

# Public Opinion, the War in Iraq, and Presidential Accountability

ERIK VOETEN

*Department of Political Science and International Affairs*

*The George Washington University*

PAUL R. BREWER

*Department of Journalism and Mass Communication*

*University of Wisconsin–Milwaukee*

---

---

How do citizens hold their leader accountable during an ongoing war? The authors distinguish between two models of accountability—the “decision maker” and “managerial” models—and investigate their implications in the context of the current war in Iraq. They employ a novel measurement model and a database of survey marginals to estimate weekly time series of aggregate beliefs about various aspects of the war. Consistent with the “decision maker” model, they find that shifts in aggregate support for the war have a greater impact on presidential approval than do equivalent shifts in perceptions of war success or approval of the president’s handling of the war. Conversely, aggregate perceptions of success are more responsive to casualties and key events than are aggregate beliefs about the war’s merits. This suggests that the link from casualties and events to presidential approval is less direct than previously assumed.

**Keywords:** *Iraq; public opinion; casualties; accountability; president*

---

---

Leaders in democracies differ from their authoritarian counterparts in the conduct of war in that they require higher levels of consent from the populace to initiate and prolong wars. This need helps to explain a broad array of empirical regularities regarding the differential war behavior of democracies. For example, democracies tend to pick fights that they are likely to win quickly even though they are more likely than authoritarian states to accept draws or even defeats as a war goes on (Reiter and Stam 2002). Such patterns suggest a broader and more continuous

---

AUTHOR’S NOTE: Earlier versions were presented at the 2004 PolMeth Conference at Stanford University, the annual meeting of the American Political Science Association in 2004, and seminars at the University of Maryland and Dartmouth College. We thank the participants in those seminars, Adam Berinsky, Richard Eichenberg, and Lee Sigelman for suggestions. We also thank Harlan Geer and Ryan Schoen for valuable research assistance. The data and further information on methodology are available from <http://jcr.sagepub.com/cgi/content/full/50/6/809/DC1> and <http://home.gwu.edu/~voeten/>.

JOURNAL OF CONFLICT RESOLUTION, Vol. 50 No. 5, October 2006 809-830

DOI: 10.1177/0022002706291054

© 2006 Sage Publications

conception of democratic accountability than one based solely on voters' use of elections to reward incumbents for triumphant wars and punish them for unsuccessful military adventures (e.g., Bueno de Mesquita et al. 2003; Gaubatz 1999). In this more expansive view, democratic leaders rely on "contemporary consent": they need high levels of public support to initiate a war and must maintain that support to carry on a war (Reiter and Stam 2002). This assumes a relation between public opinion and policy not unlike that of the "dynamic representation" model for domestic policy (Stimson, Mackuen, and Erikson 1995).

How, then, does the public judge a leader during wartime when the outcome of that war is yet unclear? The connection between news from the front and the performance of the incumbent is often ambiguous, providing elite discourse with considerable room to shape the formation of political judgments (e.g., Boetcher and Cobb 2006 [this issue]; Brody 1991; Kull, Ramsay, and Lewis 2004; Larson 1996; Larson and Savych 2005; Page and Shapiro 1992; Zaller 1992). Thus, the effects of war events and casualty reports on political judgments may flow through citizens' beliefs about the war at hand. This logic suggests that we should look beyond the direct links between events, casualties, and evaluations of leaders that scholars have traditionally examined (see Mueller 1970, 1973 for the seminal analyses; see Berinsky 2004 for an overview). It also raises an additional question: what kinds of beliefs about a war affect evaluations of a commander in chief?

We build on theories of accountability for economic policies to develop two models of public accountability during wartime: a "managerial" model and a "decision maker" model. We then evaluate some of the most important empirical implications of these models in the context of the war in Iraq, a case that has become a new focal point within the literature on public opinion and war. Gelpi, Feaver, and Reifler's (2006; see also Gelpi, Reifler, and Feaver 2005) research—which has drawn the attention of the Bush administration and the news media<sup>1</sup>—suggests, among other things, that casualties shape aggregate-level beliefs about whether the war was "worth it," presidential approval on Iraq, and overall presidential approval (see also Eichenberg and Stoll 2004; Eichenberg 2005; Eichenberg and Stoll 2006 [this issue]). The authors also argue, however, that the relationship between casualties and each of these aspects of opinion differs across three stages of the war: major combat (a positive relationship), occupation (a negative relationship), and sovereign Iraq (no significant relationship). Gelpi and his colleagues furthermore argue that beliefs about the prospects of winning the war are central to understanding casualty sensitivity (Gelpi, Feaver, and Reifler 2006) and candidate choice (Gelpi, Reifler, and Feaver 2005; Gelpi, Feaver, and Reifler 2006) at the individual level.

We revisit the Iraq case by analyzing a newly created database on aggregate public opinion toward the war. Unlike previous studies, we focus on how war support and perceptions of war success differ in their aggregate-level causes and consequences. Our findings indicate that events and casualty reports affect perceptions of the war's success but have relatively little direct impact on support for the

1. See, for example, Shane (2005).

war. Conversely, shifts in aggregate war support have stronger effects on overall presidential approval ratings than do shifts in beliefs about the war's success or even beliefs about whether President George W. Bush "is doing a good job" in handling Iraq. We take this as evidence that citizens primarily held President Bush accountable for his perceived decision-making qualities rather than his managerial qualities.

## ACCOUNTABILITY AND WAR

Most theoretical analyses of how the public holds incumbents accountable for their deeds in office have focused on economic performance (e.g., Barro 1973; Ferejohn 1986; Keech 1995). The main assumption in these models is that, all else being equal, citizens prefer incumbents who are competent economic policy makers to those who are not. Competency is, however, not directly observable. Instead, citizens use their appraisals of discernable economic outcomes to indirectly hold incumbents accountable for their aptitude in managing the economy.

Precisely what perceptions about the economy are politically consequential is a matter of fierce debate in the literature (e.g., MacKuen, Erikson, and Stimson 1992; Keech 1995). These debates turn around a variety of questions. Do citizens form political judgments based on the state of their personal financial circumstances or the macroeconomy? Do they use prospective or retrospective evaluations? Are they able to tell favorable economic outcomes created through genuine craftsmanship from those produced through temporary manipulation of economic policy instruments?

Similar issues arise in holding leaders accountable for their competency in security issues. Leaders, and certainly American presidents, have considerable discretionary powers to wage and manage wars. All else being equal, citizens should prefer a leader with proven abilities to make good decisions in times of crisis and to manage wars well. The competency of a leader is, however, not directly observable and can only be inferred from outcomes. In the long run, mounting casualties and battlefield defeats ought to undermine beliefs among the public that an incumbent is a competent war leader. In the short run, however, the implications of individual outcomes for a leader's competence may be ambiguous and open to competing interpretations.

Hence, the accountability process is likely to be indirect: citizens should form impressions about the war and then use those in their evaluations of the incumbent. This notion calls into question a traditional mode of analysis that examines the direct impact of events and casualties on presidential evaluations (e.g., Mueller 1973; Gelpi, Feaver, and Reifler 2006). Instead, we expect that after controlling for subjective appraisals of the war, casualties and important events will have no direct impact on aggregate support for the president. A similar argument has been advanced and validated empirically in the literature on economic voting: the impact of changes in objective economic outcomes on changes in presidential evaluations disappears after controlling for changes in subjective evaluations of the state of the economy (MacKuen, Erikson, and Stimson 1992).

But what kinds of appraisals of a war are politically relevant? One logical place to begin is with perceptions of the war's success. Gelpi, Feaver, and Reifler (2006, 7; see also Gelpi, Reifler, and Feaver 2005) and Eichenberg (2005, 171) highlight the importance of success in shaping public opinion regarding wars (as the former puts it, "success matters"; as the latter puts it, "victory has many friends").<sup>2</sup> Continuing our analogy with economic voting, we may suspect that perceptions of how well the war is going shape presidential approval. Drawing on such assessments, citizens may conclude that the incumbent is (or is not) sufficiently capable of managing wars.<sup>3</sup> A potential downfall of this model is that the link between immediate war success and presidential competence is ambiguous. Although the president, as the commander in chief, is ultimately responsible for managing a war successfully, the actual success of the war effort depends on other factors as well, such as the efforts by other administration officials, the military, foreign coalition partners, and fortune.

The president alone, however, is undeniably responsible for making the decision to launch a war. The obvious form of evidence for whether a leader is competent in this regard is his or her decision-making record. Accordingly, citizens may adjust their evaluations of the incumbent if they change their beliefs about whether a past decision to go to war was right and has been worth it (we refer to such evaluations of the war's merits collectively as *war support*). Building on this notion, the "decision maker" model holds that citizens, above all else, seek a leader whom they perceive as a wise maker of decisions about whether the use of force enhances the security of the country.

Such a model fits an important empirical regularity—namely, that democracies are disproportionately successful in wars not so much because they have greater resources or pick better strategies but mostly because they generally initiate wars against targets that they can defeat relatively quickly (Reiter and Stam 2002). This may be because public opinion constrains leaders to pick their wars wisely. Nevertheless, a leader has considerable private information when initiating a war. As a war goes on, elites may debate whether new information about progress in the war validates the wisdom and sincerity of this initial decision. As the balance of elite discourse shifts, so does aggregate support for the war (Larson 1996; Larson and Savych 2005; Zaller 1992) and, if the "decision maker model" holds, aggregate support for the leader (see also Brody 1991, 1994).

2. Again, Gelpi, Reifler, and Feaver (2005) and Gelpi, Feaver, and Reifler (2006) argue that expectations of success play key roles in shaping casualty sensitivity and candidate choice in the context of the Iraq War. Examining a broader range of cases, Eichenberg (2005, 174) argues that "both the principal policy objective and the success or failure of a military operation are crucial factors determining the level of citizen support in its aftermath."

3. A variant to this "managerial" account is that what matters are prospective evaluations of success, that is, the perceived likelihood that a military operation will succeed (Feaver and Gelpi 2004; Gelpi, Feaver, and Reifler 2006). It is, however, unclear why prospective assessments would differ from contemporaneous evaluations of war success or whether they do so. The literature on prospective voting in economics is motivated by the notion that voters form rational expectations about future economic performance based on the cyclical nature of the economy (e.g., Keech 1995). There is no equivalent theory about future war success.

Elite discourse about any given war is, of course, to some degree constrained by the actual things that happen. Yet some types of evaluations of a war are related in a less ambiguous way to events than are others and hence may be less easily manipulated by elites. Reports of mounting casualties or battlefield losses and wins should have relatively straightforward implications for perceptions of how well a war is going, but their implications for war support and presidential evaluations are less clear. For example, Democratic challenger John F. Kerry was quick to call into question President Bush's abilities as commander in chief based on the disappearance of 380 tons of high explosives in Iraq, whereas the White House immediately sought to absolve Bush from any responsibility by claiming that the weaponry may have disappeared before the arrival of American troops.<sup>4</sup> No one, however, could have credibly claimed that the loss of such large quantities of weaponry to resistance forces was evidence of success in the war, even if elite messages could (and did) differ on the extent to which it showed incompetence.

Further obscuring the relationship between news from the front and conclusions about decision-making skills are the incentives that leaders face to start wars for opportunistic reasons. A leader may want to divert attention from a poorly performing economy or some other domestic crisis (e.g., Stoll 1984; Russett 1990; Downs and Rocke 1994). A leader may also engage in strategic brinkmanship by inducing a crisis to signal competence in security affairs.<sup>5</sup> This signal may be credible if the risky military adventure is such that only a competent executive would undertake it, thereby constraining a candidate whose competence is questionable. For example, President Bush presented himself in the 2004 election campaign as a leader with the courage and strength to go it alone.<sup>6</sup> The war in Iraq lent credibility to his message. The point here is not to suggest that the war was fought for opportunistic reasons but rather that citizens may be forced to assess the sincerity of a decision to go to war. This provides elites with incentives to question—and, in turn, defend—the motives of the incumbent.

To summarize, we expect that, compared to perceptions of its success, support for the war in Iraq is shaped to a greater extent by relatively unconstrained elite discourse and hence to a lesser extent by events and casualty reports. As a consequence, aggregate opinion about the war's success is likely to be more volatile than opinion about its merits. On the other hand, we anticipate that swings in the former are relatively inconsequential politically in that we expect shifts in aggregate support for the president to follow shifts in aggregate war support more closely than they follow shifts in aggregate perceptions of war success. Moreover, if citizens are truly first and foremost interested in the decision-making skills

4. See Halbfinger (2004).

5. This rationale is derived from models of signaling competence through fiscal policy. It is based on the criticism that traditional political business cycle accounts take insufficient account of the rational expectations voters have regarding preelection tax cuts, increased spending, or (by extension) wars (see also Hess and Orphanides 1995).

6. See, for example, Bennett (2003).

rather than the management skills of a leader, then we also expect that shifts in war support predict shifts in overall presidential job approval more strongly than do shifts in a more direct measure of perceived managerial abilities: perceptions of Bush's handling of Iraq. This is a stringent test given the virtual identical phrasing of items that measure overall Bush job approval and his job performance on Iraq.

## DATA AND METHOD

In testing these claims, we face a choice between individual- and aggregate-level analyses. Numerous studies have applied the former approach to the study of economic voting (e.g., Kinder and Kiewiet 1979). Others, however, have warned that in the absence of panel data, such analyses carry the potential for endogenously induced bias (e.g., Kramer 1983; Erikson 2004). The same issue extends to the domain at hand. At the individual level, war support and perceptions of war success are difficult to disentangle (Berinsky 2004). Projection effects tend to frustrate efforts to establish why citizens hold certain opinions. For example, an individual who likes the president for other (partisan) reasons may be more likely both to believe that the war is going well and to support it, even without much information about the war (on the latter point, see Kull, Ramsay, and Lewis 2004).

We circumvent this issue by analyzing changes in aggregate beliefs and controlling for the most obvious alternative explanation that may lead to shifts in assessments of the president: perceptions of the economy. To be sure, our approach opens us to some potential problems characteristic to analyses of aggregate public opinion data and prohibits us from drawing inferences about individual variation in opinion formation (see Berinsky 1999, 2002). A sizable body of recent research has argued, however, that analyses of aggregate public opinion data can shed light on how various aspects of public opinion are related to events, to one another, and to public policy (e.g., Page and Shapiro 1992; Erikson, MacKuen, and Stimson 2002). Moreover, aggregate-level analysis suits our substantive research objectives well: politicians care primarily about the aggregation of individual opinions and the extent to which movement in aggregate opinion carries over to evaluations of their personal performances. Finally, our approach allows us to analyze very detailed (weekly) fluctuations in public opinion.

An important departure from previous studies is that we estimate a model that allows us to use multiple indicators to measure the same concept. This approach increases the points in time for which we have observations. Moreover, it makes the analysis less dependent on idiosyncrasies such as the question wording of the most frequently asked item and less likely to mistake outliers due to sampling error for temporary shifts in opinion. Our approach does, however, require the assembly of a database of survey marginals and the development of a measurement model to estimate a common time series from these diverse indicators.

TABLE 1  
Overview of Data Used in Analysis

<i>Concept and Representative Question</i>	<i>Observations</i>	<i>Items</i>	<i>Weeks</i>	<i>Start Date</i>
<i>War support</i>				
"Do you think the situation in Iraq is/was worth going to war over, or not?"	420	25	187	8/6/2002
<i>Perceptions of war success</i>				
"In general, how would you say things are going for the U.S. in Iraq: very well, moderately well, moderately badly, or very badly?"	183	18	154	3/25/2003
<i>Bush approval on Iraq</i>				
"Do you approve or disapprove of the way George W. Bush is handling the situation with Iraq?"	246	10	164	1/16/2003
<i>Overall Bush approval</i>				
"Do you approve or disapprove of the way George W. Bush is handling his job as president?"	482	7	187	8/7/2002

#### A DATABASE OF SURVEY MARGINALS

The assembly of response frequency data from assorted surveys has been facilitated by organizations that maintain online archives of polling results, often sorted by issue area. We used two such sources: The Polling Report<sup>7</sup> and the *National Journal's* Poll Track.<sup>8</sup> The main criteria for including a question in our database were that it had to provide an unambiguous measurement of one of our concepts of interest, war support, and perceptions of war success and that the question be asked at least four times. For each question, we recorded dates, survey organization, question wording, answer categories, and marginals for each answer category. Most surveys were conducted over a period of three to four days. We assigned the survey to a week by the mid-date of fieldwork. The data ran until the first week of May 2006.<sup>9</sup> Table 1 provides an overview of the data that we collected for each time series.

We used time series on twenty-five different survey questions that tap the public's support for the decision to invade Iraq. The most frequently asked item (forty-seven times) was the CNN/Gallup question, "All in all, do you think the situation in Iraq

7. [www.pollingreport.com](http://www.pollingreport.com).

8. <http://nationaljournal.com>.

9. An earlier version, which analyzed data until the 2004 election, is available from <http://home.gwu.edu/~voeten/>.



is/was worth going to war over, or not?" The primary instruments for the other main organizations are as follows:

- "Do you think the U.S. made the right decision or the wrong decision in using military force against Iraq?" (Pew,  $n = 38$ )
- "All in all, considering the costs to the United States versus the benefits to the United States, do you think the war with Iraq was worth fighting or not?" (ABC,  $n = 36$ )
- "Looking back, do you think the United States did the right thing in taking military action against Iraq, or should the U.S. have stayed out?" (CBS,  $n = 36$ )
- "Do you support or oppose U.S. military action to disarm Iraq and remove Iraqi President Saddam Hussein?" (FOX,  $n = 23$ )
- "Do you think that the United States should or should not take military action to remove Saddam Hussein from power in Iraq?" (NBC,  $n = 15$ )

In all, the twenty-five questions generated 420 data points. We discuss below how we estimate a common trend from this diverse set of questions.

The data set for perceptions of war success comprised 183 observations for 154 weeks. The most frequently asked questions were the following: "How would you say things are going for the U.S. in its efforts to bring stability and order to Iraq?" (CBS,  $n = 37$ ), "How well is the U.S. military effort in Iraq going?" (Pew,  $n = 31$ ), and "In general, how would you say things are going for the U.S. in Iraq?" (CNN/Gallup,  $n = 25$ ).

The time series for approval of Bush on Iraq began in mid-January 2003 as too few organizations asked pertinent questions before that time. Most of our 246 observations came from the question, "Do you approve or disapprove of the way George W. Bush is handling the situation in Iraq?"<sup>10</sup>

In addition, we collected data on overall evaluations of Bush's performance as president. Overall job approval is measured using the following standard question: "Do you approve or disapprove of the job George W. Bush is doing as president?" This question was asked 482 times by seven survey organizations.<sup>11</sup>

## METHOD

Seemingly innocuous variations in question wording can significantly influence survey responses, and the same question administered by different survey organizations can yield different marginals due to variation in sampling and survey procedures. The methodological task at hand is to disentangle genuine temporal changes from the effects of question wording and random noise. This task is complicated by the fact that the various items are not measured at the same points in time, necessitating a statistical model rather than simple standardizing as a way to separate these sources of variation (see also Stimson 1999). The intuition behind the method that

10. CNN/Time asked a variant: "Do you think President Bush is doing a good job or a poor job handling the situation in Iraq?" *Newsweek* originally asked, "Do you approve or disapprove of the way Bush is handling policies to deal with the threat posed by Iraq and its leader Saddam Hussein?"

11. Gallup/CNN/USA Today, ABC/Washington Post, CBS/New York Times, *Newsweek*, FOX News, NBC, and Pew.



we develop is that we seek to estimate the common trend underlying the various survey questions by factoring in that each question may have a different bias compared to some yardstick question and that some questions may more successfully tap our concept of interest than do others. A more formal exposition follows.

Let  $Y_{jt}$  denote the observed percentage supporting the war (or perceiving it to be successful, and so forth) on question  $j$  at time  $t$ .<sup>12</sup> Our objective is to estimate the changing level of support (or perceived success)  $\theta(t)$  that presumably underlies these observations. Each observation provides a noisy measurement of the public's true aggregate position at a specific time. The following model describes the observations:


$$Y_{jt} = \alpha_j + \beta_j \theta_t + \varepsilon_{jt}. \quad (1)$$

In equation (1),  $\alpha_j$  represents the "item bias," a parameter that corrects for the extent to which a question understates or overstates war support in comparison to an anchor item. For example, the FOX News question ("Do you support or oppose the United States having taken military action to disarm Iraq and remove Iraqi President Saddam Hussein?"), on average, generated an estimated 11–percentage point higher level of support for the war than did the Gallup question from Table 1, which mentions neither disarmament nor Saddam Hussein.

$\beta_j$  is a "weight" parameter that adjusts for the extent to which a question taps our main concept of interest. It is equivalent to a factor loading. Theoretically, we could estimate this parameter for each item. In practice, however, this would be difficult given the limited amount of data for each individual item. We therefore assume that this parameter is a constant except for items for which we have substantive grounds to suspect otherwise. The data set for war support includes questions that specifically mention ground troops and casualties. These questions may tap into different aspects of support for the war than do questions that make no explicit mention of either, which we account for by estimating separate weight parameters. In addition, the data set on perceptions of war success includes thirty-eight observations that tap expectations of future success—most notably, NBC's question "Do you feel more confident or less confident that the war in Iraq will come to a successful conclusion?" ( $n = 10$ ). Our measurement model allows for the possibility that trends in prospective evaluations behave differently from trends in contemporaneous assessments of war success.

The latent level of support follows a random walk process:  $\theta_t \sim N(\theta_{t-1}, \tau^2)$ , with a diffuse prior on  $\theta_t$ . This imposes dynamic structure on the estimation. Our specific assumption about persistence is motivated by arguments in the literature that in finite samples, series such as public mood or presidential approval can be expected to mimic the behavior of a unit root process (e.g., DeBoef and Granato 1997). The

12. The data are centered around zero. If there were multiple positive answer categories, we collapsed these. A focus on support implies the assumption that trends in opposition largely mirror approval. To check the validity of this assumption, we examined whether there were noticeable trends in the "don't know" categories, which there were not.



**Figure 1: Beliefs about the War**

random portion of the measurement error is described by  $\epsilon_{jt} \sim N(0, \sigma^2)$ . The models are estimated using Markov chain Monte Carlo (MCMC) methods and implemented in WinBUGS.

The models fit the raw data very well.<sup>13</sup> Moreover, the estimates correlate well with estimates from an alternative procedure: WCALC (Stimson 1999).<sup>14</sup> The main advantage of our approach over WCALC is that we obtain a series that is simultaneously smooth and allows for sudden shocks, whereas WCALC obscures sudden breaks in the data (and, thus, the short-term impact of events) or leads to seemingly random zigzagging (when the smoothing option is turned off). A Web appendix provides more detail on the method and comparisons to WCALC.

## EVENTS, CASUALTIES, AND BELIEFS ABOUT THE WAR

Figure 1 compares war support with Bush job approval on Iraq and perceptions of war success. All three series are scaled to represent the marginals from the

13. In an ANOVA analysis on the raw data controlling for fixed effects of the items, the adjusted  $R^2$  is over .90 for all series and around .95 for most series.

14. Bivariate correlations are .97 for war support, .96 for Bush Iraq approval, .94 for perceptions of success, and .99 for overall job approval.

main Gallup series.<sup>15</sup> The percentage of Americans who supported the war hovered at around 55 until late January 2003. Thus, it appears that the initial efforts by the administration to persuade the public of the war's necessity had little effect.<sup>16</sup> Support rose sharply starting around the time that Colin Powell gave the speech in the United Nations Security Council that signaled the beginning of the diplomatic endgame.<sup>17</sup> The speech itself did not lead to a one-time shift in support for the war. Instead, such support grew steadily as it became increasingly clear that the United States was relatively isolated in its insistence on the use of force. The initial apparent success of the war effort drove support even further upward, to a maximum of almost 70 percent. After that, support dropped gradually, although not smoothly, to between 35 and 40 percent in early 2006. The gradual decline was interrupted by what appeared to be pieces of good and bad news—in particular, the capture of Saddam Hussein and the revelations of the prison scandal in Abu Ghraib.

As can be expected, aggregate war support and perceptions of war success followed roughly the same pattern. Theoretically, perceptions of war success and war support clearly overlap at least to some degree. Our argument suggests, however, that they have subtly (but consequentially) different dynamics. In particular, we expect perceptions of war success to be more volatile and responsive to events than war support, but we expect war support to be more determinative of presidential evaluations. The patterns in the graph speak to the former expectation. The bivariate correlation between the weekly differences in war support and perceptions of war success is moderate (.52). Moreover, the relative levels of instability in the series followed our expectations. Although the means of the series are roughly equivalent, their variances differ considerably and as expected.<sup>18</sup> It is evident from glancing at the graph that perceptions of the war's success were more responsive to events than were evaluations of the war's merits. Within little more than a year, the estimated percentage of respondents believing that things are going very or moderately well in Iraq varied from a high of 78 percent in the first weeks after the start of the war to a low of 43 percent when the revelations about Abu Ghraib were prominently in the news and Nicholas Berg was beheaded (May 8, 2004). Even two and a half years after the initial invasion, aggregate perceptions of war success proved more sensitive to short-term considerations, such as the Iraqi elections in January 2005 and December 2005, than did aggregate war support.

15. For perceptions of war success, the Gallup question is "In general, how would you say things are going for the U.S. in Iraq?" The marginal represents the percentage of respondents who answered "very well" or "moderately well."

16. Much of the initial support for an invasion was apparently in place in the immediate aftermath of September 11, 2001. For example, on October 4, 2001, 73 percent of respondents answered favorably to the question "Do you favor or oppose the United States taking direct military action in Iraq?" ( $N = 1,005$ , Gallup and the *Los Angeles Times*).

17. "Remarks to the United Nations Security Council," Secretary Colin L. Powell, New York City, February 5, 2003.

18. Standard deviations are 9.2 (perceptions of success), 8.9 (Bush approval on Iraq), and 7.5 (war support).

Evaluations of Bush's performance in Iraq were similarly variable in the initial stages of the war, ranging from a high of 73 percent in April 2003 to 42 percent in April 2004. It appears that fluctuations in evaluations of Bush's performance steadied after the 2004 elections; overall, however, they appear to have been more responsive to short-term considerations than do evaluations of the war's merits.

### CASUALTY SENSITIVITY

A prediction derived from Mueller (1973) and pursued by others is that support for a war decreases with cumulative American casualties (deaths). Thus far, however, analysts have not examined the possibility that casualties may have differential effects on war support and perceptions of war success. In fact, most studies directly regress casualties on presidential approval without studying intermediate beliefs. Moreover, analyses tend to focus on highly aggregated time units—at best monthly data—rather than the more fine-grained weekly data that we study here.

A methodological problem in evaluating the impact of cumulative casualties on aggregate beliefs about the war is that the time series are nonstationary.<sup>19</sup> Hence, it is not advisable to regress the levels of these series on each other, as spurious results are likely to occur. Instead, we may analyze the differenced series. This, however, ignores the possible long-run relationships between variables: it may be that there is no immediate effect of casualties on beliefs about the war but rather a long-term equilibrium relationship where beliefs eventually become more negative as casualties increase. To evaluate both immediate and long-term equilibrating effects, we estimate a single-equation error correction model, which has become the model of choice in analyses of presidential approval and similar time series (e.g., Beck 1993; Clarke, Stewart, and Whiteley 1998; DeBoef and Kellstedt 2004).<sup>20</sup> The immediate effect of casualties is measured by the marginal change in the number of cumulative casualties from the preceding week (see Gartner and Segura 1998).<sup>21</sup> If casualties have a long-term equilibrium-correcting effect, then the lag of the cumulative casualties should also have an impact. As suggested by Mueller (1973), we take the natural log of cumulative casualties.

To model the time series properly, we need to include indicators for events that were unrelated to American deaths but that may have influenced war support and perceptions of war success. First, we include events that were politically significant (see Eichenberg and Stoll 2006 [this issue], for an alternative specification).

19. In this case, this is true by the definition of the measurement model.

20. In the bivariate case,  $\Delta Y_t = \beta_0 - \beta_1 (Y_{t-1} - \beta_2 X_{t-1}) + \beta + \Delta X_t + \epsilon_t$ . The single-equation model is often theoretically preferable to (Beck 1993) and more efficient than (DeBoef 2001) the two-step Engle-Granger method (Engle and Granger 1987). Moreover, it is less directly linked to the analysis of cointegrated series, as it can be derived from the autoregressive distributed lag model. Results from Prais-Winsten regressions are available from the authors. The main results are robust across different specifications.

21. The casualty data come from <http://icasualties.org>, which tallies CENTCOM and Department of Defense releases of combat and noncombat deaths.

Based on an evaluation of other literature and news summaries, we include indicators for the release of the Kay report, the capture of Saddam Hussein, the Abu Ghraib scandal, the election for a transitional National Assembly (January 2005), and the election for a full parliament (December 2005). Second, we include indicators for the most significant insurgency attacks that yielded the largest number of Iraqi casualties: the March 2004 suicide bombings at the Shi'a festival of Aashurah, the August 2005 stampede that killed nearly 1,000 Shi'a pilgrims, and the February 2006 bombing of a revered Shiite shrine in Samarra. All events are coded as dummy variables that take the value of 1 for a three-week period following each event, under the assumption that each event was so significant that it would have dominated debate for at least three weeks.<sup>22</sup>

Third, we include dummy variables to indicate the three periods of the war identified by Gelpi, Reifler, and Feaver (2005): "major combat" (until the first week of May 2003), "occupation," and "sovereignty," the period after the formal transfer of sovereignty to the interim Iraqi government in June 2004. Finally, we include a second-degree orthogonal (Hermite) polynomial to capture unmodeled trends, such as the gradual phasing out of the rally effect and the declining belief in a quick victory following from the passing of time.

Table 2 presents the results of regressions on weekly changes in war support and perceptions of war success as well as in presidential approval ratings (on Iraq and in general). Given lags, the model starts in the first week of April 2003. The lagged level of the dependent variable is included because high levels of support in the previous period may be correlated with a negative change in the following period due to equilibrium correction.

The results in Table 2 offer no evidence for an immediate relationship between weekly changes in casualties and beliefs about the war or the president. There is, however, some evidence for a long-term equilibrium relationship between casualties and perceptions about the war, although the coefficient is significant at conventional levels only for perceptions of war success. Moreover, the coefficient for lagged cumulative casualties on perceptions of war success is nearly three times as large as its coefficient on any of the other dependent variables. This confirms our hypothesis that perceptions of the war's success would be more responsive to casualties and events than would beliefs about the war's merits. Moreover, events and casualties account for 40% of the variation in perceptions of war success, whereas they explain only 25% of the temporal variation in war support.<sup>23</sup> This pattern suggests that swings in support for the war may have been determined to a greater extent by shifts in elite discourse that were relatively unconstrained by the happenings of the war.

It is also noteworthy that whereas events and casualties account quite well for variation in evaluations of Bush's performance on the Iraq issue, the model performs poorly at explaining variation in Bush's overall job approval ( $R^2_{adj} = .15$ ).<sup>24</sup> This suggests that the happenings of the war may have influenced overall evaluations of the president indirectly by shaping beliefs about the war.

TABLE 2  
Casualties, Events, and Evaluations of the War and the President ( $n = 152$ )

<i>Dependent Variable</i>	<i><math>\Delta</math> War Success</i>	<i><math>\Delta</math> War Support</i>	<i><math>\Delta</math> Bush Approval Iraq</i>	<i><math>\Delta</math> Overall Bush Approval</i>
Constant	42.15** (15.07)	13.95 (0.93)	13.56 (12.85)	9.03 (12.27)
Lagged level dependent variable	-0.08 (0.022)	-0.16** (0.04)	-0.09** (0.02)	-0.09** (0.03)
$\Delta$ Casualties	0.01 (0.01)	-0.01 (0.01)	-0.00 (0.00)	0.02 (0.01)
In Lagged cumulative casualties	-5.98** (2.09)	-2.10 (1.28)	-1.99 (1.78)	-1.42 (1.70)
Kay report	0.42 (0.65)	-0.20 (0.40)	-0.02 (0.55)	0.33 (0.51)
Saddam capture	4.62** (0.64)	1.86** (0.40)	3.20** (0.54)	1.74** (0.52)
Abu Ghraib	-0.13 (0.56)	-0.99** (0.35)	-1.22** (0.46)	-0.82 (0.44)
Iraqi election 1	1.57* (0.62)	0.66* (0.38)	0.17 (0.52)	0.06 (0.54)
Iraqi election 2	1.69** (0.64)	0.37 (0.40)	0.64 (0.54)	0.54 (0.51)
Festival	-1.09 (0.66)	-0.26 (0.41)	-0.08 (0.56)	0.06 (0.54)
Stampede	-0.16 (0.63)	-1.13** (0.39)	-1.23* (0.53)	-0.30 (0.50)
Mosque	0.44 (0.81)	0.19 (0.50)	0.03 (0.67)	0.31 (0.64)
Combat	2.67** (0.71)	0.84* (0.42)	1.68** (0.57)	0.77 (0.55)
Sovereign	-0.24 (0.47)	-0.31 (0.29)	0.36 (0.40)	0.33 (0.38)
First time counter	0.05* (0.02)	0.01 (0.01)	0.01 (0.02)	0.01 (0.02)
Second time counter	-1.8E-05** (0.64)	-2.8E-05 (4.07)	-5.1E-05 (5.39)	-4.3E-05** (5.17)
$R^2_{adj}$	0.406	0.252	0.320	0.152

NOTE: Standard error in parentheses.

\* $p < .05$ . \*\* $p < .01$ . All tests are two-tailed.

Abu Ghraib's strongest effect was on evaluations of Bush's job performance on Iraq, which declined by an estimated 6.1 percentage points as a consequence of

22. We coded Abu Ghraib as having a five-week effect given the prolonged discussion in the media.

23. We realize, of course, that these comparisons of fit statistics should be interpreted with some caution given that these are different dependent variables.

24. Note that, overall, the fit of models with first differences as dependent variables tends to be lower than the fit of models in which we regress directly on levels using the lagged dependent variable.

the scandal.<sup>25</sup> This may be an instance where an event revealed something about competence without having a discernable impact on perceptions of war success.

Gelpi, Reifler, and Feaver (2005) suggest that casualties may have had differential effects during the various periods of the war effort. We tested this by estimating separate models for the “occupation” and “sovereignty” periods.<sup>26</sup> There are several interesting results from this analysis.<sup>27</sup> First, we find—consistent with the argument made by Gelpi and his colleagues—that our model fits the data much better during the occupation period than during the sovereignty period. During the former period, the  $R^2_{adj}$  falls between .63 for perceptions of war success and .22 for Bush approval, whereas these numbers are .18 and .04 for the sovereignty period (the relative ordering remains unchanged). Perhaps this is because the public increasingly “tuned out” news about the war or because stronger priors rendered beliefs less susceptible to events and casualty reports after more than a year of war news.

We also find that during the occupation period, cumulative casualties were not a significant explanatory variable for any of the attitudes, whereas weekly changes in casualties had an impact on perceptions of war success (but not on war support or on evaluations of Bush’s performance). By our estimate, thirty-five casualties led to a 1–percentage point drop in aggregate perceptions of war success (the mean number of weekly casualties was fourteen, with a maximum of forty-five in this period). During the sovereignty period, weekly changes in casualties did not predict shifts in beliefs about the war or the president. There is, however, evidence for a significant long-term equilibrium relationship between the cumulative level of casualties and beliefs about both the war and the president. These effects should be interpreted with some caution: if we introduce a simple linear time counter into the model, the effect of casualties is no longer significant for any of the dependent variables. Thus, we cannot distinguish the effect of casualties from other dynamics that may lead to downward trends in evaluations of the war and the president.

### BELIEFS ABOUT THE WAR AND SUPPORT FOR THE PRESIDENT

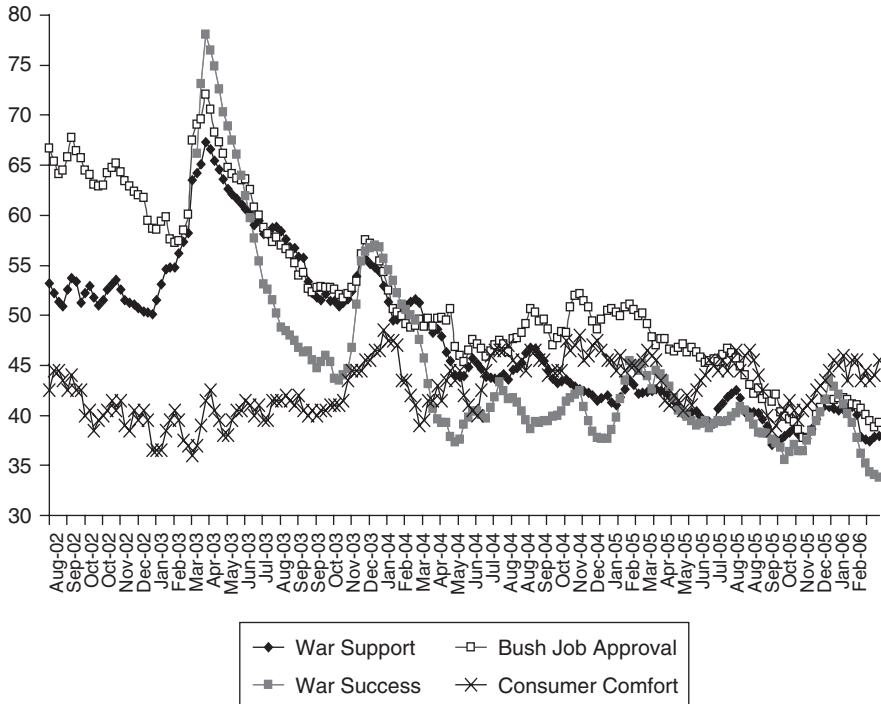
Did changes in beliefs about the war move overall presidential approval? If so, what sort of beliefs mattered most? To address these questions, we need to consider alternative accounts of presidential popularity. The most common explanation has been that evaluations of the president’s job performance are largely driven by subjective assessments of the economy, with objective economic conditions mattering only to the extent that they affect such assessments (e.g., MacKuen, Erikson, and Stimson 1992).

25. Abu Ghraib is coded as an event that lasted from weeks 17 to 21 in 2004.

26. The “combat period” yields only six weeks of data. We also estimated models with interactive effects between the periods and the casualty variables. The substantive implications were identical to those discussed here.

27. All robustness checks are reported in the Web appendix.





**Figure 2: Beliefs about the War, Consumer Comfort, and Presidential Approval**

Figure 2 graphs overall presidential job approval, war support, and a Consumer Comfort Index (CCI) constructed by *ABC/Money* magazine, which is available on a weekly basis (unlike other commonly used measures for subjective assessments of the economy).<sup>28</sup> Bush's overall job approval jumped by about 7.7 percentage points at the start of the war. This "rally around the flag" effect was not exceptional compared to the 9-percentage point average jump recorded for major responses to international crises (Oneal and Bryan 1995).<sup>29</sup> The visual evidence suggests that presidential job approval and war support tracked each other closely. Consumer

28. This is a rolling weekly average based on telephone interviews with about 1,000 adults nationwide each month. Fieldwork is by ICR Survey Research of Media, Pennsylvania. The published index ranged from +100 (everyone positive on all three measures) to -100 (everyone negative on all three measures), but we rescaled to the 0 to 100 interval for comparability. The three questions used to calculate the index are the following: "Would you describe the state of the nation's economy these days as excellent, good, not so good, or poor?" "Would you describe the state of your own personal finances these days as excellent, good, not so good, or poor?" "Considering the cost of things today and your own personal finances, would you say now is an excellent time, a good time, a not so good time, or a poor time to buy the things you want and need?" We also estimated a model that substituted Bush economic job approval as an independent variable. This generated similar results.

29. Some caution is warranted here in that we measured the change from the week before to the week of the actual intervention, whereas previous studies have used a much cruder time frame.

TABLE 3  
Weekly Change in Overall Presidential Approval by Beliefs about  
the War and the Economy (Since Start of War,  $n = 152$ )

	<i>Model 1</i>	<i>Model 2</i>	<i>Model 3</i>	<i>Model 4</i>	<i>Model 5</i>
Constant	0.390 (0.344)	-0.256 (0.200)	-0.339 (0.210)	0.400 (0.339)	-0.429 (0.419)
Bush approval $t - 1$	-0.066** (0.025)	-0.074* (0.034)	-0.079** (0.030)	-0.070** (0.024)	-0.168** (0.041)
$\Delta$ War support	0.706** (0.084)	—	—	0.596** (0.095)	0.380** (0.103)
Lag war support	0.097* (0.049)	—	—	0.104* (0.048)	0.005 (0.058)
$\Delta$ Bush approval Iraq	—	0.470** (0.060)	—	—	0.286** (0.068)
Bush approval Iraq $t - 1$	—	0.122** (0.038)	—	—	0.163** (0.045)
$\Delta$ War success	—	—	0.284** (0.051)	0.124* (0.052)	0.103* (0.051)
War success $t - 1$	—	—	0.028 (0.020)	—	-0.033 (0.020)
$\Delta$ CCI	0.013 (0.026)	0.028 (0.026)	0.016 (0.029)	0.004 (0.026)	-0.001 (0.024)
CCI $t - 1$	0.007 (0.016)	-0.006 (0.016)	-0.002 (0.017)	-0.009 (0.015)	-0.010 (0.014)
Time counter 1	0.003 (0.003)	-0.001 (0.002)	-0.003 (0.002)	0.004 (0.003)	0.000 (0.004)
Time counter 2	-3.9E-005* (1.9E-005)	-3.5E-005* (1.5E-005)	-1.2E-005 (1.4E-005)	-0.6E-005 (1.1E-005)	-3.5E-005* (1.8E-005)
$R^2_{adj}$	0.347	0.338	0.203	0.367	0.449
SE estimate	0.728	0.734	0.805	0.717	0.669

NOTE: Standard error in parentheses. CCI = Consumer Comfort Index.

\* $p < .05$ . \*\* $p < .01$ . All tests are two-tailed.

comfort, on the other hand, followed a different path. In fact, the correlations between consumer comfort and overall presidential approval are negative and significant, although this relationship is weaker in the period after the war's start.<sup>30</sup> It is plausible that the anticipation of the war depressed consumer confidence.<sup>31</sup>

Table 3 presents the results of a regression analysis using error correction models similar to the ones from Table 2. The dependent variable is the weekly change in presidential approval ratings. The lagged level of presidential approval is included to capture equilibrium corrections. The other independent variables are weekly changes in various beliefs about the war and consumer confidence. The substantive

30. The Pearson  $R = -.41$  for the entire period and  $-.24$  after the war's start.

31. See DeBoef and Kellstedt (2004) on the point that politics shapes economic perceptions.

results are robust to the inclusion of evaluations of Bush's job performance on the economy as an alternative control for economic perceptions. To ensure comparability, the analysis is limited to the period after the war started.

Most important, weekly shifts in war support had a strong immediate effect on changes in presidential approval. A 1–percentage point increase in support for the war led to a .71–percentage point immediate increase in overall presidential approval. By contrast, the corresponding effects for perceptions of war success and Bush approval on Iraq were .28 and .47, respectively. There is also some evidence for significant long-term equilibrium-correcting relationships between presidential job approval and war support. In addition, the model that includes war support explains the data better than do the other two models: model 1 has the highest  $R^2_{adj}$  and yields the lowest standard error in its estimate of weekly changes in presidential job approval.

This difference is especially clear when we compare models 1 and 3. Given that weekly shifts in war support and perceptions of war success are only moderately correlated and appear to have somewhat different determinants (see Table 2), we estimated a model (model 4) that includes both variables in the same model. Weekly shifts in perceptions of war success had a significant independent effect on presidential job approval, but the effect was very small (.12) in comparison to that for shifts in war support (.60). We also present the results of an equation that includes all the variables (model 5). A concern here is that the high bivariate correlations between the lagged levels of the various belief variables,<sup>32</sup> in combination with the relatively small number of observations, make the interpretation of the lagged terms unreliable. The estimates on the first differences confirm the findings of the other models. Together, these results confirm our expectation that shifts in aggregate beliefs about the war's merits would be more consequential politically than would shifts in beliefs that the war was going well, although the latter were more responsive to events and casualties.

The finding that shifts in war support also relate more strongly (if slightly) to presidential approval than do shifts in approval for Bush's job performance on Iraq is surprising given that the wording of the questions for overall job approval and job approval on Iraq differed only slightly. We take this as strong evidence for the consequences of shifts in aggregate beliefs about the merits of the war.

A model that includes only the event and casualty variables (see Table 2) explains little variation in weekly approval ratings in comparison to the models presented in Table 3. Moreover, the event variables are jointly insignificant when we include events and casualties in the same model as aggregate war support.<sup>33</sup> These findings support our claim that we should look beyond the standard approach of linking events and casualty figures directly to approval ratings if we seek to capture the political impact of war.

32. The correlation is .95 for Bush Iraq approval and war support; it is .89 for war support and perceptions of war success. As mentioned earlier, the bivariate correlations between first differences are around .5 for these series.

33. Results available from authors.

All of this implies, of course, that a next logical step would be to estimate a multiple-equation model in which aggregate war support is a function of events and, more important, some measure of elite discourse. Or, we may stipulate that events shape perceptions of war success, which then, together with elite discourse, shape war support, which is ultimately a potent predictor of presidential approval. We leave the precise specification of such a multiple-equation model and the development of an accurate measure for elite discourse as a future exercise. We did test whether endogeneity renders the ordinary least squares (OLS) estimates from Table 3 inconsistent but found no evidence for this in a Durbin-Wu-Hausman test.<sup>34</sup>

Controlling for beliefs about the war, changes in consumer confidence had the expected positive impact on presidential approval ratings, but these effects were not significant at conventional levels (see also Eichenberg and Stoll 2006 [this issue]). When we replace this measure with evaluations of Bush's job performance on the economy, the main result stands: shifts in war support were more important in explaining shifts in presidential approval ratings than were shifts in perceptions of the economy.

## CONCLUSIONS

Our two main findings are (1) that events and casualties accounted for considerable variation in perceptions of the war's success but did not explain as much variation in support for the war and the president and (2) that shifts in war support accounted for shifts in presidential approval ratings better than did shifts in perceived war success or even shifts in Bush's job approval on Iraq. These findings suggest that the public changed its collective mind on whether the war was going well in response to events and casualties, but the results also suggest that opinion on this topic was of relatively little consequence for presidential support. In contrast, changes in public support for the war had immediate and considerable consequences for the president (a finding that fits well with individual-level research on the relative consequences of war support and perceptions of war success for vote choice;<sup>35</sup> see Gelpi, Reifler, and Feaver 2005). Yet here we found weaker relationships between key events and casualty reports, on one hand, and war support, on the other.

How, then, does the public evaluate and reevaluate presidential decisions to go to war? Like other observers, we suspect that war support is shaped in part by the signals, or cues, that elites send and that these signals are only partly constrained by reality. Research on previous wars—including World War II, the Korean War, the

34. We ran a Durbin-Wu-Hausman test (Davidson and MacKinnon 1993) for the two-equation model that regresses differences in war support on presidential approval and where differences in war support are a function of events. We could not reject the hypothesis that the ordinary least squares estimates are consistent,  $F(1, 79) = .08$ ,  $p = .78$ . They are similar for changes in war success,  $F(1, 79) = 2.0$ ,  $p = .16$ .

35. On the other hand, Gelpi, Reifler, and Feaver (2005) argue that expectation of success is more important than war support in shaping casualty tolerance.

Vietnam War, and the 1991 Gulf War—indicates that the nature of elite discourse about a war is a key determinant of changing patterns in both war support itself (Berinsky 2004; Larson 1996; Larson and Savych 2005; Zaller 1992) and presidential approval during wartime (Brody 1991, 1994). Of particular relevance to our findings is the notion that the extent of elite consensus shapes public support for a war. A plausible interpretation of the trend in support for the war in Iraq is that increasing elite polarization regarding the war produced not only polarization among the public (with Democratic and Republican partisans diverging on its wisdom; see Berinsky 2004) but also declining support for the war in the aggregate.

Another question is whether our conclusions can be extended to other wars, particularly those with differing principal policy objectives (see, e.g., Eichenberg 2005). In the case of the 1991 Gulf War, the short (six-week) span of the war and the limited variance in public opinion about the war and the president during this span (see Mueller 1994) render our aggregate time-series approach impractical. For earlier wars of longer durations, the relative paucity of data may do the same. Still, our analyses of the war in Iraq suggest, at the very least, additional caution in judging claims regarding the unmediated effects of casualties and events on support for some of these previous wars (for a complementary argument based on individual-level analysis, see Berinsky 2004).

Even taking into account its limitations, our study has advanced the literature on the relationship between accountability and war considerably. First, we have presented direct evidence regarding the role of aggregate beliefs about a war in shaping aggregate support for a commander in chief. The results for the Iraq war suggest that previous analyses may have underestimated this impact to a considerable degree given that they tended to rely on casualties and events as proxies for aggregate beliefs about the war. Second, we have distinguished between two models for understanding presidential accountability during wartime and tested these models. In doing so, we have produced evidence that the “decision maker” model of accountability works better than the “managerial” model. These insights should lead to further theoretical and empirical inquiry into what kinds of subjective beliefs about a war are politically consequential and what shapes these beliefs. Moreover, our competing accountability models have implications for the burgeoning literature on democracies and war: although a wide range of results suggests that democracies behave differently in choosing and fighting wars, there is still a wide gap in understanding the accountability mechanisms that generate these results (see Reiter and Stam 2002). Third, we have developed a new measurement model for estimating time series from survey marginals. This model may prove useful for other applications, especially when data are sparse, given that it allows one to combine data from various survey instruments. Finally, our analyses provide insights regarding the dynamics of public opinion about a war that is interesting in and of itself—one that is ongoing and for which questions regarding presidential accountability have been and will continue to be of great interest in the scholarly community and beyond.

## REFERENCES

- Barro, Robert. 1973. The control of politicians: An economic model. *Public Choice* 14 (1): 19-42.
- Beck, Nathaniel. 1993. The methodology of cointegration. In *Political analysis*, vol. 4, edited by John R. Freeman, 237-48. Ann Arbor: University of Michigan Press.
- Bennett, James. 2003. In debate on foreign policy, wide gulf or splitting hairs? *The New York Times*, September 30, A1.
- Berinsky, Adam J. 1999. The two faces of public opinion. *American Journal of Political Science* 43 (4): 1209-30.
- . 2002. Silent voices: Social welfare policy opinions and political equality in America. *American Journal of Political Science* 46 (2): 276-87.
- . 2004. Assuming the costs of war: Events, elites, and the American public. Unpublished manuscript, Massachusetts Institute of Technology, Cambridge.
- Brody, Richard A. 1991. *Assessing the president: The media, elite opinion, and public support*. Stanford, CA: Stanford University Press.
- . 1994. Crisis, war, and public opinion. In *Taken by storm: The media, public opinion, and U.S. foreign policy in the Gulf War*, edited by W. Lance Bennett and David L. Paletz, 210-27. Chicago: University of Chicago Press.
- Boettcher, William A., III, and Michael D. Cobb. Echoes of Vietnam? Casualty framing and public perceptions of success and failure in Iraq. *Journal of Conflict Resolution* 50 (6): 831-54.
- Bueno de Mesquita, Bruce, Alastair Smith, Randolph M. Siverson, and James D. Morrow. 2003. *The logic of political survival*. Cambridge, MA: MIT Press.
- Clarke, Harold D., Marianne Stewart, and Paul Whiteley. 1998. New models for new Labour: The political economy of Labour Party support, January 1992–April 1997. *American Political Science Review* 92 (3): 559-75.
- Davidson, Russell, and James G. MacKinnon. 1993. *Estimation and inference in econometrics*. New York: Oxford University Press.
- DeBoef, Suzanna. 2001. Modeling equilibrium relationships: Error correction models with strongly autoregressive data. *Political Analysis* 9 (1): 78-94.
- DeBoef, Suzanna, and Jim Granato. 1997. Near-integrated data and the analysis of political relationships. *American Journal of Political Science* 41 (2): 619-40.
- DeBoef, Suzanna, and Paul M. Kellstedt. 2004. The political (and economic) origins of consumer confidence. *American Journal of Political Science* 48 (4): 633-49.
- Downs, George W., and David M. Rocke. 1994. Conflict, agency, and gambling for resurrection: The principal-agent problem goes to war. *American Journal of Political Science* 38 (1): 362-80.
- Eichenberg, Richard C. 2005. Victory has many friends: U.S. public opinion and the use of military force, 1981-2005. *International Security* 30 (1): 140-77.
- Eichenberg, Richard C., and Richard J. Stoll. 2004. *The political fortunes of war: Iraq and the domestic standing of President George W. Bush*. London: The Foreign Policy Centre.
- . 2006. War president: The approval ratings of George W. Bush. *Journal of Conflict Resolution* 50 (6): 783-808.
- Engle, Robert F., and Clive W. J. Granger. 1987. Co-integration and error correction: Representation, estimation, and testing. *Econometrica* 55 (2): 251-76.
- Erikson, Robert S. 2004. Macro vs. micro-level perspectives on economic voting: Is the micro-level evidence endogenously induced? Paper presented at the annual meeting of the Society for Political Methodology, July, Stanford, CA.
- Erikson, Robert S., Michael B. MacKuen, and James A. Stimson. 2002. *The macro polity*. New York: Cambridge University Press.
- Feaver, Peter, and Christopher Gelpi. 2004. *Choosing your battles: American civil-military relations and the use of force*. Princeton, NJ: Princeton University Press.

- Ferejohn, John. 1986. Incumbent performance and electoral control. *Public Choice* 50 (1): 5-26.
- Gartner, Scott Sigmund, and Gary M. Segura. 1998. War, casualties and public opinion. *Journal of Conflict Resolution* 42 (3): 278-300.
- Gaubatz, Kurt Taylor. 1999. *Elections and war: The electoral incentive in the democratic politics of war and peace*. Stanford, CA: Stanford University Press.
- Gelpi, Christopher, Peter D. Feaver, and Jason Reifler. 2006. Success matters: Casualty sensitivity and the war in Iraq. *International Security* 30 (3): 7-46.
- Gelpi, Christopher, Jason Reifler, and Peter D. Feaver. 2005. Iraq the vote: Retrospective and prospective foreign policy judgments on candidate choice and casualty tolerance. Unpublished manuscript, Duke University, Durham, NC.
- Halbfinger, David M. 2004. Kerry attacks Bush over loss of explosives. *The New York Times*, October 27, A17.
- Hess, Gregor D., and Athanasios Orphanides. 1995. War politics: An economic, rational-voter framework. *American Economic Review* 85 (4): 828-46.
- Keech, William. 1995. *Economic politics: The costs of democracy*. Cambridge, UK: Cambridge University Press.
- Kinder, Donald R., and D. Roderick Kiewiet. 1979. Economic grievances and political behavior: The role of personal discontents and collective judgments in congressional voting. *American Journal of Political Science* 23:495-515.
- Kramer, Gerald. 1983. The ecological fallacy revisited: Aggregate vs. individual-level findings on economics and elections, and sociotropic voting. *American Political Science Review* 77 (1): 92-111.
- Kull, Steven, Clay Ramsay, and Evan Lewis. 2004. Misperceptions, the media, and the Iraq War. *Political Science Quarterly* 118 (4): 569-98.
- Larson, Eric V. 1996. *Casualties and consensus: The historical role of casualties in domestic support for U.S. military operations*. Santa Monica, CA: RAND.
- Larson, Eric V., and Bogdan Savych. 2005. *American public support for U.S. military operations from Mogadishu to Baghdad*. Santa Monica, CA: RAND.
- MacKuen, Michael B., Robert S. Erikson, and James A. Stimson. 1992. Peasants or bankers? The American electorate and the U.S. economy. *American Political Science Review* 86 (3): 597-611.
- Mueller, John E. 1970. Presidential popularity from Truman to Johnson. *American Political Science Review* 64 (1): 18-34.
- . 1973. *War, presidents, and public opinion*. New York: John Wiley.
- . 1994. *Policy and opinion in the Gulf War*. Chicago: University of Chicago Press.
- Oneal, John R., and Anna Lillian Bryan. 1995. The rally 'round the flag effect in U.S. foreign policy crises, 1950-1985. *Political Behavior* 17 (4): 379-401.
- Page, Benjamin I., and Robert Y. Shapiro. 1992. *The rational public: Fifty years of trends in Americans' policy preferences*. Chicago: University of Chicago Press.
- Reiter, Dan, and Allan C. Stam. 2002. *Democracies at war*. Princeton, NJ: Princeton University Press.
- Russett, Bruce. 1990. *Controlling the sword*. Cambridge, MA: Harvard University Press.
- Shane, Scott. 2005. Bush's speech on Iraq War echoes voices of an analyst. *The New York Times*, December 4.
- Stimson, James A. 1999. *Public opinion in America*. 2nd ed. Boulder, CO: Westview.
- Stimson, James A., Michael Mackuen, and Robert S. Erikson. 1995. Dynamic representation. *American Political Science Review* 89 (3): 543-65.
- Stoll, Richard. 1984. The guns of November: Presidential reelections and the use of force, 1947-82. *Journal of Conflict Resolution* 28 (2): 231-46.
- Zaller, John R. 1992. *The nature and origins of mass opinion*. New York: Cambridge University Press.