

Using Markets to Inform Policy: The Case of the Iraq War

Justin Wolfers

Wharton, CEPR, IZA and NBER

jwolfers@wharton.upenn.edu

www.nber.org/jwolfers

Eric Zitzewitz

Stanford GSB

ericz@stanford.edu

<http://faculty-gsb.stanford.edu/zitzewitz>

Abstract

Market prices incorporate large amounts of information, and our aim in this paper is to demonstrate that prediction markets can help extract this information, prospectively allowing this aggregated expertise to inform policy decisions in real-time. We provide a case study, exploiting data from a market trading in contracts which paid off if Saddam Hussein was removed as leader of Iraq, to learn about financial market participants' expectations of the consequences of the 2003 Iraq war. We conducted an ex-ante analysis, which we disseminated before the war, finding that a 10 percent increase in the probability of war was accompanied by a \$1 increase in spot oil prices that futures markets suggested was expected to dissipate quickly. Equity prices movements implied that the same shock led to a 1½ percent decline in the S&P 500. Further, the existence of widely-traded options allows us to back out the entire distribution of market expectations of the war's near-term effects, finding that these large effects reflected a negatively skewed distribution, with a substantial probability of an extremely adverse outcome. The flow of war-related news through our sample explains a large proportion of daily oil and equity price movements. Subsequent analysis suggests that these relationships continued to hold out-of-sample. Our analysis also allows us to characterize which industries and countries were most sensitive to war news, and when the war turned out somewhat better than ex-ante expectations, these sectors recovered, confirming these cross-sectional implications. We highlight the particular features of this case study that make it particularly amenable to this style of policy analysis, and discuss some of the issues in applying this method to other policy contexts.

October 2004

This paper incorporates material from "What Do Financial Markets Think About War in Iraq?" that will not be published separately. Special thanks to Andrew Leigh for his contributions throughout this project. We would also like to thank Tim Bresnahan, Jeremy Bulow, John Campbell, John Cochrane, Steve Davis, Macgregor Duncan, Douglas Elmendorf, Martin Feldstein, Peter Henry, Dean Karlan, Ilan Kremer, Lars Lefgren, Steve Levitt, Ulrike Malmendier, Alfredo Marti, John McMillan, Paul Oyer, Charles Ploser, Charles Plott, Jim Poterba, Jon Reuter, Andy Skrzypacz, Scott Stern, Robert Topel, Hal Varian, Romain Wacziarg, Barry Weingast, Jeff Zwiebel and seminar participants at the AEA annual meetings, Berkeley, Columbia, Harvard, Kansas, the NBER, the New School, Oregon, Stanford, and UT-Austin for stimulating discussions. John Delaney and William Saletan provided insights into their proprietary data. Daphne Chang, Jan Driscoll, Doug Geysler, Catherine Haga, Allison Kavanagh, Peter Latusek, Pauline Mena, Paul Reist, Nora Richardson, Rosalind Young, and Simone Yu provided outstanding research assistance. Wolfers gratefully acknowledges the support of a Hirtle, Callaghan & Co. – Arthur D. Miltenberger Research Fellowship, and the support of the Zull/Lurie Real Estate Center.

1. Introduction

When a government weighs a policy decision with important economic consequences, investors and other economic agents struggle to assess its likely effects. Normally, this struggle is conducted individually or in small groups and its outcome is not credibly observable to economists or policymakers. Experts can disseminate estimates of the likely effects of the policy, but their professed assessments suffer from a cheap talk problem and may be tainted by political expediency.

This paper presents a method for extracting some of the rich information embedded in financial market prices using a prediction market—a small-scale market in securities whose payoffs are explicitly contingent on the policy being implemented. The prices of these prediction market securities provide an indication of the participants' estimates of the probability of the contingency occurring. Examining the correlation of these prices with securities tracking other economic or financial variables can, under specific identifying assumptions, yield insight into investors' beliefs about the impact of the policy. In markets, investors' beliefs matter only to the extent that they back them with money, and this helps resolve the cheap talk problem.

The specific example we examine is the 2003 war with Iraq. We exploit a financial instrument that we refer to as the “Saddam Security”—an asset whose payoffs depend on the ousting of the Iraqi leader by a specific date. These securities traded on an online betting exchange similar to the better-known Iowa Electronic Markets. The Saddam Security traded more widely and with greater liquidity than the typical contract on the Iowa market, and Wall Street is well-represented among market participants. Evidence summarized in section 2 suggests that these data seem to reflect underlying war probabilities with reasonable accuracy.

During the period leading up to the war, spot oil prices were strongly positively correlated with the probability of war, and equity prices were negatively correlated. If we assume that the source of this covariance is news about the probability of war affecting financial markets, as we argue below appears justified in this case, its magnitude is informative about war's expected effects.

We constructed an ex-ante estimate of war's expected effects about a month prior to the war and circulated a working paper version a week thereafter (Leigh, Wolfers, and Zitzewitz, 2003). We found that a 10 percentage point increase in the probability of war was associated with a \$1 per barrel increase in the spot oil price. Long-run oil futures suggested that the effect on oil was expected to be temporary (largely over by the end of 2003) and that the expected long-run effect was a slight reduction in oil prices.

Turning to equity markets, we found that the U.S. stock market was extremely sensitive to changes in the probability of war. A 10 percentage point rise in the probability of war was accompanied by a 1.5 percent decline in the S&P 500, implying, under our identifying assumptions, an anticipated average effect of -15 percentage points. To understand this large expected average effect, we examined S&P option prices to back out the distribution of the expected effects of the war. This distribution was negatively skewed. It implied a 70 percent probability that war would have a moderately negative effect of 0 to -15 percentage points, a 20 percent probability of a -15 to -30 percentage-point effect, and a 10 percent chance of even larger declines.

Presumably, a 30 percentage point reduction in the market value of U.S. firms' equity would require some of the dire pre-war predictions about nuclear terrorism or the wholesale destruction of the Iraqi, Kuwaiti, and Saudi oil industries to be borne out. Fortunately, none were. As the three-week war progressed, the S&P 500 rallied by about 4 percent, reversing about one quarter of the estimated pre-war discount. In the first week of the war, as coalition forces captured Iraqi oil fields largely intact, the oil price fell by about \$7, reversing about two-thirds of the pre-war estimated effect. Of course, with a single war and thus a single observation, we cannot assess whether the discount built into the market was irrationally large, as some have argued, or whether we simply got a relatively benign draw from the distribution. What we can say with more certainty is that war had investors worried.

This analysis is of potential long-term interest for three reasons. First, it is the first attempt that we know of to use prediction and financial market data to better understand the consequences of a prospective policy decision—a U.S.-led invasion of Iraq—in real time. A financial-market-based analysis complements expert opinion in several ways. Expert opinions tended to vary widely and

vary fairly directly with the ideological predisposition of the predictor. Markets aggregate opinions, and by requiring a trader to put “your money where your mouth is,” they lessen the cheap-talk problem and create incentives for individuals to reveal their true beliefs. In addition, whereas experts tended to focus on the most analyzable costs of a policy (e.g., Nordhaus, 2002), financial markets are forced to price the harder-to-assess but potentially much larger general equilibrium and political effects on the global economy. Financial markets do not simply evaluate the cost of war today, but also incorporate the effect of this war on the number and intensity of future conflicts.¹ Moreover, whereas previous studies of the effect of political events on financial markets are necessarily retrospective, our analysis is (or was) prospective.² Analyses like ours could conceivably be used to inform decision-making in real time.³

Second, analyses such as ours may provoke a reassessment of the extent to which stock market movements can be explained by news. The conventional wisdom is that identifiable news events explain only a small portion of market movements, suggesting a role for behavioral theories of stock price movements. The wisdom is based in part on studies like Cutler, Poterba, and Summers (1989), who find that first-order news events (including political and military developments) explain only a small portion of market movements. In contrast, we find that over 30 percent of the variation in the S&P and 75 percent of the variation in spot oil prices between September 2002 and February 2003 can be explained econometrically by changes in the probability of war (and, for oil, the Venezuelan crisis).⁴

A possible reconciliation of our results with previous narrative studies of financial markets is that many political events are not surprises by the time they happen. This was the case for the 2003 Iraq

¹ They thus address Davis, Murphy, and Topel’s (2003) critique that many cost of war estimates were not compared against the costs of the alternative policies.

² Examples of retrospective analyses of financial market effects of war include studies of British bond prices during World War I by Elmendorf, Hirschfield, and Weil (1996), of Swiss and Swedish government bonds during World War II by Frey and Kucher (2000) and Waldenström and Frey (2002). Retrospective analysis of the market effects of U.S. Presidential candidates’ fortunes include Slemrod and Greimel (1999) and Knight (2003). There is also a large number of event studies examining stock market responses to regulatory or tax policy changes. See Wolfers and Zitzewitz (2004a) for a more complete review of the literature on prediction markets.

³ Two caveats are in order. First, oil and equity prices are obviously imperfect proxies for social welfare. While one could imagine prediction market securities tracking better proxies, such as GDP, in their absence we use the best proxies for war’s effect that we have access to. Second, if policy makers did start basing decisions on analyses of market expectations about outcomes, this could create manipulation and endogeneity issues, complicating the analysis. We return to this issue in the final section.

⁴ The figures referred to are the R-squared from longest difference regressions reported in Table 1.

war: by the time the final ultimatum to Saddam was issued on March 17, the market's assessment of the probability of war was already above 90 percent, and its expected stock market effects were largely priced in. It was even partly the case for a "surprise" like Pearl Harbor, since the likelihood of war with Japan and Germany was surely somewhat foreseeable even on December 6, 1941. Without securities that quantify the news content of the political narrative, the true impact of these events on the markets is almost impossible to assess.⁵

Third, our analysis highlights the likely utility of political securities in aggregating information about specific risks in financial markets, which may in turn improve the efficiency of these markets. If some of the uncertainty about the value of the S&P prior to the war reflected uncertainty about war and its impact, pricing the likelihood of war in a publicly observable market should enhance the efficiency with which information about the non-war component of value is incorporated into prices. Analogously, Yuan (2002) finds that the introduction of sovereign debt markets increases the liquidity of emerging-market corporate bonds. If this study upwardly revises our beliefs about the extent to which political uncertainty contributes to uncertainty about asset values, then it should also revise our beliefs about the utility of political risk securities, perhaps even including the so-called "terrorism futures" proposed in the Summer of 2003 by the Defense Advanced Research Projects Agency (DARPA, 2003). Indeed, the Saddam Security is the first real test of using such markets to price geopolitical risks.

The remainder of the paper is divided into four sections. The next section describes the Saddam Security and outlines our methodology. The third section reviews our ex-ante estimates of the effect of war on oil and equity prices, respectively. The fourth section compares these ex-ante estimates against actual market movements during the four-week war. We conclude by discussing the issues involved in applying this method to other policy decisions.

⁵ A partial exception to this is Rigobon and Sack (2003), who construct an estimate of the proportion of market volatility explained by war news using a methodology that compares the volatility on days with and without identifiable war news. This very different method yields results that are roughly consistent with our own. Without a means of quantifying the news content, however, Rigobon and Sack can not obtain an estimate of war's expected effects, although they can speak to the relative sensitivity of various financial instruments.

2. Background

The innovation of this paper is to exploit the Saddam Security, a contingent security that paid \$10 if and only if Saddam Hussein was removed from office by a certain date; we focus on the June security, which paid if Saddam was out by June 30, 2003. Out-of-office was defined as no longer controlling the center of Baghdad.⁶ We use this as a reasonably close, albeit imperfect, proxy for the probability of a war against Iraq in the winter of 2002-3.⁷ Doing so requires assuming that it was widely believed before the war that: 1) that Saddam was unlikely to perish or relinquish power unless war was imminent and 2) that war, if undertaken that winter, would be successful by the end of June. The second assumption is clearly not exactly true: when coalition forces crossed Iraq's border on March 19, the June Saddam traded at 95, suggesting that traders placed a 5 percent probability on Saddam holding out. For simplicity, we do not make the small adjustment implied by this probability.

Saddam Securities trade on Tradesports.com, an on-line betting exchange.⁸ Tradesports is an electronic exchange that is similar in many respects to the better known Iowa Electronic Markets. The market operates as a continuous double auction, with free entry and the possibility to place limit as well as market orders and take either long or short positions. Differences exist, however. Tradesports is a for-profit company and charges a 0.4% commission. Setting up an account on Tradesports is arguably easier, and Tradesports does not limit individual investments to \$500 as Iowa does. The most important difference, however, is clientéle. Tradesports offers primarily contracts on sporting events, but it markets itself as a trading exchange in contingent commodities

⁶ Although Tradesports did define what it meant by out of office, this outcome is clearly less contractible than is typical in futures markets. Had Baghdad fallen close to the expiry date of a security, a difficult judgment call might have been required. Partly to ensure confidence in the integrity of any such calls, Tradesports commits to not hold any positions in securities traded on its exchange (in contrast to a traditional bookmaker). A challenge in designing prediction securities is to balance contractibility with capturing the uncertainties in which participants are interested; a Tradesports contract on whether the UN would authorize force against Iraq, an arguably more contractible outcome, attracted little volume.

⁷ Formally, the price of the Saddam Security is a state price, which may be different from the subjective probability belief of the marginal investor if marginal utilities of wealth differ in the peace and war states. In what follows, we assume that Saddam Security traders are not using them to hedge the systematic component of war risk and thus prices can be interpreted as probability beliefs of the marginal investor.

⁸ Many thanks to the CEO of TradeSports, John Delaney, for generously sharing these data and his expertise throughout this project.

rather than as a sportsbook. Reflecting this, many of their active traders are based on Wall Street or in the City of London.⁹

During the ex-ante period of our analysis, monthly Saddam Security trading volumes were just over \$10,000, comparable to the Iowa presidential winner-take-all markets and about 6-7 times average volumes in the vote share markets. Volumes increased substantially in the last month before the war, and over the entire period that the markets were open, approximately \$1.2 million was traded.¹⁰ Numerous studies have suggested that the Iowa markets are liquid enough to provide meaningful prices.¹¹ Our tests of the efficiency of the pricing of the Saddam Security yield results consistent with these tests.¹² At the same time, the Saddam Security prices do exhibit bid-ask bounce and some slowness in their incorporation of information, and we take steps outlined below to mitigate the resulting attenuation bias in our estimates.

As an additional check of the plausibility of the Saddam Security prices, we compare them to the only other quantitative expert assessment we could find. William Saletan, as part of a weekday Slate.com column called “The Odds of War”, provided a daily quantitative probability assessment he referred to as the “Saddameter.” Figure 1 plots the Saddameter and the daily prices for the Saddam securities expiring in December 2002 and March and June 2003. The direction of movements is broadly consistent, especially on days with important war-related news reported in the *New York Times*.¹³ The Saddameter is consistently higher than even the June Saddam, conceivably

⁹ We base this statement on a conversation with and media comments by John Delaney, who is privy to traders’ contact telephone numbers and emails. Without access to this information ourselves, we of course cannot independently confirm it, but in an informal survey we ran on *forum.tradesports.com*, over half of respondents had trading or finance-related jobs.

¹⁰ At the time, this was the largest volume traded in a single prediction market that we are aware of. It has since been surpassed by the \$8 million traded on Tradesports’ Bush 2004 reelection contract.

¹¹ Specifically, prices follow a random walk, there is no evidence that exploiting polling data yields profitable trading strategies, and the final predictions from even the less liquid vote share markets have tended to outperform opinion polls. Berg, Forsythe, Nelson, and Reitz (2001) and Berg, Nelson, and Reitz (2001) provide useful reviews of these findings. Additionally, Strumpf (2004) reports that the price impact from random \$500 trades dissipates within 24 hours, suggesting that even fairly substantial noise trading is unlikely to move prices much from their equilibrium level.

¹² These tests are reported in Section 3 of Leigh, Wolfers, and Zitzewitz (2003). To summarize, augmented Dickey Fuller tests do not reject a random walk hypothesis, and KPSS tests do reject the null that prices are trend stationary. Negative first-order autocorrelation of price changes suggests some bid-ask bounce, while longer lags suggest no further predictability of price changes based on past price movements. Reactions to identifiable news about the likelihood of war appear to suggest that some information is incorporated with a one-to-two day lag. Tetlock (2004) also conducts tests of the efficiency of Tradesports prices and concludes that non-sports contracts are roughly efficient but do exhibit bid-ask bounce.

¹³ Appendix A provides a list of these events.

because either: 1) the Saddameter is capturing the odds of war, while the Saddam Security is capturing the odds of a quick victory or 2) Saletan is calibrated differently from the marginal market participant.

We use movements in the Saddam Security to try to interpret movements in financial prices. Our empirical approach is motivated by the following very simple model. Traders trade a financial asset and a prediction market security. The prediction market security pays $p = 1$ if and only if war occurs, and the financial asset is worth η without war and $\eta + \beta$ with. Individual traders have heterogeneous beliefs about η , β , and the probability of war, p . Beliefs about the probability of war (p) are uncorrelated cross-sectionally with beliefs about its severity (β). As such, equilibrium prices in the financial and prediction market will reflect the central tendency of trader's beliefs about $\beta p + \eta$ and p , respectively, and, under most commonly assumed utility functions will closely approximate the mean expectation.¹⁴

If beliefs about the probability of war and the no-war value of the asset (η) evolve over time, and beliefs about the likely severity of war are stable, then changes in the financial asset's price will be given by:

$$\Delta P_t^{fin} = \bar{\beta} \cdot \Delta \bar{p}_t + \Delta \bar{\eta}_t$$

where $\Delta \bar{p}_t$, the change in the prediction market price, is also the change in the average belief about the probability of war. Finally, if innovations in η_t and p_t are orthogonal, then an OLS regression of the financial market on the prediction market will recover an unbiased estimate of $\bar{\beta}$, the mean estimate of the effects of war.

The assumptions we make are those typical of event studies, and several deserve further discussion. Assuming that beliefs about the probability and severity of war are uncorrelated cross-sectionally leads $\bar{\beta}$ to be a simple average belief about severity; modifying this assumption would simply lead traders' beliefs about severity to be weighted by those about probability. Assuming that $\bar{\beta}$ is

¹⁴ For example, prices would exactly equal the mean of beliefs if the relationship between the position an investor takes and her subjective expected returns is linear; this is the case for investors with log utility in a binary prediction security and for investors with CARA utility in securities with normally distributed returns.

constant over time is reasonably defensible given that news during our ex-ante sample period was primarily about whether a U.S.-led coalition would go to war with Iraq as opposed to how severe a war would be. To the extent that news about severity was correlated with news about likelihood, however, our estimator would be biased.¹⁵ The setup above yields a linear functional form, which is probably appropriate given that the right-hand-side variable is a probability.¹⁶ In addition, the assumption that changes in η and p are orthogonal implies that the direction of causality is from war to financial markets, rather than the reverse or both being caused by a third factor. For example, if news that was bearish for stocks and bullish for oil also made war more attractive to U.S. decision makers, this could also produce the correlations we observe. Temporary oil price disruptions seem unlikely to make more attractive a war that would presumably exacerbate them, but a (small) number of observers have argued that the Iraq war was launched to distract voters from the corporate scandals of 2002 (the “wag the dog” story). It seems unlikely that this variation would dominate our high frequency data.

To partially address endogeneity concerns, we can use changes in Saletan’s Saddameter as an instrumental variable for changes in the Saddam Security. Saletan based the Saddameter on concrete news about, for example, troop movements, rather than on financial market movements.¹⁷ So unless administration responses to financial price movements were rapid enough to appear to be contemporaneous in our analysis, “wag the dog” dynamics should not produce a correlation between

¹⁵ The probability of war, as measured by the June Saddam Security, varied from 80 in late September to 37 in late November to 82 by the end of our ex-ante sample. The list of events in Appendix A suggests that most of the variation was related to uncertainty about whether the UN would approve a war and whether a U.S.-led coalition would go to war without UN approval. The main possible exception was news about weapons of mass destruction (WMD), but even this news was primarily about whether WMD would be found by inspectors as opposed to about whether WMD existed that could be used against coalition forces. Even after WMD were not used by Iraq during the war, Tradesports securities on whether WMD would be found by June 2003 traded in the 70s, suggesting that traders were then fairly confident that they existed -- regardless of whether UN inspectors had found them or not. Another pre-war news event with implications for severity was Turkey’s decision to not participate, but this occurred after our ex-ante sample.

¹⁶ We conducted a fairly exhaustive set of tests for non-linearities or asymmetric responses to changes and did not find robust, statistically-significant evidence of them, although this may be due to insufficient statistical power.

¹⁷ In private correspondence (2/11/03), Saletan expanded on the information set underlying the Saddameter: “I read 4 papers a day (NYT, WP, WSJ, LAT), but for the Saddameter, I soon began to rely on the AP and Reuters wires, because I wanted the facts unfiltered. I never looked at op-ed pages. I never looked at stock markets or oil markets. The only stories I gave weight to, other than stuff directly related to Iraq, were stories about North Korea. Also, I did give weight to polls early on, since a serious rise in domestic antiwar sentiment might have derailed Bush’s plans. But that sentiment never reached critical levels.” Beyond this, Saletan also noted that he was not even aware that there was betting on the likelihood of war.

changes in equity prices and in the Saddameter. In addition, instrumental variables helps with the errors-in-variables problems created by bid-ask bounce in the Saddam Security.

A related endogeneity issue would arise if financial markets incorporated news about war more quickly than the prediction market, and prediction market traders looked to financial markets to learn about war news. In the extreme case where Saddam Security traders looked only at financial markets, a representative trader making efficient inferences would update their war probability as follows:

$$E(\Delta p_t | \Delta P_t^{fin}) = \frac{\beta^2 \text{Var}(\Delta p_t)}{\text{Var}(\Delta \eta_t) + \beta^2 \text{Var}(\Delta p_t)} \cdot \beta^{-1} \cdot \Delta P_t^{fin}$$

Regressing ΔP_t^{fin} on a Δp_t that reflected these expectations would yield an estimated coefficient of

$$\hat{\beta} = [1 + \frac{\text{Var}(\Delta \eta_t)}{\beta^2 \text{Var}(\Delta p_t)}] \cdot \bar{\beta}. \text{ Intuitively, the coefficient is biased away from zero because the}$$

prediction market: 1) reacts to non-war related news and 2) underreacts to war-related news.

Instrumenting using a variable like the Saddameter that reflects only war-related news helps with the first issue but not with the second, and does not reduce the bias (recall that Saletan avoided consulting financial market prices in formulating the Saddameter).

Given the one to five day length of our estimation window, however, it seems unlikely that this bias will be large. Some traders may attempt to quickly profit from oil price movements without attempting to disentangle war from non-war news, but this seems unlikely to dominate price movements at our frequency. Furthermore, if we are willing to assume that changes in the Saddameter are correctly calibrated to changes in the war probability, we can test whether the Saddam Security is underreacting to war news as suggested above. Under these assumptions, regressing the prediction market on the true probability of war will yield a coefficient of

$$[1 + \frac{\text{Var}(\Delta \eta_t)}{\beta^2 \text{Var}(\Delta p_t)}]^{-1}, \text{ allowing us to calibrate the possible bias. When we regress the June Saddam on}$$

the Saddameter we get a coefficient of 0.9 (SE = 0.08, $R^2 = .81$), suggesting that any bias from the prediction market participants responding to financial market movements is less than 10 percent.¹⁸

¹⁸ The close correlation between Saletan's Saddameter and the Saddam Security suggests that using the Saddameter as an alternative measure of the probability of war would have yielded similar results, and this is in fact the case. The

3. Oil and Stock Price Effects

War in Iraq was expected to have major implications for oil prices, yet there was considerable disagreement among experts as to their magnitude and even their long-run sign. Two main points of disagreement were how severe the short-run disruptions would be and how rapidly and fully Iraqi's production would be restored after the war.¹⁹

Spot oil prices and the probability of war were strongly positively correlated during our sample period (Figure 2), suggesting the marginal market participant expected a substantial near-term price rise, likely reflecting disruptions to supply. This correlation is even stronger when we control for the other major source of news during our sample period: the labor and political tensions in Venezuela. In our formal analysis below, we use a simple proxy to partial out this potentially confounding effect.²⁰ The close correlation between high frequency movements in oil prices and our conflict indicator provides further cause for confidence that the Saddam Security indeed reflects the likelihood of war. Moreover, the directly interpretable scaling on our independent variable allows us to make a more precise statement about the effects of conflict on oil prices.

The left half of Table 1 reports the results of various regressions of changes in the daily closing prices in the market for West Texas Intermediate oil on changes in the relevant closing price in the market for Saddam Securities. These difference regressions raise a minor technical issue: our various Saddam Securities did not trade on every day during our ex-ante period, so day-to-day changes would risk losing much of our data. Thus we stack first differences from both the March and the June securities, analyzing changes in closing prices that sometimes extend over several days.²¹ We match these changes to the corresponding changes in the closing price of the oil. This

reason for focusing on a market-based assessment such as the Saddam Security in our analysis is that expert opinion in future settings may not be so well calibrated.

¹⁹ See Leigh, Wolfers, and Zitzewitz (2003) for a full review of pre-war opinion.

²⁰ The control variable used is the redemption yield on corporate bonds issued by PDVSA, the main Venezuelan state-run oil company. Specifically, we analyze redemption yields on 6.45% 1998 PDVSA bonds expiring on 2/14/2004 (Code 237019(RY)). We choose this variable precisely because it is a forward-looking financial instrument, likely correlated with expectations of future disruption in Venezuelan oil supply.

²¹ For example, one week might see the March security traded only on a Tuesday and Thursday, and the June security traded only on a Monday and a Wednesday. In this instance, we would have two observations for that week: the

gives us a total sample of 147 observed shifts in the probability of war, which we use to predict shifts in oil prices.²²

The first specification we run, daily first differences, yields our smallest estimate. If the Saddam Security prices exhibit slow incorporation of information or bid-ask bounce (they do), then this coefficient may be attenuated. Alternate specifications that account for these issues by including future Saddam Security changes, by instrumenting for the Saddam Security using the Saddameter, or by using longer differences, all yield roughly similar point estimates: 10 percentage point increase in the probability of war is accompanied by a \$1 rise in the spot oil price.²³

Arguably more important for the global macroeconomy than war's effect on spot prices is the effect on oil prices in the medium to long-term, and so we repeat the above analysis on the prices of oil futures. Futures contracts exist for delivery at each of the next 30 months, and long-dated futures exist for December out to seven years in the future. Thus, for each functional form we estimate 33 separate regressions – 28 for each of the monthly futures markets for delivery each month from March 2003-June 2005, plus five for the long-dated markets for December of 2005-2009. Throughout, we analyze closing prices in the oil futures market, comparing them with the price of the last trade in the relevant Saddam Security prior to the 2:30pm close of the NYMEX market.²⁴

Figure 3 shows a compact representation of these results, graphing the expected impact of war on oil prices.²⁵ The chart shows coefficients from our preferred 5-day difference specification, although other specifications yield similar results. We project our results out of sample and compare oil prices under the counterfactuals of a 0 and a 100 percent probability of war. Thus, the first point on

difference in prices of the March security from Tuesday to Thursday (and the corresponding Tuesday to Thursday oil price difference), and the change in price of the June security from Monday to Wednesday (and the corresponding Monday to Wednesday change in oil price). This leads us to use Newey-West standard errors, allowing for autocorrelation up to twice the difference period. We maintain this convention throughout this paper.

²² If we instead use only the June Saddam Securities, we get less precise, but not statistically significantly difference, estimates.

²³ Since the less liquid security is on the right-hand side of the regression here, including leads is analogous to including lags of the liquid security on the left as in Scholes and Williams (1977).

²⁴ While most futures were traded on most dates, there are some missing observations; these were imputed by applying the daily percentage change observed in the nearest shorter-term contract (and when this did not exist, the nearest longer-term contract is used.)

²⁵ Note that the confidence intervals in Figure 3 represent the statistical uncertainty of our estimate of the market's mean expectation about war's effect, not participants' uncertainty about the actual ex-post effects.

the chart shows the spot price rising by \$11. The second point refers to the March 2003 futures price, which increases by around \$9.50. The incremental effect of war then declines rather rapidly, falling to \$5 by June 2003, and around \$2 by the end of 2003. The actual evolution of oil prices since the war has largely confirmed this prediction, as oil prices did indeed decline rapidly through 2003. The effects of war virtually disappear in the medium run, and 2004 and 2005 oil futures were uncorrelated with the probability of war. Long-run oil futures were negatively correlated with the probability of war, suggesting a market perception that ousting Saddam would lead prices to eventually decline by around \$1.50 per barrel.

Thus although war news accounted for much of the movement in oil prices during our sample period, the effects of war on oil prices were expected to be smaller in magnitude and shorter-lived than past oil shocks. Although oil price movements can have non-linear and otherwise difficult to anticipate effects, running an oil shock of this magnitude through a macroeconomic model yields only small negative effects on GDP (Hamilton, 2003).

In contrast to long-run oil futures, global equity markets reacted very negatively to increases in the probability of war, suggesting that they were worried about more than an oil shock (Figure 4). The right half of Table 1 presents regressions that analyze this relationship more formally. As with oil, we match the most recent transaction in each Saddam Security to the closing value of the S&P 500.²⁶ As with oil, we find that future changes in the Saddam Security have an economically (if not always statistically) significant relationship with current changes in the S&P, and so we include several leads in the difference specification. These regressions suggest a 10 percent increase in the probability of war is associated with a 1.1 percent decline in the S&P 500. Long difference regressions imply a larger effect. Instrumenting for changes in the Saddam Security with changes in the Saddameter yields even larger, albeit very imprecisely estimated coefficients.

These estimates are economically important. The 5-day difference specification represents our central estimate, and projecting these estimates out of sample suggests that a change from a zero to a 100 percent probability of war reduces the S&P 500 by around 15 percent. By comparison, the S&P

²⁶ In addition, we are also able to construct a tick-by-tick sample, in which we match each trade in the Saddam securities to the next S&P future trade recorded on the CME. Point estimates, reported in Leigh, Wolfers and Zitzewitz, are similar to those in the daily analysis, albeit more precise.

500 fell by 7.6, 6.5, and 5.5 percent in the first two trading days after the Pearl Harbor bombing, the outbreak of the Korean War, and September 11th, respectively.

Such a large estimated negative effect seems difficult to reconcile with the modest medium-run effect on oil prices. To attempt to understand its source, we decompose it in different ways. A decline in equity values could be due to either decreases in expected future earnings or an increase in the required discount rate, which is the sum of the risk-free rate and the equity risk premium. The equity risk premium is notoriously difficult to measure, but we can use the inflation-indexed, 10-year Treasury as a proxy for the risk-free rate. Replicating the regressions Table 1 with this yield as the dependent variable suggests that war reduces the risk-free rate by about 40 basis points, which by itself should *increase* equity values by about 15 percent.²⁷ So the observed war discount must be coming from a substantial change in earnings expectations or the equity premium.²⁸

This estimated war discount in equity prices essentially reflects investor's beliefs about the average impact of a war. There were many possible war scenarios, however, and as a second decomposition, we can analyze option prices to infer investors' beliefs about the distribution of possible effects of the war. When we do this, we find that option prices suggest that the expected distribution of war outcomes was highly negatively skewed, presumably reflecting a small but substantial probability of a terrible outcome (e.g., chemical or biological attacks on the U.S. or its regional allies), which helps explain why the estimated *average* effect is so large.

A direct way of observing investors' beliefs about the probability of a terrible outcome is to examine how the value of deep out-of-the-money S&P puts varied with the probability of war. An S&P 500 put option is an option to sell the S&P 500 index at a pre-specified strike price. Buying a deep out-of-the-money put (for example, a put with a strike price of 600 when the S&P is trading at 900) would allow an investor to insure against extreme losses. Such an investor would incur losses if the S&P dropped to 600, but would not incur any beyond that point. The price of such a put option is thus a rough indicator of the likelihood that the S&P will drop to that level or below.

²⁷ Applying the earnings version of the Gordon (1962) growth model, a 40 basis point reduction in the risk-free rate should, all else equal, lower the earnings-price ratio by 40 basis points. Given the current S&P earnings-price ratio of about 2.5 percent, this would imply a 15 percent increase in equity valuations.

²⁸ A possible alternative is that the inflation-indexed Treasury bond is not only a poor proxy for the risk-free rate, but is also negatively correlated with it.

The prices of deep-out-of-the-money puts rose dramatically with the probability of war (Figure 5). This happened for three reasons. First, the decrease in the level of the S&P 500 brought the put options closer to being in-the-money. Second, war risk increased the expected future volatility of stock prices, and thus the value of all out-of-the-money options. Third, war risk raised the actual value of deep-out-of-the-money puts more than it raised their Black-Scholes (1973) value (also in Figure 5). This premium over Black-Scholes can be thought of as a measure of the extra weight put on extremely negative outcomes.²⁹ Black-Scholes assumes log-normally distributed returns, and so the fact that this premium increased with war risk suggests that the expected effects of war may have been non-normally distributed.

To investigate this more formally, we use the full range of S&P option prices to estimate the state price density. A state price is the price of an imaginary security that pays \$1 if a certain event happens. In this case, the state price for the state $S\&P = 600$ is the price of a security that pays \$1 if the S&P is equal to 600 when the option expires. If investors are risk neutral, then the distribution of state prices can be interpreted as investors' expectation about the probability distribution of future S&P price levels. (We relax this assumption below.)

We use a non-parametric approach to estimate the state price density for options 120 days from expiry on each day of our sample, and then examine how this density is affected by changes in the probability of war. The details of the method, which is a variation of the smoothed-implied volatility method described in Aït-Sahalia and Lo (1998), are given in Appendix B. But the intuition for how we estimate state prices from option prices can be gained from thinking about butterfly spreads. A butterfly involves buying a call with strike price $S - e$, selling two calls at S , and buying a call at $S + e$. This position pays nothing if the S&P at expiry is outside of the interval $(S - e, S + e)$; the profits inside this interval are given by a triangle with its peak at S . In the limit as e approaches zero, this position converges to a security that pays if and only if the S&P is equal to S .

²⁹ Especially since the 1987 crash, deep out-of-the-money puts have traded at a premium to their Black Scholes values, and the negative skewness implied by option prices has been greater than that in historical returns. There is not yet a consensus about whether this reflects a mispricing (e.g., Ederington and Guan, 2002), extremely high risk aversion at low wealth levels (Aït-Sahalia and Lo, 2000), or a "peso problem" (e.g., Aït-Sahalia, Wang, and Yared, 2001). Our approach is to take this premium as a starting point and draw inferences from its comovement with war risk.

Adopting the approach outlined in Section 2, we regress state prices on war risk. Our identifying assumptions are similar, except that the orthogonality conditions regarding the equity price must be strengthened so as to apply to the entire state price distribution. These regressions reveal the joint probability distribution $f(s,w)$ where s is the future level of the S&P 500 and w is an indicator for whether there is a war. Applying Bayes' rule, the probability of a state is $f(s) = f(s,0) + [f(s,1) - f(s,0)] \cdot \text{Prob}(w=1)$. The state price $p(s) = f(s) \cdot \lambda u'(s)$, where $\lambda u'(s)$ is the relative marginal utility of wealth in different states. Our regression can be written:

$$\begin{aligned} \Delta p_t(s) &= u'(s) \cdot [f(s,1) - f(s,0)] \cdot \Delta p_t^{\text{saddam}} + \varepsilon_t \\ \Delta p_t(s) &= [p(s,1) - p(s,0)] \cdot \Delta p_t^{\text{saddam}} + \varepsilon_t \end{aligned}$$

where we run this regression for each state price. These coefficients allow us to compute counterfactual state price distributions with high and low war risk. The state price density for the last day of our sample (February 6, when the price of the June Saddam Security was 75) was taken as a benchmark, and our estimated coefficients were applied to estimate the state price density at war probabilities of 0, 40, 80, and 100. The range of June Saddam securities in our sample is 37-80, so the predictions for war probabilities of 0 and 100 are out-of-sample predictions that should be interpreted with caution. Figure 6 shows that the higher-war-risk distributions clearly have lower means, higher variance, and more negative skew, as evidenced by the fatter left-hand tails.

The high and low-war-risk distributions can be compared to estimate the distribution of likely effects of war, under the assumption that the risk of war adds a mean-altering spread to the probability distribution under low war risk. More formally, even without war, the value of the stock market in 120 days is uncertain (reflecting non-war related sources of risk), and distributed according to the probability distribution function $f(s',0)$. The effects of war, W , are also uncertain, and the likelihood of the different scenarios playing out are given by the probability distribution function $g(W)$. Thus stock-market outcome conditional on war, is simply the sum of s' and W and its distribution is a convolution of these two distributions. That is, by Bayes Rule:

$$f(s,1) = \int f(s',0) \cdot g(s - s') \cdot ds'$$

Intuitively this says that the probability that we observe, say, a moderately bad value of the S&P under high war risk reflects the probabilities of the various permutations of the stockmarket being

exogenously weak with a good war outcome, through to moderate outcomes on both, and even a good stockmarket draw coupled with a bad war outcome. Since we construct the related state price distribution, $p(s, l)$ and $p(s, \theta)$ (in Figure 6), we can apply the Consumption CAPM to convert these state prices into probabilities, and use this set of restrictions to solve for $g(W)$, the probability distribution of likely effects of war on the stock market.

As a simple benchmark, we assume a representative consumer with different levels of constant relative risk aversion. We also assume that this representative consumer holds all her wealth in the S&P 500, an assumption that leads us to overstate the effects of risk aversion. Figure 7 graphs our numerical solutions for $g(W)$ for the three parameterizations of risk aversion.³⁰ As one extreme, we assume zero risk aversion, which would imply that state prices can be interpreted as probabilities. We also estimate $g(W)$ for CRRAs of one and two, the estimates given by Arrow (1971) and Friend and Blume (1975), respectively.³¹ Regardless of the risk aversion assumed, the bulk of the probability mass is concentrated in the 0 to -20 range, with a long left tail; the risk aversion assumption affects primarily whether the estimated probability of an outcome worse than a 20 percent drop is ten or twenty percent.

Assuming a CRRA of one, we estimate that there is a 70 percent probability that changing from a 0 to a 100 percent chance of war reduces the S&P 500 by 0 to 15 percent. There is a small probability of a large negative impact: a 20 percent probability of a 15 to 30 percent decline and a 10 percent probability an even larger decline. The mean expected near-term effect of war of -15 percent was thus less than both the median (-12 percent) and the mode (-10 percent). If the near-term effect of war took on its modal value, then one would expect the S&P 500 to rally roughly 5 percent as the uncertainty created by the war, particularly the small probability of an extremely bad outcome, declined. This is in fact what we observed during the four-week war, as we discuss in the next section.

³⁰ We set $g(W)$ point-by-point to minimize the mean-squared error of our resulting estimate of $f(s, l)$.

³¹ A CRRA of 2 would imply that the state price-to-probability ratio at S&P = 500 is four times that at S&P = 1,000. Some studies have calculated that much higher levels of risk aversion would be necessary to rationalize the equity premium (e.g., Mehra and Prescott, 1985), the purchase of very disadvantageous forms of insurance (e.g., Ciochetti and Dubin, 1994), or the high average prices of deep-out-of-the-money puts themselves (Ait-Sahalia and Lo, 2000; Jackwerth, 2000). In general these estimates were intended more as illustrations of puzzles than as serious estimates of risk aversion, so we focus on calibrating our model using the older, more direct CRRA estimates from experiments and introspection.

4. Ex-post analysis

Do the results of our ex-ante analysis appear reasonable in retrospect? We can answer this question in two different ways. First, if the correlations between financial prices and the probability of war were meaningful and not statistical artifacts, they should have persisted into the period between the end of our initial sample and when the decision to go to war was made. Second, given that most commentators agree that the immediate effects of the war were less dire than many forecast, we should observe at least a partial reversal of the war discount (or premium) that was built into prices. The sectors and countries most depressed by war risk should benefit the most from its partial resolution yielding further cross-sectional implications. Likewise, since war news explained a significant share of price variation in our ex-ante period, we would expect it to explain a significant share of variation leading up to and during the war.

Table 2 presents regressions of stock and spot oil prices on the June Saddam Security for three time periods: our ex-ante period, the period between the end of our original sample (February 6, 2003) and the last trading day before the issuance of the 48-hour ultimatum to Saddam (March 14), and the period between the start of the war and the expiry of the Saddam Security (April 11, two days after the statue of Saddam was toppled). These regressions suggest that the correlation we observed in our sample period persisted up until the war became a certainty. On March 19, the day coalition forces crossed the border, war was certain and the June Saddam closed at 95, reflecting the high probability of successfully removing Saddam in just over three months. From that day forward, all of the news affecting the Saddam securities was about the progress of the war. Before the war, an increase in the Saddam Security price reflected a higher probability of war and thus an increase in war-related uncertainty. Once war was a certainty, Saddam Security price increases implied an earlier expected end to hostilities and thus a reduction in war-related uncertainty. The correlations with equity and oil prices reversed, as one would have expected. This reversal does not imply that once the war started, markets shifted to believing that war is positive for equities; rather it suggests that given that war was a certainty, a short war was believed to be good for equities. Thus movements in the Saddam Security switched from reflecting news about the likelihood of war to

proxying for the duration of the war, and hence the coefficients before and during the war cannot be meaningfully compared.³²

Figure 8 graphs spot oil prices, the S&P 500, and the prices of the Saddam securities for March, April, May, and June (April and May contracts were introduced by Tradesports on February 5). Price movements suggest three separate periods. First, there was a period of rapid initial progress (March 14 to March 21) in which Iraqi oil fields were captured largely intact and the beginning of the war was not accompanied by terrorism against neighboring oil fields or U.S. cities, as some had feared. The April to June Saddam Security prices increased into the high 90s, equity prices rallied by about 5 percent, and spot oil prices fell about \$8, reversing much of our estimated war premium. Second, there was a period (March 22 to March 31) in which progress slowed, the welcome extended coalition forces by the population was less enthusiastic than forecast, and experts began envisioning a difficult urban battle for Baghdad with possible use of weapons of mass destruction (WMD). Most of the initial rally in equities and some of the decline in oil prices was reversed. Finally, from March 31 to April 11 the news turned positive again, as Iraqi resistance collapsed more quickly than expected and Baghdad fell without terrorism or WMD use. This narrative reinforces the impression that oil and equity markets closely followed news about the likely severity of war.

If we date the end of the war as of the toppling of the Saddam statue (April 9th), then the S&P 500 rallied by 3.8 percent during the war and spot oil prices fell by \$6.48. While non-war-related news surely also moved these markets during this four-week period, it is tempting to ignore this possibility and interpret these price movements as implying that about a quarter of the 15 percent war discount in equity prices and two-thirds of the \$10 war premium in oil prices was reversed.

A final test of our estimated ex-ante war discounts hinges on the expected cross-sectional impacts of the war. In our ex ante analysis we estimated war discounts at the sector and country level. These estimates vary with sector and country characteristics in a sensible way, and can be compared with the subsequent rally when the war turned out to be more benign than expected.

³² For example, in mid-2003, Yakov Amihud and Avi Wohl circulated a working paper that claimed to overturn the result of our ex-ante analysis, since it found a positive coefficient from regressing changes in the S&P on the Saddam Security over the entire February 7 to April 11, 2003 period. We pointed out to them that the interpretation of the Saddam Security changed when the war started, and subsequent drafts of their paper reflected this. The results in the final version of their paper (Amihud and Wohl, 2004) are consistent with what we report here.

Table 3 reports these ex-ante war effect estimates and returns during the war for 11 top-level S&P sector indices and selected sub-indices. War effects are estimated using 5-day difference regressions for both raw returns and, to adjust for sectors' normally co-movement with the market, take excess returns (or alphas) from a single-factor model. Results from both return measures, especially the alphas, accord reasonably well with intuition. War was estimated to be bad for airlines, consumer discretionary, and investment-sensitive sectors such as information technology, telecom, and finance. It was likewise viewed as being positive for gold stocks and, controlling for its negative effect on the overall market, for energy and defense stocks. As the war turned out to be more benign than feared, airline and discretionary spending stocks rallied, while energy stocks fell. In general, sectors with a bigger war discount had higher returns during the war, although there are some outliers, such as gold and IT.³³

We repeat this analysis for countries by analyzing the 43 national stock markets for which Morgan Stanley Capital International (MSCI) total return indices facilitate meaningful cross-country comparisons (for consistency with the analysis above, we use the S&P 500 for the U.S.). As above, we estimate expected effects of war using 5-day difference regressions of a country's MSCI index on the June Saddam Security. In all cases, we measure returns in U.S. dollars and use daily data, matching MSCI indices with the most recent Saddam Security trade as of the closing time of that particular market.

Figure 9 reports ex-ante war effect estimates for these countries. In 32 of the 43 countries, a higher probability of war is associated with a fall in the stock market, and in fourteen of these, the fall is statistically significant at the 5 percent level. In Austria and Indonesia, a higher probability of war has a positive effect that is statistically distinguishable from zero.

What explains the differing national effects? We constructed several variables in order to try to describe the observed pattern of cross-country effects. One would expect countries that are more exposed to global stock market movements (e.g., due to a specialization in investment goods) to be

³³ A cross-sectional regression of alpha during the war on the ex-ante estimated war effect (on alpha) yields a coefficient of -0.08 (standard error 0.035).

more negatively affected by the war. Finland, Sweden, and Germany fall into this category. Countries that are more dependent on oil imports should also be disproportionately affected. Neighboring countries (Israel, Jordan, Pakistan, Turkey) were believed to be more vulnerable to attack or unrest. In addition, one might expect countries to be differently affected if they participated in the war (Australia, the UK, and the U.S.) or supported the US effort (the eight European countries signed a pro-US letter to the *Wall Street Journal*: the Czech Republic, Denmark, Hungary, Italy, Poland, Portugal, Spain, and the UK).

The results in Table 4 confirm that financial markets have priced in larger adverse effects of war in neighboring, high-beta, and oil-importing nations. Controlling for these factors, neither the variables describing a country's commitment to the removal of Saddam, nor the dummy variable for European markets, are statistically significant.³⁴ During the war, these patterns were partly reversed, with vulnerable and oil-dependent countries' markets clawing back some of their pre-war discounts.³⁵

5. Discussion

By studying an unusual financial instrument, the Saddam Security, we were able to track shifts in the probability of war between September 2002 and February 2003, measure the correlation between war probabilities and financial market movements, and infer ex-ante market expectations of some of the consequences of war from these correlations. This analysis suggested that the war would have sharp but transitory effects on world oil markets -- a prediction that has been largely borne out. Global equity prices reflected larger discounts, and option prices suggested that this was partly due to a negatively skewed distribution of risks. Ultimately, the immediate effects of the war were a somewhat benign draw from this distribution, and the war's resolution was accompanied by a rally, especially in the sectors and countries with the largest pre-war discounts.

³⁴ Single-stage regressions, in which a panel of country-level stock returns were regressed on an interaction of the changes in war probabilities and the country-specific explanatory variables, yielded qualitatively similar results.

³⁵ The regressions include a dummy variable for the countries most affected by SARS during that time period (China, Hong Kong, Singapore, and Taiwan). Results are not materially affected if these countries are simply omitted or if Canada is included as being SARS-affected.

This analysis could conceivably be repeated to analyze the effects of other policy decisions, but our analysis was eased by some special features of our setting that should be noted. Any analysis based on regression needs: 1) sufficient variation in the independent variable, 2) for the relationship being examined to produce a high enough signal-to-noise ratio to be detectable, and 3) for the direction of causality to be reasonably clear. Our analysis benefited from the vacillation produced by the coalition's attempt to obtain UN backing: the price of the June Saddam varied from 82 in September 2002 to a low of 37 in November before returning to 75 by the end of our ex-ante sample, and there were substantial daily price changes produced mainly by the ebb and flow of the UN process. The analysis also benefited from the fact that war was viewed as a first-order determinant of oil and equity prices, from the fact that the direction of causality was arguably fairly clear, and from the fact that stock market-based assessments of the economic impact of war did not appear to be a determining factor in the coalition's decision making.

Clever design of the prediction securities can to some extent address a lack of variation in event probability or a low signal-to-noise ratio. For example, instead of examining the correlation of changes in financial market prices and the probability of war, one could run markets in securities that pay P^{fin} if War occurs, with all transactions being cancelled if it does not. The price of such a security and its complement would give an indication of $E(P^{fin}|War)$ and $E(P^{fin}|NoWar)$ and thus the correlation between war and the expected value of the financial asset.³⁶

These markets could yield estimated effects of war even when the probability of war did not vary and, if the markets are liquid enough, could provide meaningful estimates even when the expected effect of war was small relative to other factors affecting the market. These contingent securities do not help, however, when the direction of causality is unclear; in fact, determining causality can be even more difficult than with an analysis of historical price change correlations. For example, in our historical analysis, we relied on an analysis of the actual news events that affected war probabilities and argued that reverse and third-factor causation (e.g., bad economic news making war more likely) did not play a role in the coalition's decision making. If we had instead analyzed securities that tracked $E(S\&P|War)$ and $E(S\&P|NoWar)$, we would have had to assert that market participants

³⁶ Also see Hanson (1999), Berg and Reitz (2003), and Wolfers and Zitzewitz (2004b) for further discussion of these markets.

placed a zero probability on these factors being important. This is a much stronger assumption, although it need only be true on the particular dates whose prices one chooses to analyze, rather than throughout the sample period.³⁷

Another issue that we argue did not significantly complicate our analysis, but may affect future analyses, is the possibility of the results of an ex-ante policy analysis being contaminated by the fact that they may be used in decision making. This could happen in at least two ways. First, those with an interest in whether the policy is implemented could manipulate the prices of the S&P|War and S&P|NoWar securities. This sort of manipulation may be eased by the lower trading volumes in the contingent markets run thus far.

Assuming no manipulation, the securities designed to capture $E(S\&P|War)$ and $E(S\&P|NoWar)$ may do so, but the interpretation of the market measure of war's expected effects, $E(S\&P|War) - E(S\&P|NoWar)$, is not necessarily straightforward. For example, suppose that the government combines a market-based assessment of the net economic benefits of war with its own private information about economic and non-economic net benefits in making its decision. Then if the market predicts adverse effects, one might expect the government to be unlikely to undertake a war unless its private information suggested that the market was wrong. Market participants should recognize this in constructing their estimate, and thus the market-based estimate of war's "effect" should be smaller than if the government's decision to go to war was not expected to depend on the likely economic consequences. Different informational assumptions may yield different reaction effects and thus different biases. Nevertheless, as we have shown by example, in certain settings these securities can provide a useful mechanism for extracting the rich information already priced into financial assets. Moreover, these securities can be used to price not only the expected effects of a policy but also the probability distribution of possible outcomes.

³⁷ A more important shortcoming of these compound securities is that they may be too difficult for many prediction market participants to understand. Contingent securities on both Iowa and Tradesports have attracted volumes that compare favorably with previous Iowa vote share markets, albeit well below those attracted by the Saddam Security or Tradesports' Bush reelection contract.

References

Aït-Sahalia, Yacine and Andrew Lo (1998), "Nonparametric estimation of state-price densities implicit in financial asset prices" *Journal of Finance*, 53, 499-548.

Aït-Sahalia, Yacine and Andrew Lo (2000), "Non-parametric risk management and implied risk aversion" *Journal of Econometrics*, 94, 9-51.

Aït-Sahalia, Yacine, Yubo Wang, and Francis Yared (2001), "Do Options Markets Correctly Price the Probabilities of Movement of the Underlying Asset?" *Journal of Econometrics*, 102, 67-110.

Amihud, Yakov and Avi Wohl (2004), "Political News and Stock Prices: The Case of Saddam Hussein Contracts," *Journal of Banking and Finance*, 28(5), 1185-1200.

Arrow, Kenneth (1971), *Essays in the Theory of Risk Bearing*. North Holland, Amsterdam.

Berg, J.E, R. Forsythe, F.D. Nelson and T.A. Rietz (2001), "Results from a Dozen Years of Election Futures Markets Research," in Charlie Plott and Vernon Smith (Eds) *Handbook of Experimental Economics Results*

Berg, Joyce, Forrest Nelson and Thomas Rietz (2001), "Accuracy and Forecast Standard Error of Prediction Markets", *mimeo* University of Iowa.

Berg, Joyce and Thomas Rietz (2003), "Prediction Markets as Decision Support Systems", *Information Systems Frontiers*, 5(1).

Black, Fisher (1976), "Studies of Stock Price Volatility Changes," in *Proceedings of the 1976 Meetings of the Business and Economic Statistics Section, American Statistical Association*, 171-181.

Black, Fisher and Myron Scholes (1973), "The Pricing of Options and Corporate Liabilities" *Journal of Political Economy*, 81, 637-659.

Bliss, Robert R. and Nikolaos Panigirtzoglou (2002), "Testing the Stability of Implied Probability Density Functions" *Journal of Banking and Finance*, 26, 381-422.

Breeden D. and R. Litzenberger (1978), "Prices of State-Contingent Claims Implicit in Option Prices" *Journal of Business*, 51, 621-651.

Ciochetti, Charles J. and Jeffrey A. Dubin (1994), "A Microeconomic Analysis of Risk Aversion and the Decision to Self-Insure," *Journal of Political Economy*, 102, 169-186.

Cutler, David, James M. Poterba and Lawrence H. Summers (1989), "What Moves Stock Prices?" *Journal of Portfolio Management*. 15(3): 4-12

DARPA (2003), "Report to Congress regarding the Terrorism Information Awareness Program", May 20, 2003.

Davis, Steven J., Kevin M. Murphy, and Robert H. Topel (2003) "War in Iraq versus Containment: Weighing the Costs," *mimeo* University of Chicago.

Ederington, Louis and Wei Guan (2002) "Why Are Those Options Smiling?" *Journal of Derivatives*, 10, 9-34.

Elmendorf, Douglas W, Mary L Hirschfeld, and David N Weil, "The Effect of News on Bond Prices: Evidence from the United Kingdom, 1900-1920," *Review of Economics and Statistics*, 78(2), 341-44.

Frey, Bruno S. and Kucher, Marcel (2000), "History as Reflected in Capital Markets: The Case of World War II". *Journal of Economic History*. 60(2), 468-496

Friend, I. And M. E. Blume (1975) "The Demand for Risky Assets" *American Economic Review*, 65, 900-922.

Gordon, M (1962) *The Investment, Financing, and Valuation of the Corporation*. Irwin, Homewood, IL

Hamilton, James (2003) "What Is an Oil Shock?" *Journal of Econometrics*, 113, 363-398.

Hanson, Robin (1999) "Decision Markets," *IEEE Intelligent Systems*, 14(3), 16-19.

Jackwerth, Jens Carsten (2000), "Recovering Risk Aversion from Option Prices and Realized Returns" *Review of Financial Studies*, 13, 433-467.

Knight, Brian (2003), "Are Policy Platforms Capitalized Into Equity Prices: Evidence from Equity Markets During the Bush/Gore 2000 Presidential Election," *mimeo* Brown University.

Leigh, Andrew, Justin Wolfers, and Eric Zitzewitz (2003), "What do Financial Markets Think of War in Iraq?" NBER Working Paper No. 9587.

Mehra, R. and Edward Prescott (1985), "The Equity Premium: A Puzzle" *Journal of Monetary Economics*, 15, 145-161.

Nordhaus, William D (2002), "The Economic Consequences of a War with Iraq" In Carl Kaysen, Steven E. Miller, Martin B. Malin, William D. Nordhaus and John D. Steinbruner, *War with Iraq: Costs, Consequences, and Alternatives*. Cambridge, MA: American Academy of Arts and Sciences. pp. 51-86. Available at http://www.econ.yale.edu/~nordhaus/homepage/AAAS_War_Iraq_2.pdf

Rigobon, Roberto and Brian Sack (2003) "The Effects of War Risk on U.S. Financial Markets," NBER Working Paper No. 9609.

- Rubenstein, Mark (1994), "Implied Binomial Trees," *Journal of Finance*, 49, 771-818.
- Scholes, Myron. and J.T. Williams, 1977, Estimating betas from nonsynchronous data, *Journal of Financial Economics* 5, 309-327.
- Slemrod, Joel and Timoth Greimel (1999), "Did Steve Forbes Scare the Municipal Bond Market?", *Journal of Public Economics*, 74(1), 81-96.
- Strumpf, Koleman (2004), "Manipulating the Iowa Political Stock Market," *mimeo* University of North Carolina.
- Tetlock, Paul C., 2004, "How Efficient Are Information Markets? Evidence From An Online Exchange," Harvard mimeo.
- Waldenström, Daniel and Frey, Bruno S (2002), "How Government Bond Prices Reflect Wartime Events: The Case of the Stockholm Market". Institute for Empirical Research in Economics. University of Zurich. Working Paper No. 102
- Wolfers, Justin and Eric Zitzewitz (2004a), "Prediction Markets," *Journal of Economic Perspectives*, 18(2), 107-126.
- Wolfers, Justin and Eric Zitzewitz (2004b), "Experimental Political Betting Markets and the 2004 Election," *The Economist's Voice*, 1(2), October 12.
- Yuan, Kathy (2002), "The Liquidity Service of Sovereign Bonds," Michigan mimeo.

Appendix A: Major News Events

The following reflects a list of major news events from our sample period gathered from the New York Times by research assistants, who decided which events merited inclusion. The corresponding dates are shaded in grey in Figure 1.

| <i>Date</i> | <i>Event</i> |
|-------------|--|
| 9/22/2002 | Having voiced strong anti-war views during the campaign, Gerhard Schroeder wins German election. |
| 9/27/2002 | The U.S. presents a UN resolution giving Iraq 30 days to comply with inspections before use of force |
| 9/30/2002 | Iraqi officials agree to give UN inspectors unfettered access |
| 10/11/2002 | U.S. Senate backs use of force |
| 10/26/2002 | Major anti-war protests |
| 11/5/2002 | Republicans hold onto House and regain control of Senate |
| 11/8/2002 | UN Security Council votes 15-0 that Iraq must disarm |
| 11/12/2002 | Reports that Iraq bought nerve gas antidote |
| 11/13/2002 | Iraq agrees to admit weapons inspectors |
| 12/8/2002 | Iraq's report to UN claims it has no weapons of mass destruction |
| 12/13/2002 | US highlights major omissions from Iraq's WMD report |
| 12/19/2002 | Powell declares Iraq in "material breach" of UN Security Council resolution |
| 1/13/2003 | Turkey lets the US survey bases for possible troop deployment |
| 1/16/2003 | UN inspectors discover warheads in Iraq capable of carrying chemical weapons |
| 1/18/2003 | Major anti-war protests |
| 1/21/2003 | Britain mobilizes 26,000 troops |
| 1/22/2003 | Rumors circulated that some Arab leaders are angling to have Saddam exiled |
| 1/27/2003 | UN inspectors issue report giving Iraq a mixed report on compliance |
| 1/28/2003 | Bush makes case for war in State of the Union address |
| 2/5/2003 | Powell addresses the UN |

Source: New York Times, September 20, 2002-February 7, 2003.

Appendix B. Estimating state price densities from option prices

To estimate the state price density from option prices, we follow the method described in Ait-Sahalia and Lo (1998), hereafter AL, with certain modifications.³⁸ AL estimate state prices by first estimating the call-option pricing function $H(S, X, \tau, r, \delta)$, where S and X are the spot and strike prices, τ is the time to expiry, and r and δ are the interest rate and dividend yield. Rather than estimating H fully non-parametrically, which would be very data-intensive since it has five dimensions, AL estimate Black-Scholes implied volatilities non-parametrically as a function of S/X and τ , but rely on the Black-Scholes formula for equating current and future payoffs using r and δ . They estimate $\sigma(S/X, \tau)$ in the formula:

$$H(S, X, \tau, r, \delta) = H_{BS}[S, X, \tau, r, \delta; \sigma(S/X, \tau)]$$

where $H_{BS}(S, X, \tau, r, \delta, \sigma)$ is the Black-Scholes formula.

We modify this approach in several minor ways. First, our option prices are for options on futures traded on the Chicago Mercantile Exchange (CME), so we use the formula in Black (1976) for pricing options on futures, writing F for the future price.³⁹ Second, unlike AL, who are interested in estimating only the option pricing function, we are interesting in using our option pricing function to estimate how state price densities are changing from day-to-day. We therefore need to make assumptions to reduce the data-intensity of the method while preserving its flexibility with respect to the non-normality of future returns.

We do this by non-parametrically estimating the function $\sigma(z, \tau)$ for each day, where $z = [\ln(X) - \ln(F)] / [\sigma_{ATM} * \sqrt{\tau}]$ and σ_{ATM} is the average implied daily volatility for at-the-money options. The parameter z can be viewed as a z-score for that option. When we express implied volatilities as a function of z , we find that the “volatility smile” (the shape of $\sigma(z)$ for a given τ), does not change significantly with τ for τ greater than 15 days. This is convenient, since it allows us to estimate $\sigma(z, \tau)$ with limited worries about its sensitivity to our estimation method. For simplicity, we follow AL and use kernel smoothing to estimate $\sigma(z)$ for each day and τ and then use linear interpolation to generate $\sigma(z)$ for τ 's that we do not observe.

Having estimated $\sigma(z, \tau)$, we can then use this to calculate $H(F, X, \tau, r, \delta) = H_B[F, X, \tau, r, \delta; \sigma(z, \tau)]$, where H_B is Black's formula for options on futures. We then use this function to estimate state prices for each strike price. State prices can be derived from H by differentiating with respect to X (Breedon and Litzenberger, 1978). The value of a call option on a future and its derivatives can be written:

³⁸ Bliss and Panigirtzoglou (2002) compare the smoothed volatility method exemplified by Ait-Sahalia and Lo (1998) with other methods for estimating option pricing functions, finding that it more robust to perturbations in option prices.

³⁹ Our data from the CME are settlement prices for S&P 500 futures and options on futures. The CME calculates settlement prices using recent trade and quotes data and then performs a fair value adjustment for market movements since the option last traded. Each option price is matched a futures settlement price; if a future is not available for a given expiry month, we use the dividend yield and risk-free rate to estimate one from the future with the nearest expiry. Following the past literature, we use only the prices of out-of-the-money options, since these are less sensitive to any measurement error in the futures price. We convert the prices of put options into implied call option prices by applying put-call parity.

$$H(F, X, \tau, r) = e^{-r\tau} \int_0^{\infty} \max(P - X, 0) \cdot p(P, F, \tau) \cdot dP$$

$$\frac{\partial H}{\partial X} = -e^{-r\tau} \int_X^{\infty} p(P, F, \tau) \cdot dP$$

$$\frac{\partial^2 H}{\partial X^2} = e^{-r\tau} \cdot p(P, F, \tau)$$

The second derivative gives the state price function $p(P, F, \tau)$: the price of a security worth a dollar if the price at expiry equals P . The first derivative gives the delta of the option or, alternatively, the price of a security worth a dollar if the price at expiry is greater than X . If investors were risk-neutral, we could interpret the state price density as the p.d.f. of future returns and the delta as one minus the c.d.f.

Table 1. Oil and stock price changes and Saddam Security changes

| Specification | N | Spot oil price ^(a) (\$ per barrel) | | | Ln(S&P 500) | | |
|---|-----|--|----------|--------------------------|---------------------|-----------|-----------|
| | | Adj. R ² | Coeff. | Std. Err. ^(b) | Adj. R ² | Coeff. | Std. Err. |
| First differences^(c) | 147 | 0.078 | 5.38*** | (1.88) | 0.02 | -0.035 | (0.038) |
| Including leads | 143 | 0.086 | | | 0.02 | | |
| T | | | 5.62*** | (1.99) | | -0.036 | (0.040) |
| T+1 | | | 2.39 | (1.45) | | -0.012 | (0.035) |
| T+2 | | | 2.00 | (1.10) | | -0.060* | (0.030) |
| Total effect | | | 10.01*** | (2.15) | | -0.109 | (0.069) |
| Instrumental variables^(d) | 79 | | 12.03* | (6.87) | | -0.442* | (0.211) |
| Long differences | | | | | | | |
| 5 day | 139 | 0.385 | 11.24*** | (2.08) | 0.075 | -0.145*** | (0.057) |
| 10 day | 129 | 0.608 | 10.49*** | (2.32) | 0.197 | -0.197*** | (0.067) |
| 20 day | 109 | 0.779 | 13.09*** | (2.79) | 0.300 | -0.185*** | (0.055) |

Notes: ***, **, and * denote significant at 1%, 5% and 10% levels, respectively.

(a) West Texas Intermediate crude. Oil price regressions include a control for unrest in Venezuela (see footnote 19).

(b) Newey-West standard errors used for difference regressions, controlling for autocorrelation up to twice the difference period.

(c) Differences are differences in number of daily observations in which both the Saddam Security and either oil or stocks traded.

(d) Instrumental variable is Slate.com's "Saddameter"

Table 2. Out-of-sample correlation between stock and oil prices and June Saddam

| Dependent variable: S&P and Spot Oil price | | |
|--|------------------------------|-------------------|
| Time period | 5-day difference regressions | |
| | LN(S&P 500) | Spot Oil |
| Ex-ante sample (9/24/02 to 2/6/03) | -0.137** (0.062) | 10.5** (5.1) |
| Obs. | 67 | 66 |
| R ² | 0.06 | 0.26 |
| Pre-war, out-of-sample (2/7/03 to 3/14/03) | -0.257*** (0.029) | 3.8* (2.2) |
| Obs. | 25 | 25 |
| R ² | 0.57 | 0.03 |
| During the War (3/17/03 to 4/12/03) | 0.275*** (0.081) | -25.1** (10.1) |
| Obs. | 20 | 19 |
| R ² | 0.35 | 0.37 |

Notes: These regressions are 5-day difference regressions as in Table 1, including changes in only the June Saddam as the independent variable.

Table 3. Ex-ante estimates of war effect on sectors and returns during the war

| Sector | Ex-ante estimated effect of war | | | | Alpha During War (3/14 to 4/9) |
|-------------------------------|---------------------------------|---------|-----------|---------|--------------------------------------|
| | Raw return | | Alpha | | |
| | Coeff. | S.E. | Coeff. | S.E. | |
| Consumer Discretionary | -0.229*** | (0.050) | -0.076*** | (0.019) | 0.034 |
| Consumer Staples | -0.005 | (0.035) | 0.084*** | (0.029) | -0.003 |
| Energy | 0.008 | (0.041) | 0.090*** | (0.034) | -0.015 |
| Oil and Gas Drilling | 0.096 | (0.074) | 0.220*** | (0.076) | -0.055 |
| Oil and Gas Equipment | -0.072 | (0.068) | 0.038 | (0.066) | -0.003 |
| Oil and Gas Exploration | 0.101** | (0.045) | 0.170*** | (0.043) | -0.009 |
| Oil and Gas Refining | 0.100 | (0.065) | 0.168*** | (0.058) | 0.010 |
| Finance | -0.184*** | (0.064) | -0.009 | (0.021) | 0.016 |
| Health Care | -0.076** | (0.034) | 0.046*** | (0.017) | 0.029 |
| Industrials | -0.087 | (0.051) | 0.063*** | (0.013) | 0.016 |
| Aerospace and Defense | 0.053 | (0.058) | 0.171*** | (0.037) | -0.001 |
| Information Technology | -0.399*** | (0.084) | -0.145*** | (0.038) | -0.077 |
| Materials | -0.090* | (0.049) | 0.013 | (0.020) | 0.002 |
| Gold mining | 0.438*** | (0.130) | 0.430*** | (0.130) | 0.033 |
| Telecom | -0.162 | (0.102) | -0.021 | (0.062) | -0.041 |
| Transportation | -0.050 | (0.036) | 0.069** | (0.028) | 0.025 |
| Airlines | -0.399*** | (0.090) | -0.260*** | (0.072) | 0.076 |
| Utilities | 0.046 | (0.083) | 0.102 | (0.075) | 0.031 |

Notes: This table repeats the 5-day difference specification in Table 1 for the 11 S&P top-level sector indices and selected sub-indices. Alphas are estimated using a single-factor model, with the beta estimated using daily data from 1/1/1996 to 6/30/2002.

Table 4. Explaining Cross-Country Pattern of the Effects of War


| Dependent variable | Ex-ante war effect (9/24 to 2/6) | Returns during war (3/14 to 4/9) |
|--|-------------------------------------|-------------------------------------|
| Usual Co-movement with MSCI World (Beta) [#] | -.13*** (.03) | -.008 (.016) |
| Net Oil Imports (Fraction of GDP) | -.49* (.27) | .059 (.089) |
| Vulnerable to Attack or Unrest (Turkey, Israel, Jordan, Pakistan) | -.13*** (.03) | .062** (.024) |
| Europe | -.056 (.054) | .023 (.015) |
| Pro-U.S. Europe (Czech Rep., Denmark, Hungary, Italy, Poland, Portugal, Spain, UK) | -.014 (.058) | .007 (.017) |
| Troops committed (Australia, U.K., U.S.) | -.028 (.028) | .029 (.020) |
| SARS affected (China, Hong Kong, Singapore, Taiwan) | -.094*** (.034) | -.060*** (.017) |
| Adj. R ² | .57 | .57 |
| N | 43 | 43 |

Notes: ***, ** and * denote significant at 1%, 5% and 10% levels, respectively (Standard errors in parentheses).

Country-level observations weighted by the inverse of the squared standard error on the dependent variable.

[#] The beta on the MSCI World is estimated using daily data from 1/1/1996 to 6/30/2002. To account for asynchronous market opening hours, the coefficients on the current and lagged change in the World index are added together.

Figure 1: Daily Closing Prices on the Saddam Security



Note: Daily Saddam Security prices are most recent transaction price on as of 4 pm Eastern Time.

Figure 2. Spot Oil Prices and Saddam Securities




Figure 3. Effect of War on Future Oil Prices




Figure 4. Saddam security and the S&P 500




Figure 5. Probability of war and the prices of out-of-the-money puts on the S&P 500




Figure 6

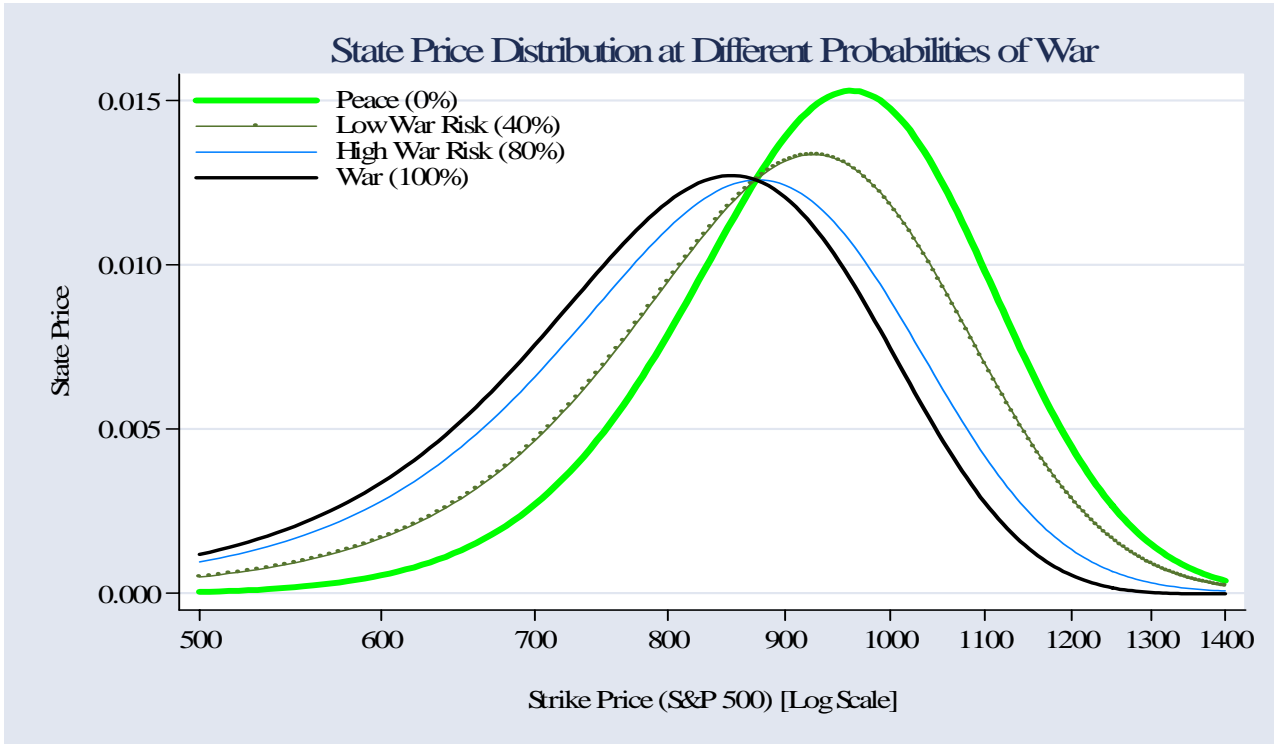


Figure 7

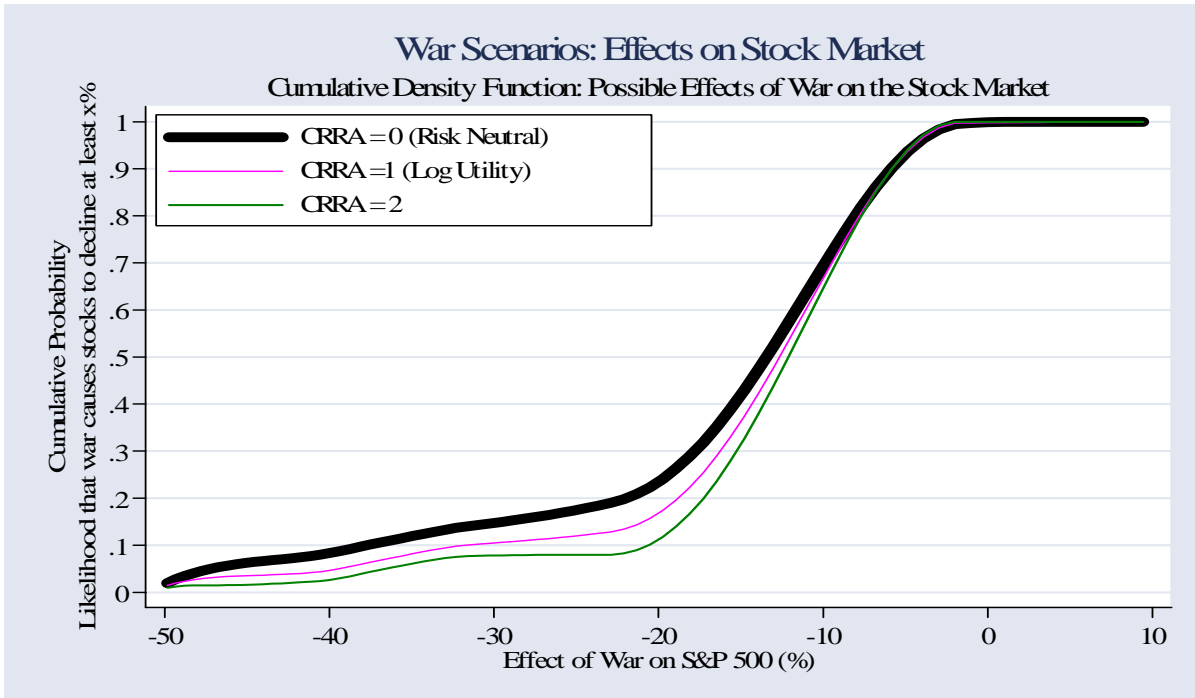


Figure 8. Saddam Security, Equity, and Oil Prices During the War




Figure 9

Estimated Effects of War on National Stockmarkets

Bars represent coefficient estimates; Lines represent 95% confidence interval

