Pr(Y \cap X)}{\Pr(X)}$$

These calculations are a little ambiguous because it is unclear what constitutes an "observation" in the data set since it is possible for a pattern to occur multiple times in a single two-day period. Consequently for a time interval T we have defined

$$\Pr_T(Y \mid X) = \frac{\sum_{t \in T} y^*_t}{\sum_{t \in T} \max(1, y^*_t)})$$

where

$$y^*_t = \begin{cases} 0 \text{ if } X_t = 0 \\ Y_t \text{ otherwise} \end{cases}$$

This has the effect of counting all of the cases where both Y and X occurs—effectively inflating Pr(X) in this situation—but only one "non-event" when X occurs but Y does not. If this adjustment is not made, and instead Pr(X) is calculated without reference to Y, some of the "probabilities" will be greater than 1. The alternative would be to adjust downwards and only count a single Y for any given occurrence of X, which would also keep the probabilities bounded by 0 and 1; we may experiment with this alternative formulation in the future.

The probabilities are computed in a moving window of that is 120 days in length (60 2-day intervals in the display). In figures 2 through 8, the labels on the X-axis show the beginning of the time interval. The roughly 26-year span of the data provides for 80 probability estimates. When a probability cannot be computed—Pr(X) = 0—there is a gap in the line, though sometimes this will be obscured by the other line, or *MS-Excel* will decide to put a line in there just because, well, *MS-Excel* decides to. In the discussion that follows, Pr(Y|X) will be referred to as the "conditional probability" and Pr(Y) will be referred to as the "marginal probability."

Conditional probability assessment of rules.

Figures 2 through 6 show the difference in the conditional and marginal probabilities for the various patterns. Four general characteristics are immediately apparent from these graphs. First, except for a small set of intervals, almost all of the differences are positive, indicating that the conditional probability is higher than the marginal: the occurrence of the behavior in one dyad increases, sometimes quite substantially, the probability of observing it in the other dyads. Second, the differences between the two dyads are highly, but not perfectly, correlated.[6] This is consistent with the general observation that event data display a high level of reciprocity; this result extends that reciprocity not just to the patterns, but their interdependence. Third, the material and verbal

---

[5] The program is in ANSI C and compiled with gcc; the source code and—shock!—reasonably coherent documentation is available from the authors. A command file gives the X-Y pairs of variables on which probabilities are computed and the window used for the calculation; it also allows X to be lagged. Calculations are nearly instantaneous. In the future we expect to integrate this into the event pattern web site itself.

[6] It is possible that this is an artifact of the calculations: we've got to do a little more work on the underlying math to figure out whether the common Pr(Y ∩ X) term in the calculation of the two lines places a lower bound on their correlation, or whether this is primarily a function of the high level of reciprocity in the event data generally.

events generally display quite difference characteristics, which we would expect to see. Finally, in all of the cases, the probabilities vary over time, and in many instances, these changes correspond closely to the major political epochs of the conflict: the invasion of Lebanon, the first *intifada*, the Oslo negotiations, and the al-Aqsa *intifada*.

Conflict

Figure 2 shows the conditional probability differences for the simple count of the conflict events. The PI line involves conflict directed from the Palestinians to the Israelis; IP is conflict directed from the Israelis to the Palestinians. These are basically probabilities of events rather than patterns.[7]

Perhaps surprisingly, these differences are generally relatively low compared to some of those in the later charts, and for material conflict, the difference goes to almost zero during the al-Aqsa *intifada* period. This initially counter-intuitive result is largely a function of the high *marginal* probability of material conflict in this dyad: because the dyad generally sees a high level of conflict throughout the period, the additional information provided by conditioning on conflict is low. An interesting example of this are the two spikes in Figure 2a, which occur in the period between Israel's withdrawal to south of the Litani in Lebanon and the outbreak of the first intifada, which is generally a very quiet period in the dyad. Only here do we see an isolated high conditional difference on the Israeli side, with Israel responding strongly to the small number of Palestinian uses of material conflict.

The verbal conflict graph shows a clear concentration of high differences during the period of the Oslo process, roughly 1992 to mid-2000. However, even during this period the differences are relatively low, in the range of 0.1 to 0.3. Atypically, the verbal conflict differences show a number of negative values in the period prior to the first *intifada* (December 1987), and the pattern generally looks consistent with random fluctuations around zero. Consequently, in contrast with material conflict, verbal conflict offers no predictive value for anticipating verbal conflict by the other side.

Cooperation

The graphs for cooperative events shown in Figure 3 are quite different than those for conflict. First, there are a far greater number of differences with high magnitude, hitting the range of 0.5 to 0.7, indicating that the conditional probabilities are substantially greater than those of the marginal probabilities. This is probably due in large part to the fact that cooperative activity in this dyad is much rarer than conflictual activity: this means that the marginal probabilities Pr(Y) and Pr(X) are low, which both raises the value of Pr(Y|X) due to the presence of Pr(X) in the denominator, and the value of the difference due to a smaller number being subtracted.

The pattern in Figure 3a is a bit counter-intuitive, since the high probabilities tend to correspond with periods of conflict—Lebanon and the two *intifadas*—rather than to the Oslo period. The two curves generally coincide, though the magnitudes of the differences for the Israeli actors are usually much higher than those of the Palestinian. Compared to some of the later graphs, we do not see a strong positive relationship during the Oslo period, despite the general anecdotal sense that this was a period of high cooperation.

The explanation for this probably lies in the nature of material cooperation in this dyad: it is primarily cease-fires and prisoner deals, which coincide with conflict. While the Oslo process was supposed to involve a lot of material cooperation, the fact that the process failed meant that relatively

[7] Specifically, this analysis differs from a conflict tit-for-tat pattern—which at first reading it seems quite similar to—in the following way: The TFT is looking at a specific, deterministic *short-term* response across a fourteen-day period, with a lag time between the "tit" by one side and the "tat" by the other. The conditional probability is looking at the *aggregate likelihood* of conflict behavior over a 120-day period, with no assumption that the behavior of one side precedes the other.

little actually occurred, and what did was generally unilateral (Israeli withdrawal from Palestinian urban areas). Furthermore the material cooperation that did occur—for example for a couple of years the two sides conducted cursory joint patrols in areas where there was shared control—these were either not newsworthy, or not picked up adequately in the event coding scheme.

The plot for verbal cooperation, in contrast, shows a strong pattern of decline, with high values at the beginning of the graph and very low values—almost zero—during the al-Aqsa *intifada* period. However, there are almost no negative values: verbal cooperation in one direction of the dyad always increases the probability of cooperation in the other. In contrast to the material cooperation, there is very little difference between curves of the two sides.

The explanation for this is quite straightforward: as the data progresses, the two sides simply talk to each other more frequently, and by the period of the al-Aqsa *intifada* these interactions have become routine and consequently carry almost no information. Verbal cooperation—typically meetings and agreement—is almost by definition symmetrical, which explains the high degree of convergence between the two lines.[8] At the beginning of the sequence, these symmetric events are rare but when they occur—for example during the period of Israel's gradual withdrawal from the area between Beirut and the Litani—it is a strong indicator that the other side will also engage in this behavior. By the end of the period, these behaviors have become routinized.

Conflict TFT

Figure 4 shows the difference between the conditional and marginal probabilities for the conflict tit-for-tat *pattern*: the probability is whether one party is using a strategy of TFT response to the other, as distinct from whether the events themselves are TFT. The "PI TFT Conf" line corresponds to Palestinians responding to Israel; the "IP TFT Conf" line is the reverse.[9] Note also that in contrast to most discussions of TFT, we are distinguishing between conflictual and cooperative behavior—in classical TFT terms, we are looking separately at DD/DD/DD and CC/CC/CC… patterns. Furthermore, we are looking at the *probability* of the TFT pattern occurring in a 120-period, not for a specific instance of TFT.

As we demonstrated in earlier papers, the extent to which TFT is found in the data varies over time; that characteristic is reflected in the conditional probability behavior. The conflict TFT is, frankly, a bit puzzling. We note first that almost all of the differences are positive—much more so than we saw with the conditional differences on the events themselves—and generally very highly correlated, even in a complex sequence of probabilities such as that seen during the first *intifada* (May-87 to May-91 in the chart). The differences are also generally much higher than those of the event measures, which suggests that the two sides are adjusting their behavior more to general strategies than to specific events.

Beyond this, the patterns, while clearly corresponding to political epochs, are a bit difficult to interpret. In the first period, corresponding to Israel's involvement in Lebanon, we see a very erratic set of spikes that probably correspond to periods of military engagement between the two sides.[10]

---

[8] We have not systematically computed correlations for the lines—we will do so before submitting this for publication—but in those instances that we have, they confirm the "eyeball test" that the two lines are correlated at a high level.

[9] In other words, while the *sequence* in the "PI TFT Conf" pattern is an Israeli action first, followed by a Palestinian response, the *label* puts the responder—the actor who is employing the strategy—first. "PSE TFT Conf" might be a less confusing label…

[10] Recall that in contrast to the most recent invasion of Lebanon, Israel's 1982 invasion was intended to destroy the "state within a state" that the PLO had established in Beirut and parts of southern Lebanon following it's expulsion from Jordan in the "Black September" of 1970. This was, in the short term, largely successful. In the long term, Israel traded the militarily ineffectual, and largely isolated PLO for the militarily sophisticated, Iranian-armed and locally popular Hezbollah. Despite the military defeat and evacuation of the PLO in the summer of

This is followed by a very clear pattern of rising conditional probability that corresponds almost exactly to the first *intifada*, which effectively ends in August 1991 when most of the West Bank and Gaza are put under curfew in response to Palestinian support for Iraq's invasion of Kuwait. The rising pattern presumably reflects the increased institutionalization and co-adaptation of the two side during the *intifada*.

The probabilities decrease to near zero following the Iraq war and then pick up again during the Oslo process. The conflict TFT during this period probably reflects the fact that while Oslo involved considerable cooperative behavior, it has always been opposed by Palestinian Islamic militant groups, and actions those groups engaged in a series of TFT exchanges with Israel of assassinations and bombings throughout the second half of the 1990s. Since we have aggregated all Palestinian groups in this analysis, we are probably seeing that behavior reflected here. Finally in the period of the al-Aqsa *intifada,* conflict by both sides becomes so frequent that the conditional probabilities provide essentially no information.

The pattern in verbal conflict is quite different. As before, there is virtually no activity prior to the Oslo period, since the two sides were essentially not talking to each other. The exception is some activity in 1988 that probably corresponds to sparring in conjunction with third-party mediation of the first *intifada*. (In the period prior to 1988, the conditional probability is actually undefined because there is no TFT behavior, but Excel has helpfully drawn a line there anyway…). The TFT picks up with the establishment and institutionalization of the Palestinian Authority, and continues at a high level even through the al-Aqsa *intifada*: in the verbal domain, we continue to see variations in the patterns of conflict, in contrast to the simple continuous violence in the material domain.

Cooperative TFT

As expected, the cooperative TFT patterns shown in Figure 5 are quite different than the conflict TFT, which we had also noticed in earlier analyses. In material cooperation, the TFT is completely absent prior to the Oslo process, and then spikes for about a two year period during which cooperation in the context of Oslo was most active (prior to the assassination of Israeli Prime Minister Yitzhak Rabin in October 1995). Following that, we actually see a consistent *negative* pattern—virtually the only such case in all of our analyses—which interestingly corresponds to the period of time when there are repeated unsuccessful attempts to implement Oslo under Rabin's successors, particularly Netanyahu. These negative differences are quite small and may essentially be zero, but the dramatic contrast between this and the Rabin period is striking. The cooperation picks up again at the beginning of the al-Aqsa *intifada*—this again probably reflects the split between the militants, who enthusiastically embraced the revolt, and the Palestinian Authority, which waffled between trying to continue negotiations with Israel and joining with the Islamic groups in the military conflict. The diminishing of the relationship in the final years of the sequence is consistent with the weakening of the PA—including both the physical isolation and deteriorating health of Yassir Arafat—during those years.

The graph for verbal cooperative TFT, in contrast, looks very similar to that of the simple verbal cooperation event counts: an irregular but very steady decline from the beginning to the end of the series. The similarity of these two—which we did not find in the other pattern/event pairs—may again be a function of the fact that the most common cooperative verbal events are meetings and agreements, which are necessarily reciprocal, and therefore reciprocal event behavior almost always matches a TFT pattern.

1983, sporatic clashes still occurred with Palestinian refugees remaining in Lebanon—most of whom did not evacuate—as well as some violence in the West Bank and Gaza.

Olive Branch

In the olive branch pattern, one actor responds to prior conflict by the other with cooperation: it is essentially the DC pattern of classical prisoners' dilemma studies. In the graphs, "PI olive" is the pattern of Palestinian cooperation in response to Israeli conflict; "IP olive" is Israeli cooperation in response to Palestinian conflict. Note that the olive branch pattern requires only some cooperation; it does not require only cooperation (though this could be easily programmed using the existing system) and in fact one frequently sees the behaviors combined.

Two things are evident from the olive branch graphs. First, in contrast to most of the behaviors we've studied, the verbal and material graphs look generally similar. At the very least, this lends a certain face validity to the pattern: conciliatory behavior in this dyad is rare (and risky, as Rabin discovered) and consequently when it is undertaken, it is reflected across the board.

Second, the conditional difference for Israel is very substantially higher in magnitude than that for the Palestinians, even though the curves themselves correlate fairly well. In other words, the probability that the Israeli will offer an olive branch conditioned on Palestinian olive branch behavior is much stronger than the Palestinian probability conditioned on Israeli olive branches. This probably reflects both the greater institutionalization and coordination on the Israeli side, and also the fact that the Palestinian side is much more split on strategy than the Israeli, and consequently responds more cautiously.

Beyond this, we find most of the high differences are concentrated during the Oslo period, which we would expect. There is a short spike of olive branch coordination during the Lebanon conflict—this is probably a reflection of the third-party mediation that resulted in the PLO evacuation from Beirut—and some limited coordination during the al-Aqsa *intifada* (where, interestingly, the magnitudes of the Israeli and Palestinian curves are roughly similar, in contrast to earlier periods). In the first *intifada,* in contrast, the olive branch relationships are *negative*—the likelihood of one side using an olive branch strategy in the presence of the other side doing this is less than the likelihood of an olive branch strategy in general. This would be consistent with a situation where olive branch behavior was interpreted as a sign of weakness, an interpretation that is certainly not inconsistent with the rhetoric of the first *intifada*.

Conditional probability assessment of rules.

A second application of conditional probability is within some of the if-then rules that we have used as patterns. Because if-then rules are themselves composed of two patterns, an antecedent and a consequent, we can use the same approach we used earlier, except that now we will decompose the pattern. Because the antecedent precedes the consequent in time, we can also now look at these in terms of prediction (though not necessarily causality). Letting X be the antecedent and Y the consequent, we are interested in looking at

a. Predictive rules: $P(Y|X) >> P(Y)$

b. Null rules: $P(Y|X) = P(Y)$

c. Incorrect rules: $P(Y|X) << P(Y)$

In Figures 7 and 8, we look at two of these rules. In the conflictual TFT rule, the antecedent is conflict in the period [t-14, t-8] and the consequent is conflict in the period [t-8, t]. In the olive branch rule TFT rule, the antecedent is conflict in the period [t-14, t-8] and the consequent is cooperation in the period [t-8, t]. In each instance, we can look at Israel and Palestinian behavior—noting that in these analyses the probabilities calculations involve whether one side is following the rule, not the interactions between the rule-following behavior of the two sides—and also at the contrast between verbal and material behavior. Because we are looking at the difference between the conditional and marginal probabilities, we are further assessing only whether the rule itself is predictive, not whether it is being followed. So, for example, in periods of high mutual conflict without any cooperation, the probability of the conflict TFT rule is very high, but so is the marginal probability of conflict irrespective of TFT, so the difference between the two is near zero.

Conflict TFT

The graphs for the conditional differences within the rules themselves look quite different from those of the conditional behavior between the rules in two ways. First, although generally the conditional differences are positive, indicating that the rules have some predictive value, the correlation between the two series is substantially less than it was in the earlier analyses. Second, there is substantially less correspondence between the patterns of probability and the political epochs of the conflict. This second result is somewhat surprising but not entirely counter-intuitive: those epochs are defined in part by the *interdependence* of the rules that the sides are using, not on the occurrence of the rules themselves.

In the material TFT, we see a fairly steady positive difference in the rules, though for most of the period the value is relatively low, less than 0.2. The Palestinian and Israeli series are generally comparable in magnitude, though overall the Palestinian is a bit stronger. The predictive rule breaks down completely after the first year on the al-Aqsa *intifada*: as usual, this is doubtlessly due to the very high frequency of conflict in general, so the conditional probability adds very little.

The verbal conflict TFT has a bit more of a pattern, at least on the Palestinian side. Here most of the positive differences are during the Oslo process period, though they are erratic even here. The period of the first *intifada* shows a couple of high positive spikes, but also four negative periods, indicating that the rule is not being followed. Israel shows a very consistent pattern of near zero response, which is consistent with interpretation that Israel was simply ignoring Palestinian criticisms during the entire period.

Olive Branch

The olive branch rule—Figure 8—is somewhat more interesting. As we noted earlier, the olive branch is not a common strategy in this conflict, and this is confirmed by the conditional differences: this is the only figure where we see at least as many negative values as positive, indicating that for much of the time, this rule is not operating. This will come as no surprise to observers of this conflict.

In addition, unlike the decomposition of the conflict TFT rule, there is some correspondence between the plots and the political epochs, though the erratic structure of the series makes this less clear than the structure of some of the earlier plots. In the material olive branch plot we see that the differences are far more likely to be positive during the Oslo process and during the 1983-1985 period when Israel was gradually withdrawing from much of Lebanon. A particularly interesting sequence occurs near the end of the Oslo process, where we first see a period where Israel is employing the olive branch but the Palestinians are not, followed by a period when the Palestinians employ it but the Israelis do not, followed by the complete breakdown of the rule in the al-Aqsa *intifada*. As with the conflict TFT rule, these differences are generally comparable in magnitude for the two sides, but they are not particularly well correlated.

The pattern in the verbal olive branch is very indistinct, and in fact probably involves little more than a random fluctuation around zero. Keep in mind that this is the difference between the conditional and marginal probability, not the probability of the olive branch itself: as shown in Figure 6, there actually was some coordination in verbal olive branch strategies during the latter part of the Oslo period. However, the olive branch rules were not strongly *predictive* even during this period. Given the fact that olive branches are politically costly (in contrast to TFT, which is generally the political norm), we would not expect it to be predictive, and it isn't.

Conclusion

The results of this exercise generally continue to confirm our sense that discrete patterns in event data can be used meaningfully in political analysis. While some of the patterns we have found here are more distinct than others, and in some cases we are probably just dealing with noise, in most cases we seem to be seeing decidedly non-random patterns.

The analysis also provides further evidence on two points that we think are important for event data analysis in general. First, one sees very substantial differences between material and verbal behavior, and between cooperative and conflictual behavior, and in many cases these differences have clear relationships with the qualitative narratives of the events. Second, it is possible to get meaningful results from techniques that are frequency-invariant—that is, methods such as this where the same results would be found even if, say, the number of events doubled while the distribution of events over time and their relative frequencies were kept the same. Frequency-invariance is very important when the density of coverage (the average number of stories per time unit) changes substantially due to shifts in the source (e.g. Reuters or AFP), or the coverage provided by a single source changes because of editorial policies ("media fatigue").

We mention these two characteristics because they generally are *not* true of the most frequent mode of event data analysis, which involves aggregating all events into a single number using a scale. This completely obscures the differences between the four types of events and, because in most scales material events have the highest magnitude weights, ends up obscuring the verbal behavior at the expense of the cooperative. Furthermore in protracted conflict, conflict events far out-number cooperative events, so the scaled aggregates essentially become measures of the level of conflict. In terms of frequency invariance, some analytical techniques will accommodate this (linear regression is the obvious one), though only for homoskedastic series, which is typically not the situation with event data.

That said, we are still a bit stuck as to what to do with these things, and will be the first to admit that at the moment we are still essentially at the descriptive level, showing the validity of the method (and, by implication, the underlying data), rather than at a theoretical or inferential level. Some of those descriptions—notably the temporal variation in the olive branch behavior—may be providing insights that are not immediately apparent from a purely qualitative understanding of the conflict, but in order to show the practical utility of this approach, we probably need to relate it more firmly to a theory.

And why don't we?—after all, between the two of us we have been bludgeoning undergraduate and graduate students with the importance of theory for more than half of a century. The problem we have here—and, we should add, this is quite arguably shared by Wolfram's work with respect to the natural sciences—is that the systematic (a.k.a. "scientific") approaches to international relations have generally not presented "hypotheses" in a language of pattern recognition. Those hypothesis are instead either correlational or, over the past twenty years, phrased in terms of the values of coefficients in models being significantly different than zero.

It would be possible to state some of our results in a hypothesis form:

> H_1: Sincere implementation of the Oslo accords are represented by periods in which conditional minus marginal probabilities of olive branch are positive; insincere implementation of the Oslo accords are represented by periods in which these differences are less than or equal to zero.

"Sincere" and "insincere" are now formally based upon the patterns we find in the data. We have thus linked a concept—sincerity of commitment to peace—to an empirical indicator which we can now use in any other conflict analysis, such as Darfur or Bosnia.

However, this is not what we did—we instead observed this inductively, *post facto*. Unfortunately, the existing theoretical literature also does not present *a priori* hypotheses in pattern-based form.

Consequently, having established (at least to our own satisfaction) the existence of complex patterns in this data, perhaps the next step in our research should be to take a existing theoretic approach, translate it into a pattern-based form, and then seek those patterns in the data, rather than running some general patterns and seeing what we get.

In addition, predictive utility is not the only purposes for which we might use these rules. For example, there is also value in deciphering what rules the two sides were using, how use of those rules evolved over time, how "successful" the rules were in eliciting response from the other side, and so forth, and all of these can be done with the existing method. Furthermore, we are doing this in a systematic fashion that could never even be attempted before except on a qualitative level. That is, we can falsify qualitative statements about rule use, non-use, non-response, and so forth, and that in turn will provide the inductive/abductive basis for theory development.

That, however, raises several additional questions that go beyond the relatively simple conditional probability analysis that we have presented here. Several of these have been suggested by some of the idiosyncratic results of this research, and might be pursued in more detail in a later exercise.

First, we would be interested in identifying the time periods when the antecedents of rules are encountered with a high frequency ($Pr(X)$) as distinct from situations where $Pr(Y|X)$ or $Pr(Y|X)-Pr(Y)$ is high—that is, distinguishing between whether a rule might be invoked because the requisite antecedent conditions are present from situations where the rule was "correct"—both the antecedent and consequent were found. The frequency and conditional probabilities are two different measures and there might be some interesting relationship between. We are also interested in rules that are encountered with a low frequency: for example we would not be surprised to find that behavior involves a combination of high-frequency "standard operating procedures" that account for most behavior and low-frequency "crisis behavior" that also occurs predictably but only in exceptional circumstances.

In the process of discussing this possibility, however, we have also encountered another issue: when is a rule "interesting?" That is, there are undoubtedly a number of trivial rules that have high predictive conditional probabilities, but predict behaviors that are routine either in the sense that they occur frequently in the data set, or they are uninteresting for substantive reasons. "Interesting" rules, in contrast, probably involve a combination of novelty (the rule predicts a pattern that has not occurred frequently earlier in the data set) and substantive utility (the rule predicts events with a clear theoretically-relevant interpretation, for example the escalation or de–escalation of the conflict, rather than something pattern that is rare but has no obvious meaning.). The conceptualization of "non-trivial" rules is an important next step.

A second major possibility of future research—albeit one that is more of methodological than theoretical interest—is the evaluation of actual rules against random data and random rules against actual data: what is the probability that we are simply finding these patterns by chance? These assessments are comparable to the probabilities of Type I and Type II error in statistical analysis. Specifically, we would want to assess

a. What is the probability that rules we have specified based on the qualitative and theoretical literature will be found in a sequence of events generated randomly? This assessment can be done on various sets random data sharing increasing levels of structure with the true data, for example by using Monte Carlo methods to generate data sets with the same marginal distribution with respect to the number of events by dyad but with a uniform distribution across event types; then adding the additional restriction that the marginal distribution of event types correspond to the actual data; then adding the additional restriction that the marginal distribution of complementary event pairs correspond to the actual data and so forth.

b. What is the probability that randomly generated rules will be found in the actual data? In other words, to what extent are there "rules" in any data set? This is a somewhat more difficult problem since it requires developing the concept of a "random rule" but by modifying

Wolfram’s methods for specifying a space of discrete rules, it should be possible to do this systematically.

The work we have done so far is still tentative and has only begun to explore the possibilities of a pattern-based approach to event-data analysis. Nonetheless, we find even these first steps to be very promising: the data show patterns that are credible from the perspective of our qualitative and theoretical understandings of the conflict, but also enable us to characterize the event stream in a more systematic fashion than we could with other tools. Automated coding has allowed us to generate far more detail on sub-state actors than was found in earlier, human-coded data, and this in turn should give us greater insights into the nuances of this and other conflicts. We are on the threshold of developing a new social science tool that may offer rigor in the assessment of non-continuous, agent-produced variables, the success of which endeavor may have far-reaching methodological ramifications for many fields.

Bibliography

Abelson, Robert P. 1973. "The Structure of Belief Systems." In *Computer Models of Thought and Language*, edited by R.C. Schank and K.M. Colby. San Francisco: Freeman.

Anderson, John R. 1983. *The Architecture of Cognition*. Cambridge: Harvard University Press.

Axelrod, Robert. 1984. *The Evolution of Cooperation*. New York: Basic Books

Bickerton, Ian J. and Carla L. Klausner. 1998. *A Concise History of the Arab-Israeli Conflict*. (3rd edition) Upper Saddle River, NJ: Prentice-Hall.

Bond, Doug, J. Craig Jenkins, Charles L. Taylor, and Kurt Schock. 1997. “Integrated Data for Events Analysis (IDEA): An Events Typology for Automated Events Data Development.” *Journal of Peace Research* 40: 727-745.

Bond, Doug, Joe Bond, J. Craig Jenkins, and Charles L. Taylor. 2003. “Mapping Mass Political Conflict and Civil Society: The Automated Development of Event Data.” *Journal of Conflict Resolution* 41: 553-579.

Dixon, William J. 1986. "Reciprocity in United States-Soviet Relations: Multiple Symmetry or Issue Linkage." *American Journal of Political Science* 30:421-45.

Gause, F. Gregory. 1999. "Systemic Approaches to Middle East International Relations." *International Studies Review*1,1:11-32.

Gerner, Deborah J. 1994. *One Land, Two Peoples: The Conflict over Palestine*. (2nd ed.) Boulder, CO: Westview Press.

Gerner, Deborah J., Philip A. Schrodt, Ronald A. Francisco, and Judith L. Weddle. 1994. “The Machine Coding of Events from Regional and International Sources.” *International Studies Quarterly* 38: 91-119.

Gerner, Deborah J., Philip A. Schrodt, Ömür Yilmaz and Rajaa Abu-Jabr. 2002. “Conflict and Mediation Event Observations (CAMEO): A New Event Data Framework for the Analysis of Foreign Policy Interactions.” Paper presented at the American Political Science Association, Boston, August 2002.

Gerner, Deborah J., Philip A. Schrodt, Ömür Yilmaz and Dennis Hermrick. 2005. “The CAMEO (Conflict and Mediation Event Observations) Actor Coding Framework.” Paper presented at the American Political Science Association, Washington, September 2005.

Goldstein, Joshua S., 1992. “A Conflict-Cooperation Scale for WEIS Events Data.” *Journal of Conflict Resolution* 36(2): 369-385.

Goldstein, Joshua S., and John R. Freeman. 1990. *Three-Way Street: Strategic Reciprocity in World Politics*. Chicago: University of Chicago Press.

Goldstein, Joshua S., and Jon C. Pevehouse. 1997. "Reciprocity, Bullying and International Cooperation: A Time-Series Analysis of the Bosnia Conflict." *American Political Science Review* 91,3: 515-530.

Goldstein, Joshua S., Jon C. Pevehouse, Deborah J. Gerner, and Shibley Telhami. 2001. “Dynamics of Middle East Conflict and US Influence.” *Journal of Conflict Resolution* 45, 5: 594-620.

Holland, John H., Keith J. Holyoak, Richard E. Nisbett and Paul R. Thagard. 1986. *Induction: Processes of Inference, Learning and Discovery*. Cambridge: MIT Press.

Hudson, Valerie M, Philip A. Schrodt and Ray D. Whitmer. 2004. "A New Kind of Social Science?: Applying Wolfram-Like Models to the Analysis of International Behavior." Paper presented at the International Studies Association, Montreal, March 2004.

Hudson, Valerie M, Philip A. Schrodt and Ray D. Whitmer. 2005. "A New Kind of Social Science?: Moving Ahead with Reverse Wolfram Models Applied to Event Data." Paper presented at the International Studies Association, Honolulu, March 2005.

Huxtable, Phillip A., and Jon C. Pevehouse. 1996. "Potential Validity Problems in Events Data Collection." *International Studies Notes* 21, 2: 8-19.

Khong, Yuen Foong. 1992. *Analogies at War*. Princeton: Princeton University Press.

King, Gary and William Lowe. 2004. "An Automated Information Extraction Tool for International Conflict Data with Performance as Good as Human Coders." *International Organization* 57,3: 617-642.

Kohonen, T. 1984. *Self-Organization and Associative Memory*. New York: Springer-Verlag.Newell,

Margolis, Howard. 1987. *Patterns, Thinking and Cognition: A Theory of Judgement*. Chicago: University of Chicago Press.

McClelland, Charles A. 1976. *World Event/Interaction Survey Codebook*. (ICPSR 5211). Ann Arbor: Inter-University Consortium for Political and Social Research.

Mintz, Alex and Philip A. Schrodt. 1988. "A Conditional Probability Analysis of Regional Interactions in the Middle East." *American Journal of Political Science*. 32,1: 217-230.

Newell, Allan and Herbert Simon. 1972. *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall.

Political Psychology. 2003. "Special Issue: Neuroscientific Contributions to Political Psychology." *Political Psychology* 24,4.

Purkitt, Helen E. 1991. "Artificial Intelligence and Intuitive Foreign Policy Makers Viewed as Limited Information Processors." In *Artificial Intelligence and International Politics*, edited by Valerie Hudson. Boulder: Westview.

Reber, Arthur S. 1993. *Implicit Learning and Tacit Knowledge*. New York: Oxford University Press.

Schrodt, Philip A., and Deborah J. Gerner. 1994. "Validity assessment of a machine-coded event data set for the Middle East, 1982-1992." *American Journal of Political Science* 38: 825-854.

Sylvan, David, Stephen Majeski and Jennifer L. Milliken. 1991. "Theoretical Categories and Data Construction in Computational Models of Foreign Policy." In *Artificial Intelligence and International Politics*, edited by Valerie Hudson. Boulder: Westview.

Simon, Herbert A. 1982. *The Sciences of the Artificial*. Cambridge: MIT Press

Tessler, Mark. 1994. *A History of the Israeli-Palestinian Conflict*. Bloomington, IN: Indiana University Press.

Ward, Michael and Sheen Rajmaira. 1992. "Reciprocity and Norms in U.S.-Soviet Foreign Policy." *Journal of Conflict Resolution* 36,2: 342-368.

Wolfram, Stephen. 2002. *A New Kind of Science*. Champaign, IL: Wolfram Media

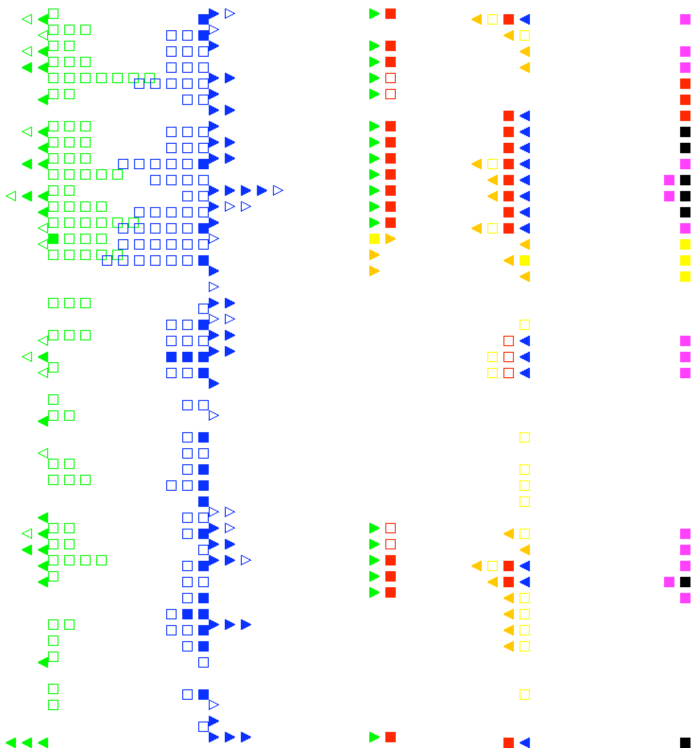

Figure 1a: Pattern Display for 28 September 1995 to 28 December 1995

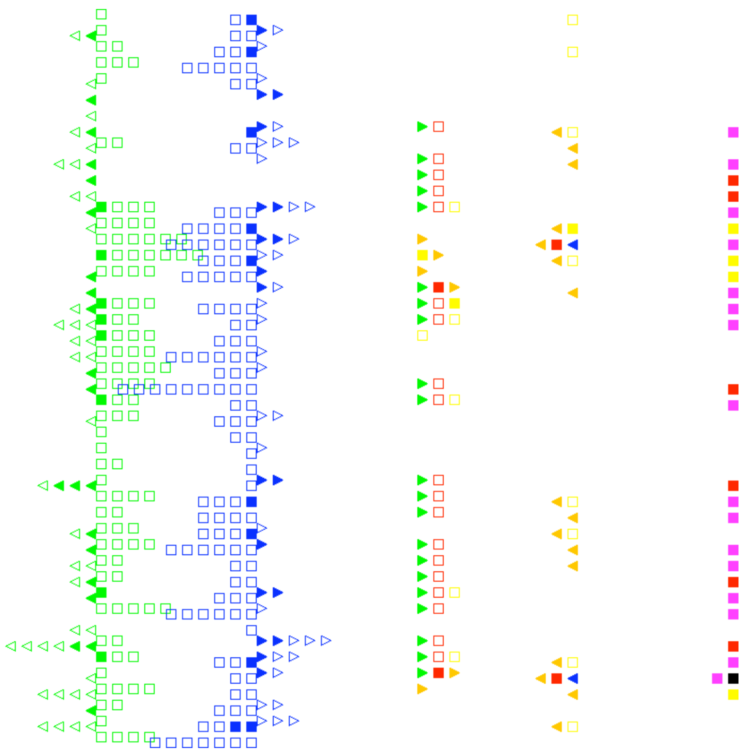

Figure 1b: Pattern Display for 11 July 2000 to 10 January 2001

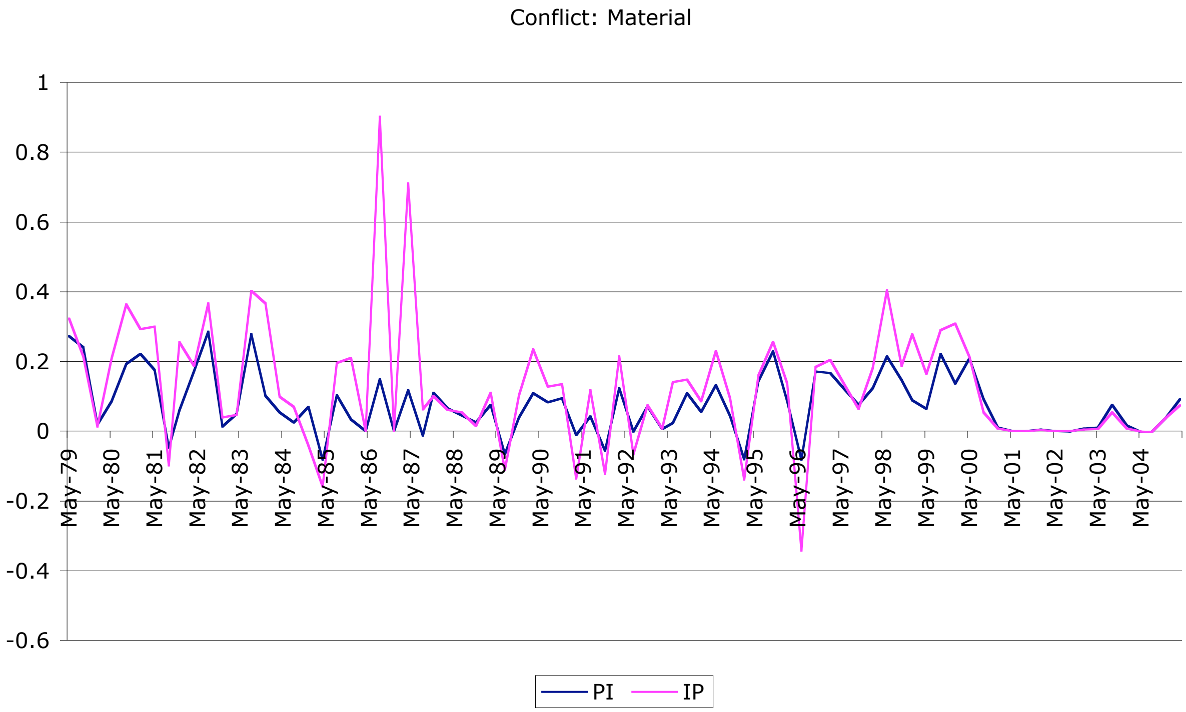

Figure 2a: Difference between conditional and marginal probabilities for material conflict

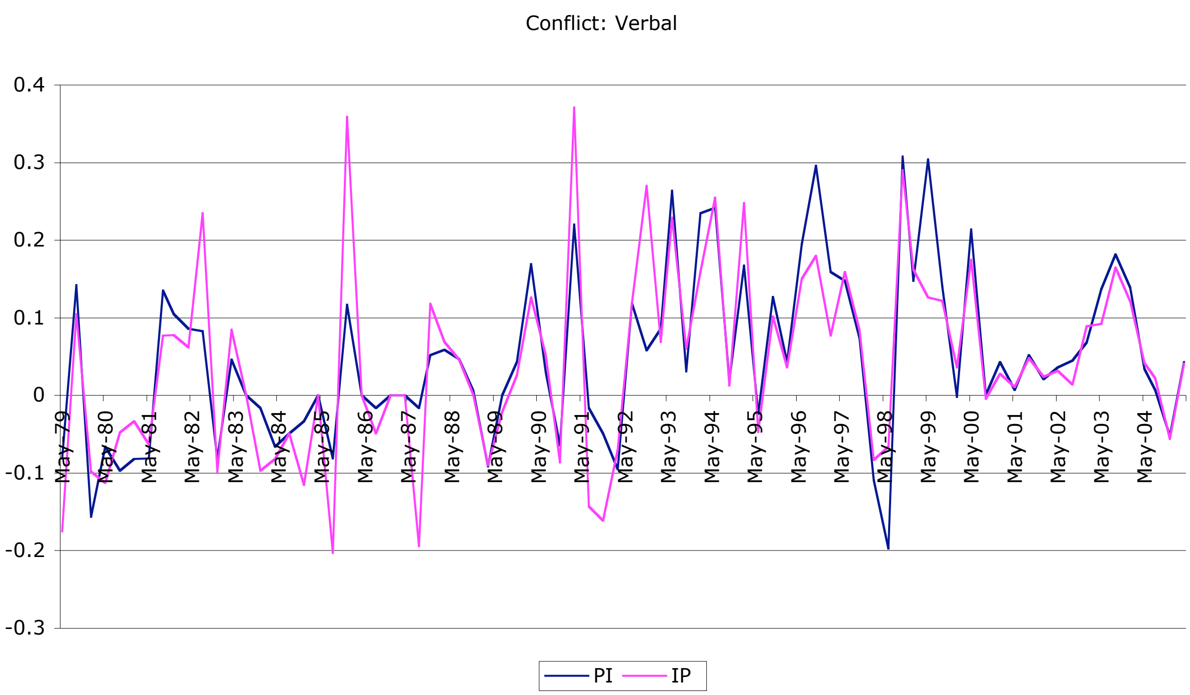

Figure 2b: Difference between conditional and marginal probabilities for verbal conflict

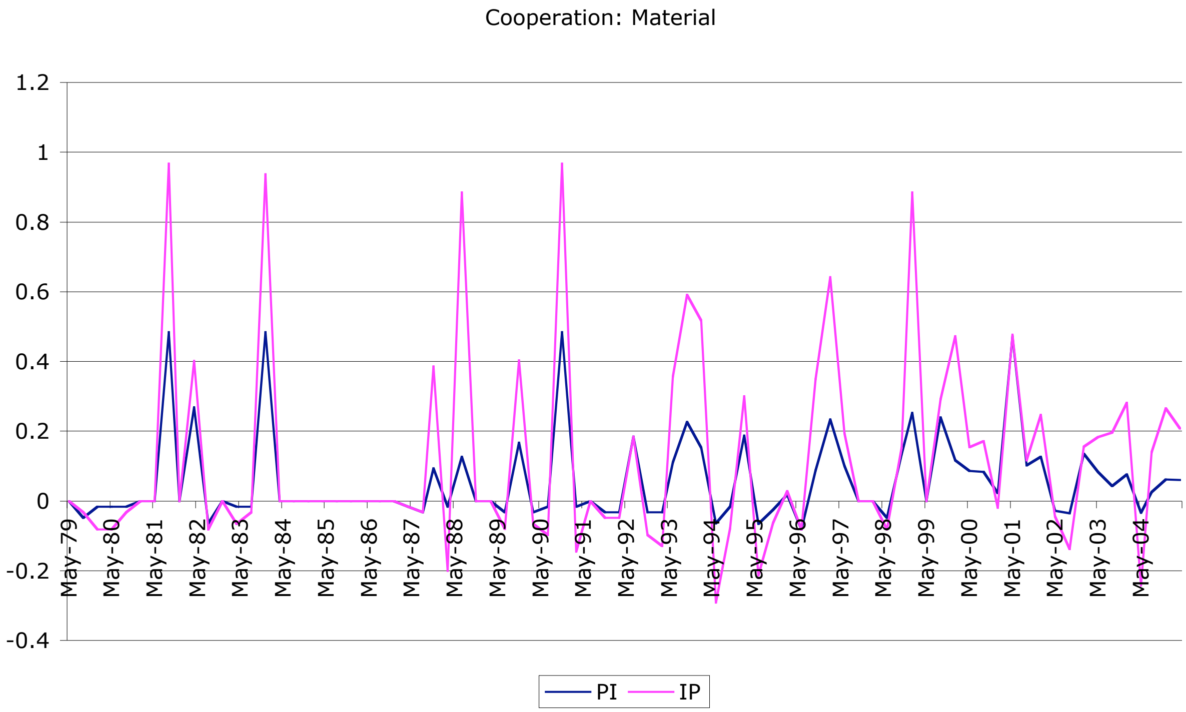

Figure 3a: Difference between conditional and marginal probabilities for material cooperation

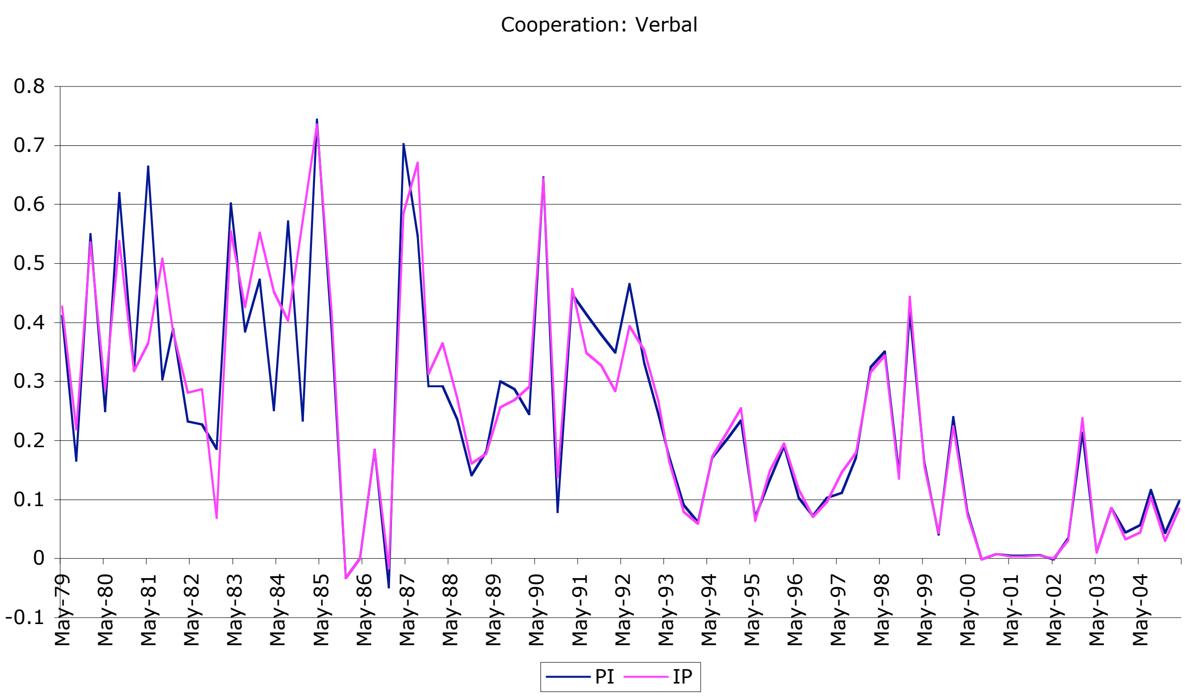

Figure 3b: Difference between conditional and marginal probabilities for verbal cooperation

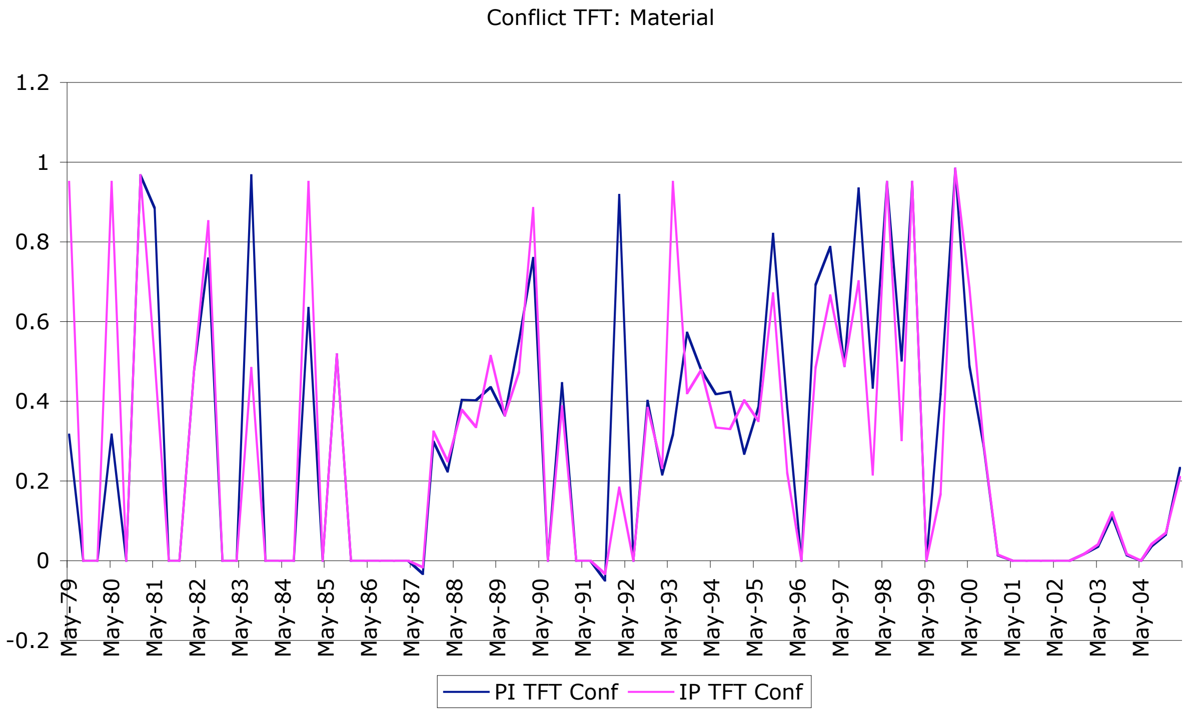

Figure 4a: Difference between conditional and marginal probabilities for material conflict TFT

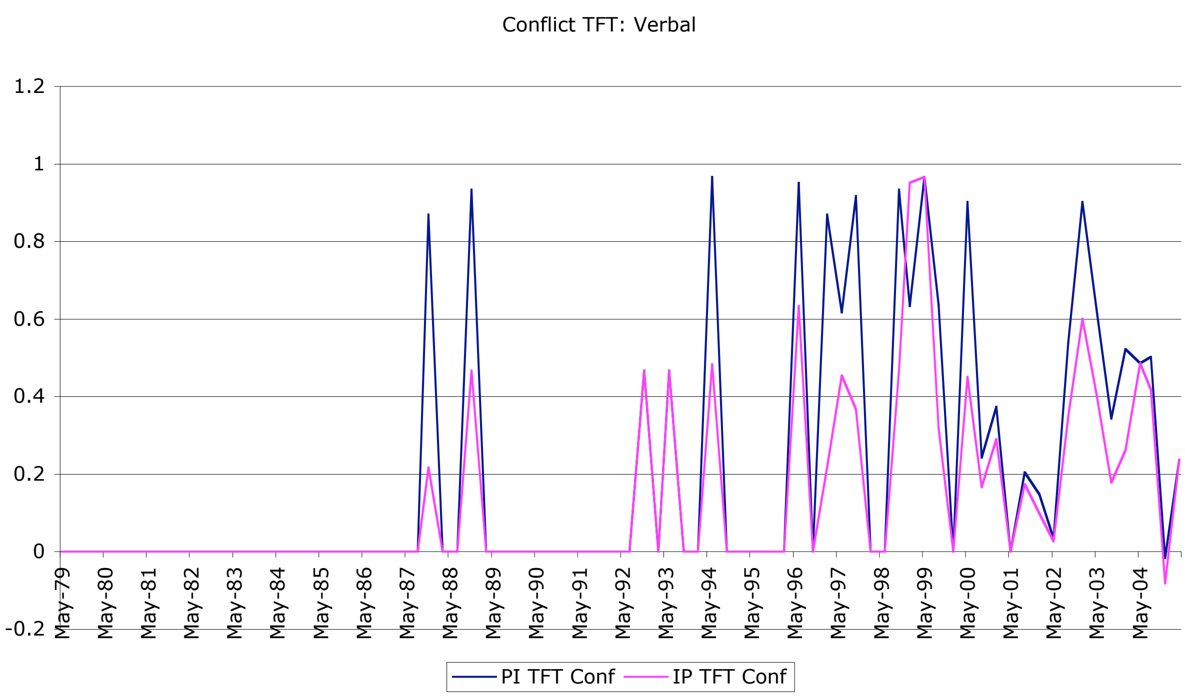

Figure 4b: Difference between conditional and marginal probabilities for verbal conflict TFT

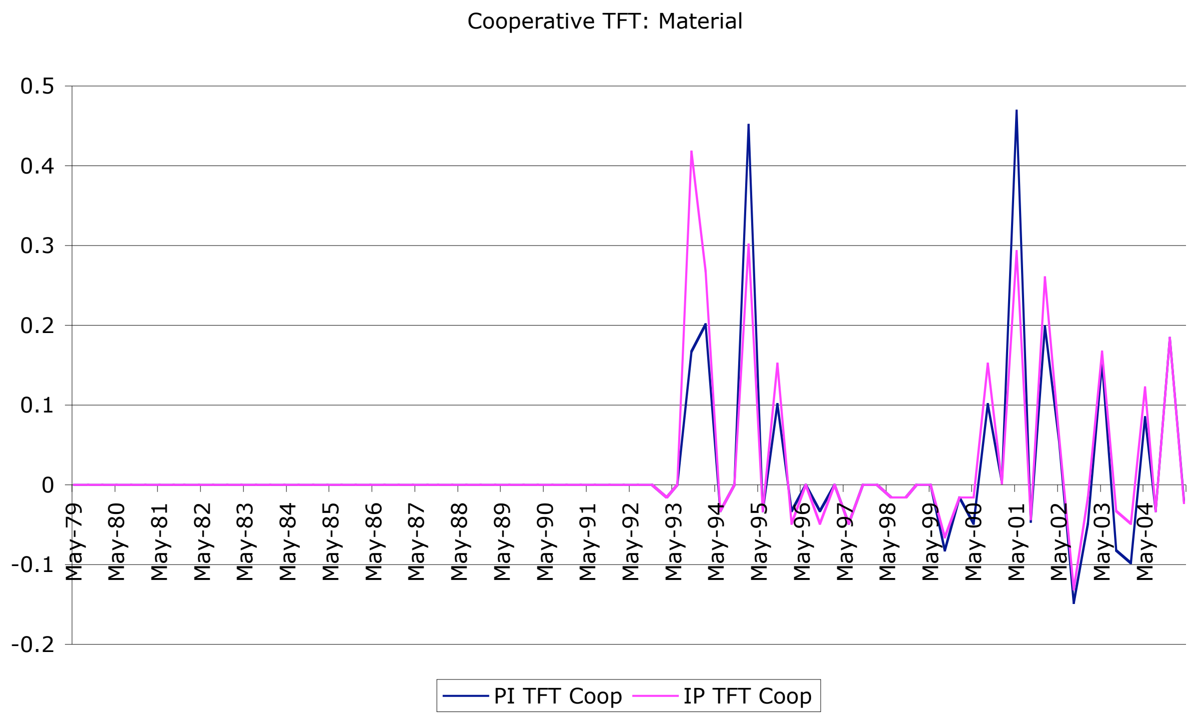

Figure 5a: Difference between conditional and marginal probabilities for material cooperative TFT

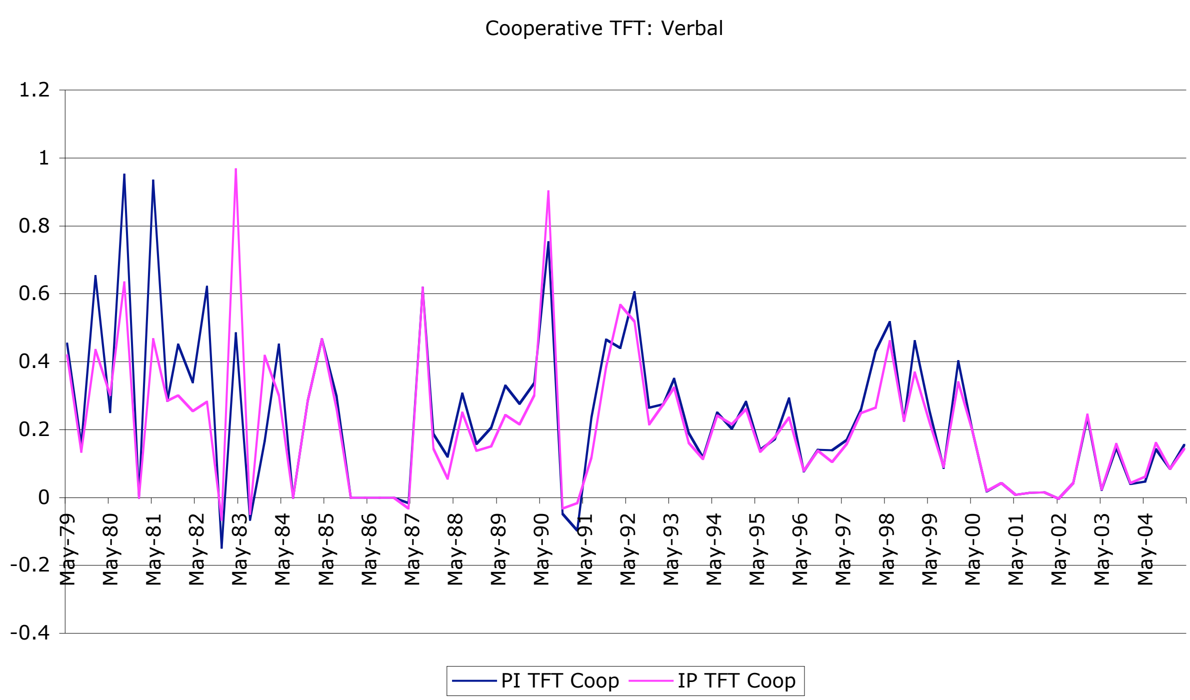

Figure 5b: Difference between conditional and marginal probabilities for verbal cooperative TFT

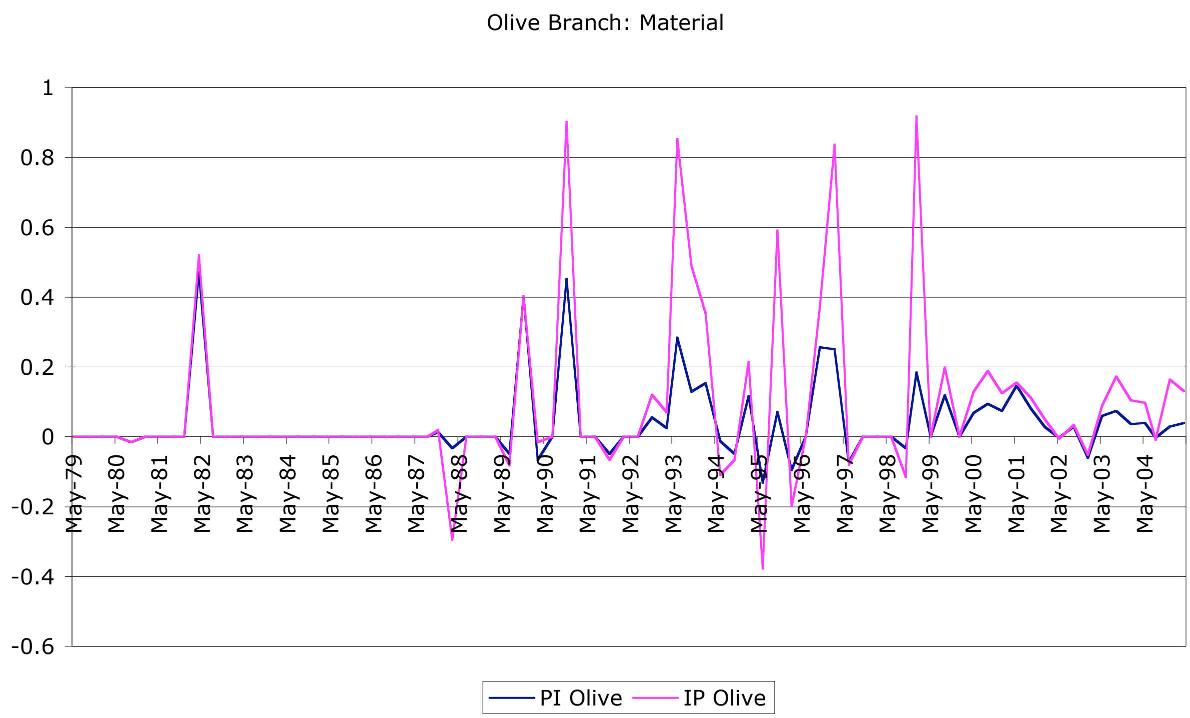

Figure 6a: Difference between conditional and marginal probabilities for material olive branch

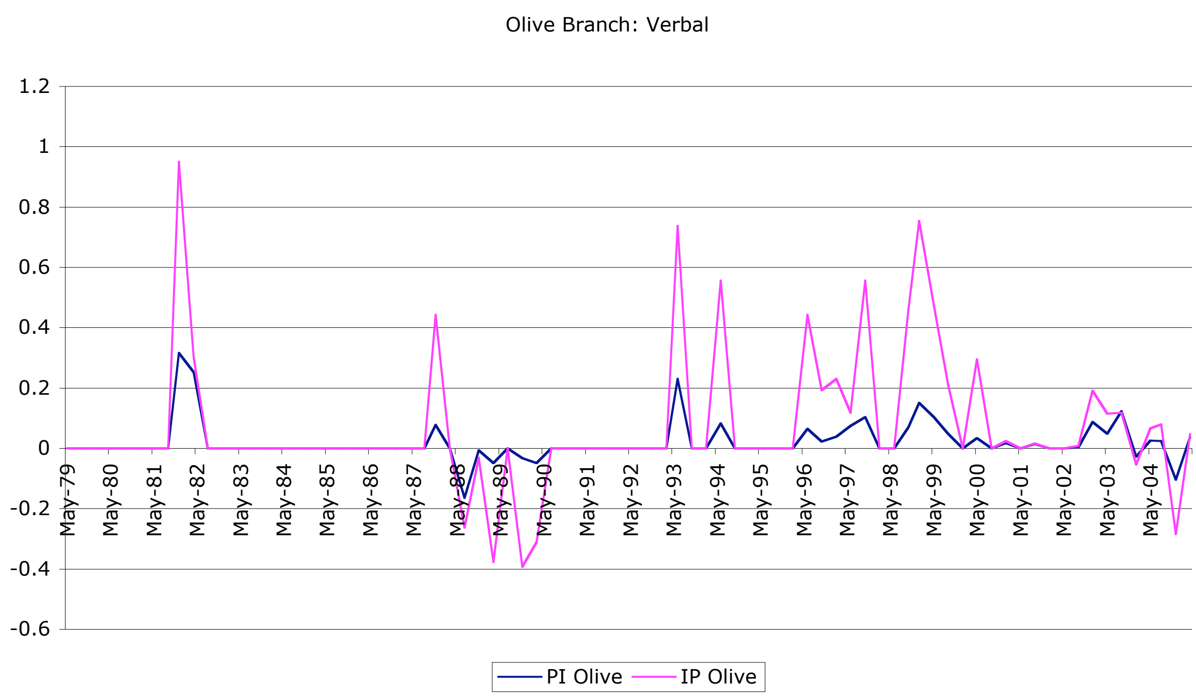

Figure 6b: Difference between conditional and marginal probabilities for verbal olive branch

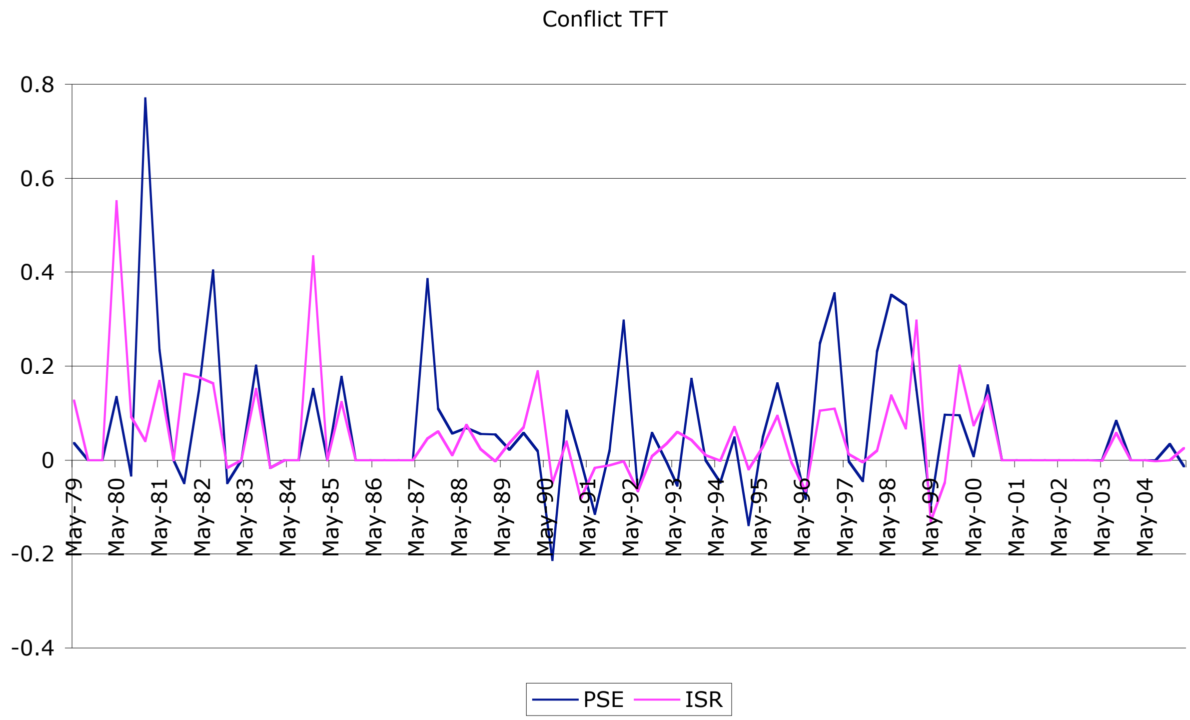

Figure 7a: Difference between conditional and marginal probabilities for antecedent and consequent in material conflict TFT rule

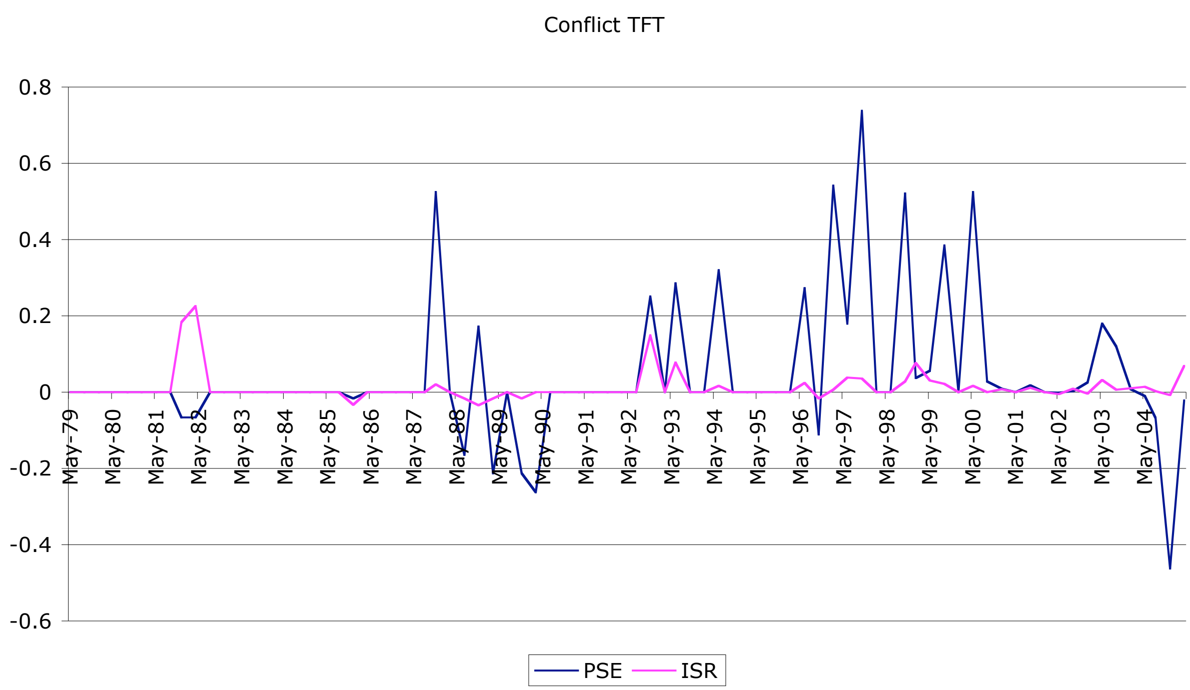

Figure 7b: Difference between conditional and marginal probabilities for antecedent and consequent in verbal conflict TFT rule

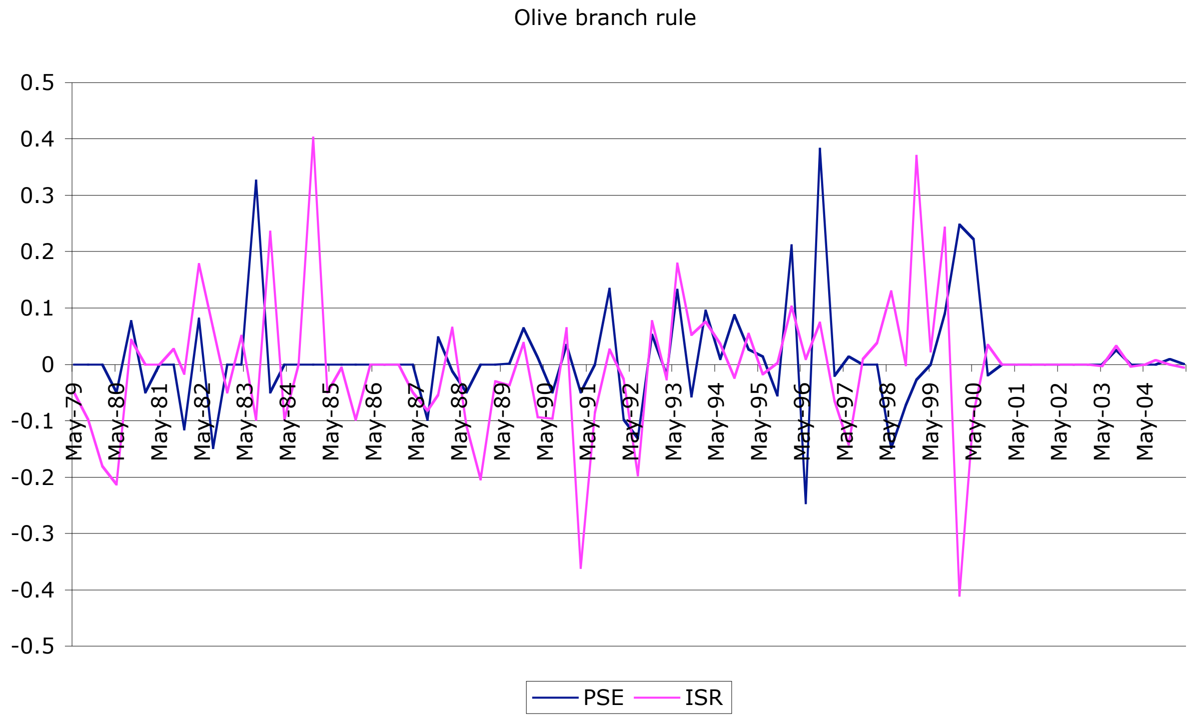

Figure 8a: Difference between conditional and marginal probabilities for antecedent and consequent in material olive branch rule

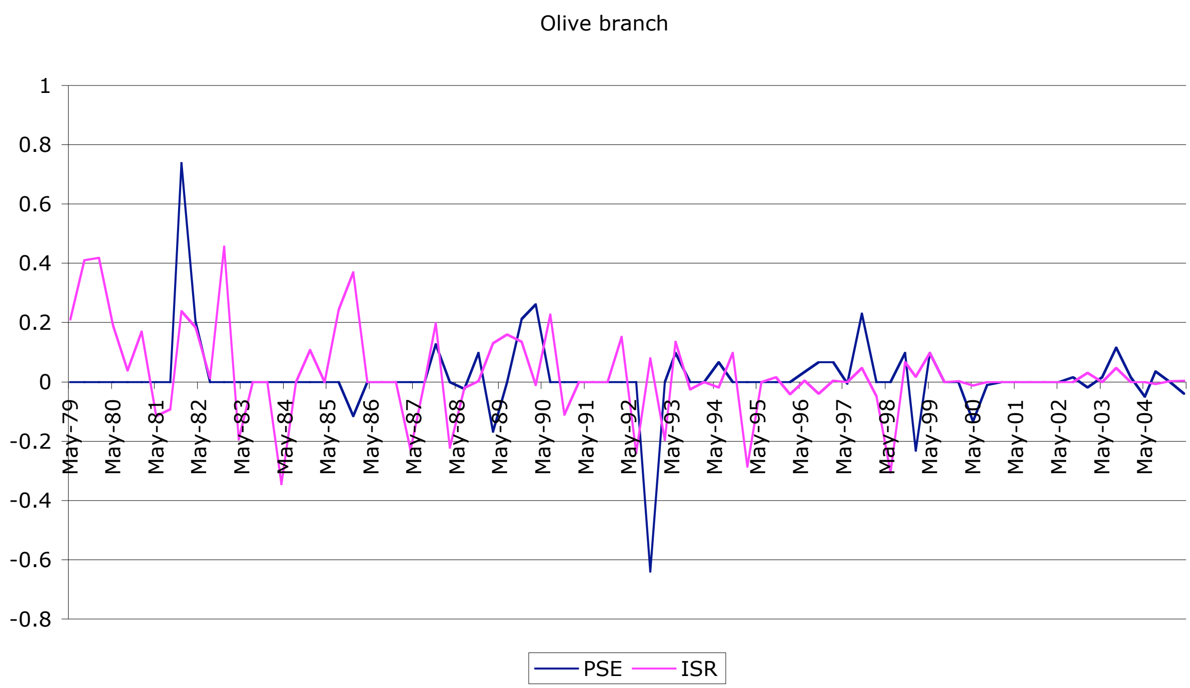

Figure 8b: Difference between conditional and marginal probabilities for antecedent and consequent in verbal olive branch rule