

Peaceful Uncertainty: When Power Shocks Do Not Create Commitment Problems

Muhammet A. Bas*

Robert Schub†

Accepted for Publication at *International Studies Quarterly*

February 11, 2017

Abstract

Does a shock to the balance of power cause the favored actor to exploit its newfound advantage by initiating conflict? The main assumption in the modeling literature on commitment problems as a source of war is that states know and anticipate power shifts. We relax this assumption such that states must estimate future power shifts by looking at past and present capabilities of themselves and their adversaries. We incorporate these estimates and their attendant uncertainty into a model of war and find that while commitment problems are a source of war, the existing models overpredict war by ignoring this dynamic. States continuously updating their estimates and accounting for uncertainty promotes peace. One implication is that an apparent window of opportunity in which the power balance is suddenly favorable is less problematic than previously theorized. This result has applications to nuclear proliferation dynamics as well as shifting power and conflict generally. We find empirical support for the model in tests analyzing power shifts and interstate wars.

*Department of Government, Harvard University.

†Nuffield College and Rothermere American Institute, University of Oxford.

Introduction

Why does bargaining fail when the alternative, war, is costly? The bargaining theory of war identifies power shifts and the resulting commitment problems as one rationalist explanation for such bargaining failures (Fearon, 1995; Powell, 2004, 2006). A rising adversary cannot credibly commit to maintaining the existing division of a contested resource because its newfound power will enable it in the future to demand a revised division in its favor. The declining state, fearful of the bargains it must suffer in the future, may prefer attacking today before the balance of capabilities becomes less favorable. In informal terms, the declining state has a window of opportunity in which to attack (Levy, 1987; Van Evera, 1999).

A state's estimate of the future power balance is central to the commitment problem account of conflict and international politics more generally. Many factors affect the power balance and thus make it difficult to predict with certainty. These factors include considerations partly exogenous to international contestation, such as differential long-run economic growth rates (Gilpin, 1981), domestic coups, revolutions, deaths of monarchs (Blainey, 1988), electoral uncertainty (Garfinkel, 1994), economic troubles (Clark, Fordham and Nordstrom, 2011), natural disasters, or technological hurdles to developing new capabilities (Bas and Coe, 2012). Endogenous arming choices can also entail uncertainty due to imperfect observability of whether an adversary invests in new capabilities (Debs and Monteiro, 2014) or of the progress of an adversary's weapons program (Bas and Coe, 2016). Despite these multifaceted estimation impediments, actors in models with shifting power typically either know the future balance of capabilities or at least know the distribution that determines the future balance of capabilities. In the former case, the declining actor is certain of both the magnitude and timing of its decline (Fearon, 1995). In the latter case, the potentially declining actor may be uncertain about when it will decline as power changes stochastically, but knows that it will do so in expectation (Acemoglu and Robinson, 2000; Fearon, 2004; Krainin and Wiseman, forthcoming). Either case assumes that actors know a good deal about the trajectory of the power balance.

Actors, such as state leaders, of course do not know the precise power balance in the future or even the distribution that determines that power balance. Recognizing this limitation in the information environment, we relax these assumptions. Given the multiple factors affecting future power balances, we posit that states seek observable indicators that are predictive of capabilities and use the past relationship between these indicators and capabilities to predict the latter's future values. If the actual realizations of the military balance diverge from their predictions, states perceive this as a shock and potentially as a window of opportunity in which to strike before the balance reverts back. That is, the power balance changes exogenously and actors do not know the distribution guiding these changes but instead must estimate the future power balance based on its past and present values.¹ We also allow for structural breaks where the historical realizations of the power balance are no longer informative for estimating its future values. This framework, we suggest, offers

¹This differs from studies in which (1) power shifts by a known amount if a state makes a partially unobservable investment in new technologies (Debs and Monteiro, 2014) or (2) actors know the probabilities of stochastic shifts occurring of known magnitudes (Krainin and Wiseman, forthcoming).

a more natural analogue to the challenges leaders confront and the approaches they adopt.

We show that the uncertainty introduced by having actors estimate the future balance of capabilities can promote peace.² This is especially apparent when allowing for structural breaks in the power balance. To develop the logic we focus on when power shocks lead to conflict. A power shock is a highly favorable realization in the balance of capabilities in a given period. If actors know that the power balance is likely to revert to its pre-shock level, then they rationally attack as their adversary cannot credibly commit to providing favorable bargains in the future. The window of opportunity is closing. Introducing uncertainty changes the dynamic. With uncertainty, states may estimate that power shocks are indicative of a more favorable future balance of power rather than a fleeting moment of relative strength. Given the possibility that the newfound advantage will endure, the commitment problem is less likely to bind. The window of opportunity may remain open.

Relating the theory to the existing literature helps clarify the logic. In canonical commitment problem accounts (e.g., Fearon, 1995), a declining state preventively attacks. In our account, as in Fearon (2004), a state that experiences a favorable power shock is the possibly declining state. The newly advantaged actor may anticipate mean reversion in the power balance—the power shock yields only a temporary advantage—and attacks preventively. Departing from prior work, we establish that greater uncertainty over the long run power balance and the possibility of structural breaks undercut the mean reversion expectation. The newly advantaged actor may not attack as it no longer expects a rapid decline. More generally, the theory speaks to the ramifications of *realized* changes in the power balance, not just *expected* changes. Prior studies address whether realized changes lead to conflict if those changes mark a power transition (Organski, 1967; Organski and Kugler, 1980; Houweling and Siccama, 1988; Kim and Morrow, 1992) or if the advantaged state makes onerous demands couched in a rhetoric of power, rather than alternative legitimization strategies (Goddard, 2015). We suggest realized changes in relative power lead to conflict conditional on expectations about, and uncertainty over, future changes (e.g., due to mean reversion).

Empirical tests support the theory. Building from the approach in Bell and Johnson (2015), we construct estimates of a dyad's military balance for periods t and $t + 1$. Comparing the estimate for t and its attendant uncertainty to the realized capabilities in t provides a measure of when power shocks occur. As previously theorized but not established empirically, power shocks increase the probability of interstate war. However, this paper's theory highlights two important refinements that are borne out in the empirical analysis. First,

²Wolford, Reiter and Carrubba (2011) analyze a model that combines asymmetric information about resolve (costs of war) and shifting power dynamics to evaluate their differential effects on the likelihood of conflict. They show that this type of uncertainty can be peaceful, and its resolution can lead to war continuation rather than termination. We similarly conclude that uncertainty in the presence of shifting power can facilitate peace, but we aim to isolate the role of a different, previously unstudied, type of uncertainty about exogenous parameters of the model that directly generate power shifts and lead to commitment problems. To differentiate our account from existing studies on asymmetric information about military capabilities and conflict (e.g., Morrow (1989a); Wagner (1994); Slantchev and Tarar (2011); Bas and Schub (2016)), we assume that the two states share the same information and jointly learn about the exogenous process generating power shifts by observing past and present realizations of capabilities.

the relationship between power shocks and war is conditional. When states that experience a power shock in t estimate that their relative capabilities will only further improve in the future ($t + 1$), a power shock today has limited effect on the probability of war. Second, not all year-over-year shifts in the power balance of a given size have similar implications for war occurrence. The extent of uncertainty matters. A relative shift of a fixed magnitude in a state's probability of military victory is unproblematic if the new balance is similar enough to long range expectations which become wider with greater uncertainty. That same shift, however, is more likely to cause war when it falls outside of the expected range of the power balance. For both refinements, uncertainty about the future distribution of capabilities mitigates commitment problems and promotes peace.

Consider some scenarios that help fix ideas. Imagine two rivals, one of which acquires a powerful new technology such as nuclear weapons. This constitutes a power shock. If the actors believe future power balances are determined by a fixed distribution, then the nuclear state attacks its rival before the latter acquires a similar capability and the power balance reverts. In contrast, if actors must estimate the future balance and believe a structural break has occurred whereby the nuclear state will enjoy a sustained advantage, then the impetus to attack diminishes. Alternatively, consider states engaged in a territorial dispute where one actor suffers a rapid decline in its strength. If we assume a quick reversion to pre-decline levels then the suddenly favored state prefers to attack during its window of opportunity. However, if we allow for the possibility that the decline will persist, negotiated settlements become viable. Chinese territorial disputes in Central Asia plausibly fit this account at the time the Soviet Union dissolved. China confronted much weaker bargaining opponents and likely anticipated they would remain weak.

Our approach and results illuminate implicit assumptions made in past work and suggest some possible refinements. For instance, Gleditsch and Ruggeri (2010) show that irregular leader transitions increase the probability of civil conflict and suggest this occurs because the new leader is temporarily weak before she consolidates power. Outsiders have a window of opportunity before the new leader strengthens her hold on power at which point bargains will be unattractive to the outsiders. This framework assumes that the outsiders are certain the power balance will revert in favor of the leader—that is, that consolidation will occur. Similarly, Nielsen et al. (2011) posit that foreign aid shocks generate commitment problems as a temporarily weak government cannot promise to share the spoils once foreign aid returns. Again, this assumes the actors know that the power shock is transitory rather than a reflection of continuously diminished aid levels. Acemoglu and Robinson (2000) assume that the “elite” and the “poor” know the distribution determining the relative attractiveness of revolution in each period. If the poor become strong but in expectation will return to being weak, the commitment problem binds and the elite extend the franchise. Our argument presents an alternative possibility. If the actors do not know the distribution determining the future attractiveness of revolution, then a favorable realization for the poor may cause them to expect to sustain this favorable position. They believe a subsequent reversion in bargaining power is less likely and the poor can expect to receive

attractive bargains in the future and thus do not demand extension of the franchise.³ Our contention accords with work showing that the effect of economic shocks on democratization is conditional on whether those shocks are transitory or permanent (Kim, 2016). The transitory versus permanent distinction is similarly vital to our account, though we do not assume that states can perfectly delineate between these categories but instead must estimate it with uncertainty.

The remainder of the paper proceeds as follows. The following section presents the theory, discussing various ways actors may estimate the future power balance based on past realizations and how these estimative processes affect the probability of peace. The next section states empirical implications, operationalizes key concepts, and describes the research design. We then present results and robustness tests while addressing concerns due to confounding and links to the arms race literature (Morrow, 1989*b*; Rider, Findley and Diehl, 2011). A final section concludes.

A Stochastic Model of Shifting Power

Commitment problems generated by anticipated shifts in the balance of military power between states is one of the main rationalist explanations for war (Fearon, 1995; Powell, 2006). In the formal IR literature, the mechanism for war in such models is that states that are expected to get stronger in the future cannot commit to not using their future strength to obtain better bargaining outcomes. This makes the declining state better off attacking the opponent in present.⁴ A standard assumption in these models is that the direction of the shift in power and future values of the balance of power are known by all states, or states have perfect information about the probability distribution determining the balance parameter in future periods. In this section, we show that relaxing these assumptions to allow for uncertainty about the magnitude and direction of future shifts improves the prospects for peace.

For presentational purposes, consider a stochastic model of shifting power between states.⁵ The game involves repeated bargaining between two states A and B, over an issue represented by the unit interval. In each discrete time period $t=1, 2, \dots$, the players use the “take it or leave it” bargaining protocol, where A can immediately attack B to end the game, or makes an offer for a division $(x, 1-x)$ to B, which B can either accept, or reject and go to war.⁶ If B accepts A’s offer, it becomes the division of the benefits for that round and players move to the next period. War is modeled as a game ending costly lottery. State A wins the war

³This could cut the opposite direction where the poor are strong in expectation but believe they are more likely to be weak in the long run, become strong in a given period and demand the franchise even when negotiated transfers would suffice.

⁴Commitment problems are also a key impediment to war termination, which third-party enforcement and international organizations help mitigate (Walter, 1997; Shannon, Morey and Boehmke, 2010).

⁵A more general version of this model with complete information is presented in the online appendix in Bas and Coe (2012). To this model, we incorporate imperfect information about the process that determines the balance of capabilities and power shifts.

⁶We represent bargaining offers by A’s share x .

with probability p_t , representing the military balance for period t , and B wins it with $1 - p_t$. The winner receives the benefits for the rest of the game, while the loser receives nothing. A and B's per-period costs of war are c_A and c_B , respectively. States discount the future with a common discount factor $\delta \in (0, 1)$. We assume that the states' utilities are linear in their share of benefits, i.e., for any share of y_i , $U_i(y_i) = y_i$. We assume that states prefer accepting an offer over war when they are indifferent. Finally, we assume that players' preferences and all the exogenous parameters are common knowledge.⁷

In this game, the balance of power shifts stochastically. If bargaining in a given round t ends with a peaceful division of resources, Nature draws the next period's balance parameter p_{t+1} from a cumulative distribution function $F(p)$ defined over $[0, 1]$ with mean p , and with density $f(p)$. Draws close to 1 represent rounds in which A is stronger, whereas B has a military advantage when p_t is closer to 0. For simplicity, in the first version of the model without any uncertainty, we assume that this is a stationary process; the same distribution $F(p)$ generates the realizations across different time periods, and a given period's draw is independent from realizations from previous rounds. The choice to model power changes as exogenous follows prior work (Fearon, 1995; Powell, 1999; Krainin and Wiseman, forthcoming). Though outside the scope of this piece, power can shift endogenously, for instance due to arming or military threats (Chadefaux, 2011; Slantchev, 2011; Debs and Monteiro, 2014; Bas and Coe, 2016) or as a consequence of domestic political conditions (Chapman, McDonald and Moser, 2015). If present at all, uncertainty in models with endogenous arming is asymmetric as the arming state naturally has an informational advantage concerning its arms investments. In contrast, uncertainty about the stochastic element of power shifts in our account is mutual.

First, following Bas and Coe (2012), we characterize the equilibria of this game under perfect information about $F(p)$. We focus on the most efficient subgame perfect equilibria in stationary Markov strategies (MPE) that minimize the probability of war. MPE are subgame perfect equilibria that are based on strategies that condition behavior only on payoff relevant information, summarized by a *state* variable (Fudenberg and Tirole, 1991). In our model, the payoff relevant state in any given period is that period's balance parameter determined by Nature's move, and states A and B reach a peaceful settlement or fight based on their strategies that condition their behavior on that period's military balance. The stationary MPE takes a particularly simple form, in which players use cutpoint strategies defined based on the balance parameter to attack their opponent. In particular, A prefers attacking B in periods when A is significantly stronger than the long term average, while B prefers doing the same in periods where B is advantaged. When war occurs, it is due to the standard commitment problem as power is expected to shift back to the long term average, ending the temporary advantage one state enjoys. Intuitively, A prefers war in equilibrium when the expected payoff to war exceeds the expected payoff from bargaining. The latter quantity includes A receiving the most generous bargain ($x=1$) in the current period, the expected decline in bargaining leverage in the future due to

⁷We will relax some of these assumptions later when we introduce uncertainty about the probability distribution determining the balance parameter.

diminished relative capabilities, and the likelihood of future war on possibly weaker terms. The following lemma describes the stationary MPE and characterizes the cutpoint balance parameters \bar{p} and \underline{p} that states use in their decisions to attack the opponent or not:

Lemma 1. *In the stationary MPE, states condition their behavior on the current period's balance parameter p_t . After observing Nature's draw of p_t from $F(p)$ in any given period t ,*

- *if $p_t > \bar{p}$, A attacks and ends the game; B accepts any offer $x \geq 0$;*
- *if $p_t < \underline{p}$, A prefers making any offer $x \geq 0$ to war; B rejects any offer and attacks;*
- *if $\underline{p} \leq p_t \leq \bar{p}$, A offers $x^* = \min\{1, \bar{x}\}$; B accepts any offer $x \leq \bar{x}$,*

where \underline{p} is the value of the balance parameter at which B is indifferent between war and accepting the most generous offer A can give at that period ($x=0$); \bar{p} similarly represents the balance at which A is indifferent between war and the most generous feasible share in that period ($x=1$); and \bar{x} is the bargaining offer that would leave B indifferent between accepting and rejecting. These parameters are formally defined and derived in the Supporting Files.

The above model shares an important assumption with many existing models that focus on commitment problems as a source of war: states can anticipate with certainty the expected direction of the shift in power, or what the exact future distribution for the balance of power will be. In practice leaders are not privy to this information and instead likely must estimate the future power balance based on observable predictors and past values. We can relax the model's assumptions in a few ways to make it more realistic. For instance, we can assume instead that states do not know the distribution of p_t , $F(p_t)$, but approximate it based on the past and present realizations of the parameter. An important question this new assumption introduces is how this type of uncertainty affects states' incentives to use force when they enjoy a temporary advantage. Alternatively, if states allow for the possibility of structural breaks in their estimation of the distribution for the balance parameter, i.e., the possibility that the balance shifts significantly in certain periods to a new distribution, would this make war more or less likely? In other words, how do \underline{p} and \bar{p} change under each scenario? Below, we compare these scenarios, and using computer simulations, discuss their implications for the likelihood of peace in models of war.⁸

Various Levels of Uncertainty

This section compares various forms and levels of uncertainty with respect to the distribution of the balance parameter. First, as a baseline, we start with the complete information case presented above, which represents the standard model in the literature—i.e., when states know the true distribution and the trajectory of

⁸Since the non-linear system of equations in the lemma does not have a closed form solution, in our simulations we employ numerical optimization using the BB package in R to calculate the cutpoints for a given set of exogenous parameters.

the balance of power (Scenario 0). In any given period, both states know $F(p)$ and this is common knowledge. To fix ideas, consider the dyadic power balance between, say, the US and USSR (or Russia, depending on the time period), which fluctuated over time due to perturbations that were in part exogenous to bilateral contestation. These perturbations could arise from domestic demands for a shift of spending toward butter at the expense of guns, unanticipated technology hurdles, or fiscal shortfalls due to economic downturns or reconstruction needs following natural disasters. Scenario 0 assumes that leaders in the US and USSR know the probabilities of all possible perturbations, and thus the probabilities of all possible power balances, and that these probabilities are constant across time. In other words, it assumes leaders enjoy a wealth of information.

With this assumption, cutpoints \bar{p} and \underline{p} do not depend on time as states use constant cutpoints to determine whether to attack. This model is described above and the equilibrium behavior is summarized in Lemma 1.

Scenario 1 relaxes this assumption. Instead of knowing $F(p)$, states estimate this distribution in any period t using past and present realizations of the balance parameter from periods 1 through t . Put differently, Scenario 1 assumes that US and Soviet leaders knew that perturbations—e.g., technological breakthroughs—could cause the power balance to shift over time. Moreover, it assumes that leaders knew that the probability of various shifts occurring remained constant over time. However, US and Soviet officials did not know the actual probabilities with which perturbations would occur. Instead, they had to estimate the likelihood of future exogenous changes based on their observations of past changes.

In this scenario, the equilibrium concept becomes a Markov Perfect Bayesian Equilibrium (MPBE) as the states have uncertainty about Nature's distribution that generates the balance parameter and they can only observe draws from it. The Markov state variable becomes the vector $P = \{p_1, \dots, p_t\}$ of realizations from the distribution $F(p)$. Nature as a non-strategic player starts the game by selecting a distribution $F(p)$ from a space of distributions Ω , and makes draws from this distribution at each period. In this scenario, we assume that the existence of a single distribution is common knowledge. While different forms of belief updating are possible under this equilibrium, we assume that states form their beliefs by estimating the distribution using the Method of Moments estimator based on the observed values of the balance parameter, as described below. With this assumption, and given the fact that both states observe the same draws, in each period states hold the same belief about the underlying distribution. We also assume that the prior belief at the beginning of the game before any observed draws from the distribution is equal to the mean of $F(p)$, which is $p_0 = .5$ in our simulations.⁹

For our simulations, we assume $F(p)$ is $Beta(\alpha, \beta)$ with probability density function $f(p) = \frac{p^{\alpha-1}(1-p)^{\beta-1}}{B(\alpha, \beta)}$ defined over the $[0,1]$ interval, where $B(\alpha, \beta) = \int_0^1 p^{\alpha-1}(1-p)^{\beta-1} dp$ is the Beta function. We choose the Beta distribution for its flexibility in incorporating a variety of unimodal or bimodal shapes and skewness,

⁹The effect of the prior on the cutpoints is minimal, as its weight in states' beliefs about the distribution gets smaller with more observations from the process. Not surprisingly, our simulation results are robust to alternative prior specifications.

capturing different asymmetries of power.¹⁰ We assume that, based on past and present realizations of p , states estimate α and β in each period with $\hat{\alpha}_t$ and $\hat{\beta}_t$. In particular, in any period t , upon observing the sequence of balance parameter values $P = \{p_1, p_2, \dots, p_t\}$ up to that period, and the corresponding sample

mean $\mu_t = \frac{1}{t} \sum_{m=1}^t p_m$ and the standard deviation $\sigma_t = \sqrt{\frac{1}{t-1} \sum_{m=1}^t (p_m - \mu_t)^2}$, states calculate $\hat{\alpha}_t$ and $\hat{\beta}_t$ as

$$\hat{\alpha}_t = \mu_t \left(\frac{\mu_t(\mu_t - 1)}{\sigma_t^2} - 1 \right) \quad (1)$$

$$\hat{\beta}_t = (1 - \mu_t) \left(\frac{\mu_t(\mu_t - 1)}{\sigma_t^2} - 1 \right) \quad (2)$$

While the results are robust to using alternative estimators for the shape parameters of the Beta distribution, we assume that states use the Method of Moments estimator due to its simple form and intuitiveness.¹¹ Our central point is that actors must form estimates about the future power balance, not that they necessarily do so using the specific estimator we describe here.

Scenario 1 assumes that, to determine the cutpoint value for attacking, states take into account both the historical values and the current period t 's realization. In other words, states allow for the possibility that past balance may not provide a complete characterization of the future balance, and use the current period to update their estimate of the underlying distribution. Accordingly, a state's future expectations become more favorable following a favorable realization in a given period. The state revises its cutpoint which reduces the likelihood that the current favorable balance is sufficient to incite an attack. The Supporting Files show that the distinction between using only past realizations versus using past and present realizations to estimate the distribution is an important factor for the overall probability of peace.¹²

Finally, in determining the future distribution, states may also take into account the possibility that there is a structural break in the stationary distribution of the balance parameter $Beta(\alpha, \beta)$, in the sense that a new Beta distribution with different shape parameters α^* and β^* generates the balance parameter (Scenario 2). In

¹⁰The choice of Beta also follows when state military capabilities are viewed as non-negative random variables. Suppose that states A and B have capabilities $m_A \geq 0$ and $m_B \geq 0$ that are independently distributed $Gamma(\alpha, \theta)$ and $Gamma(\beta, \theta)$, respectively. The balance between the two states, expressed in the ratio form, $p = \frac{m_A}{m_A + m_B}$ will then be distributed $Beta(\alpha, \beta)$. To maintain the link between the theoretical and the empirical model, we employ the commonly used ratio form to calculate the dyadic balance of capabilities in our empirical analyses. Moreover, the observed power balance often approximates a Beta distribution. For example, historical "realizations" of the Egypt-Israel power balance are distributed roughly $Beta(33,21)$.

¹¹An alternative is the Maximum Likelihood Estimator (MLE) for the parameters of the Beta distribution, which, unlike the Method of Moments estimator, does not have a closed form solution.

¹²For our analyses, we also assume that states take all the available history into account in estimation. Our substantive results get stronger if instead we allow states to only focus on more recent periods (e.g. the most recent k periods) or temporally weight the observations to estimate α and β . As the weight of the present period's draw in estimating the distribution increases, the peaceful effect of uncertainty is also amplified.

each period, states evaluate if that period’s realization together with the historical record is consistent with such a break. If states are convinced that a break has occurred at some point t' in the past, they only take into account the periods after the identified structural break in t' in approximating the future distribution.¹³ For instance, if a state enjoys a favorable balance for a given number of consecutive periods, it could conclude that the distribution shifted permanently with a more favorable mean, and the opportunity it enjoys is not necessarily temporary.¹⁴ Returning to the running example, Scenario 2 assumes US and USSR officials still must use past and present realizations to estimate the probabilities with which various exogenous shifts to the power balance will occur. But now they also must allow for the possibility that structural breaks, akin to the dissolution of the Soviet Union, could cause the process governing these perturbations to change. Upon recognizing that a break occurred, officials deem Soviet-period perturbations uninformative for estimating perturbations in the future US-Russia dyadic balance. Compared to prior scenarios, we believe this one provides a better distillation of the knowledge leaders have and the estimative challenges they confront.¹⁵

In sum, we consider three scenarios with varying degrees of uncertainty. These are:

- **Scenario 0:** Actors know the true distribution, $F(p)$, determining the balance parameter.
- **Scenario 1:** Actors use past and present realizations of the balance parameter to estimate $F(p)$ and preclude the possibility of structural breaks.
- **Scenario 2:** Actors use past and present realizations of the balance parameter to estimate $F(p)$ and allow for the possibility of structural breaks.

Figure 1 compares these three scenarios with respect to the level of knowledge about the underlying distribution for the balance parameter. The plot is based on an example 80-period history of realizations

¹³More formally with this scenario, Nature starts the game by drawing $F(p)$ and $B(p)$ from Ω to represent the distributions before and after the structural break, and selects b from $t = \{0, 1, 2, \dots\}$ to determine the timing of the break. We assume that states do not know $F(p)$, $B(p)$, or b , and their beliefs under the MPBE take a particular form: by observing realizations across periods, they detect the structural break and estimate the distribution after the most recent break to use in their equilibrium strategies.

¹⁴In our simulations, we use the `strucchange` package in R to estimate structural breaks. Following Bai (1994), given the vector of observed past and present values $P = \{p_1, \dots, p_t\}$ in each period t , the code identifies \hat{b} that minimizes the residual sum of squares (RSS) of the process $p_i = \beta_1 + u_i$ for $i < b \leq t$, and $p_i = \beta_2 + u_i$ for $b \leq i \leq t$ ($\hat{b} = t$ indicates no structural break is detected).

¹⁵Alternatively, consider environmental shocks as the mechanism driving power shifts. In this case, Scenario 0 equates to the actors knowing with certainty the probability of environmental perturbations of varying magnitudes in any given period. For Scenario 1, actors use the historical record of such events to estimate the likelihood of future events. This is a more plausible assumption for many disasters, such as earthquakes or volcanic eruptions. Finally, for Scenario 2, the underlying probability of environmental events of varying magnitudes could shift over time. This might be particularly relevant for events linked to broader patterns of climate change which is roughly akin to a structural break. If climate change does affect the frequency and severity of such events, as the Intergovernmental Panel on Climate Change (IPCC) predicts it does, then the observed historical rates may no longer be sufficient for estimating the future likelihood of environmental shocks. We note later in the paper that environmental disasters that would be large enough to cause significant shifts in relative capabilities are likely rare. However, the IPCC’s conclusions suggest that this may change in the future.

from a Beta distribution. From period 1 through 40, a symmetric $Beta(\alpha = 2, \beta = 2)$ with a mean $p = .5$ is the data generating process for the balance parameter. There is a structural break in favor of state A after period 40 and the balance parameter is then generated from $Beta(\alpha = 4, \beta = 2)$ with a mean $p = .67$. Dotted lines in Figure 1 represent the distribution mean before and after the structural break. Since Scenario 0 assumes that states know the stationary distribution with certainty, the estimated mean does not depend on actual realizations of the balance parameter and does not change across time except for the structural shift at period 40. Next, the dashed lines represent the dynamically estimated mean $\mu_t^p = \frac{\alpha_t}{\alpha_t + \beta_t}$ in each period based on Scenario 1, which does not take into account the possibility of a structural break. The solid line belongs to Scenario 2, in which states not only estimate the distribution, but also dynamically evaluate the possibility of a structural break, and if they detect one, ignore the realizations before the most recent break in their estimation of the distribution. The three dimensional plots are the densities of the Beta distribution that the states estimate in each period according to Scenario 2. Finally, the points represent the actual realizations of the balance parameter in this example 80-period history.

First, note that the uncertainty about the underlying distribution implies that states base their decisions on potentially different Beta distributions at each period for Scenarios 1 and 2 versus Scenario 0. Also, unlike Scenario 1, Scenario 2 estimates quickly capture the new distribution with a short delay after the structural break occurs. A delay occurs with Scenario 2 because states might need a sufficient number of draws from the new distribution before they detect a structural shift. In contrast to Scenario 2, the Scenario 1 estimate slowly drifts away from the old distribution, but does not converge to the new one's mean as it still uses draws from the old one, which should be irrelevant for the estimation procedure. One implication is that the war likelihood increases immediately after the structural break as states may mistake lasting advantages as fleeting ones. The risk is particularly acute for Scenario 1 which cannot quickly detect the structural break.

How do various levels of knowledge in Scenarios 0-2 affect the cutpoints states use for attacking the opponent? Figure 2 compares these scenarios for an example parameter configuration.¹⁶ The figure is based on a 40-period history, which involves a structural break at period 15 that shifts the distribution in A 's favor. In particular, for periods 1-14, the balance parameter is drawn from a $Beta(4, 4)$, which then shifts to $Beta(4, 2)$ at and after period 15. In each period of a given history, we calculate states' cutpoints using the various levels of uncertainty specified in Scenarios 0-2. We repeat the history 1000 times and plot the median values from these repetitions for each scenario.

Black horizontal dotted lines represent the Scenario 0 cutpoints before and after the structural break. For each segment, the top line represents A 's equilibrium cutpoint – for any value of p above the line, A prefers attacking. Likewise, B attacks for balance values below the bottom horizontal line. Since this scenario assumes that states know the distribution exactly, the cutpoints are constant across the periods except for the shift after the structural break at period 40.

¹⁶We set $\delta = .9$, $c_A = .3$, $c_B = .3$.




Figure 1: Comparison of estimated balance parameter means for different types of uncertainty. Points represent the realized balance parameter in each period. Dotted lines depict the mean for Scenario 0 in which actors know the underlying distribution determining the power balance and that a structural break occurred after period 40. Dashed line represents the dynamically estimated mean for Scenario 1 which does not allow for structural breaks. Solid line represents the dynamically estimated mean for Scenario 2 which does allow for structural breaks. Density curves represent the estimated distribution under Scenario 2.

The main difference between Scenario 1 and Scenario 2 is whether they allow for the possibility of structural breaks. Thus, the two scenarios produce very similar results before the structural break. Solid red lines represent their cutpoints. After the break, they diverge, and green dashed lines represent Scenario 1. As in Scenario 0, the top and bottom lines represent A and B's cutpoints for attacking, respectively. Before the structural break, the average cutpoints for both Scenarios 1 and 2 are more extreme compared to Scenario 0, due to states using the present period's realization to update the approximation for the distribution. In other words, each state is less likely to attack overall in Scenario 1 (and 2) than Scenario 0 before the break. This occurs because favorable realizations cause the states to adjust the estimated distribution in their favor which reduces the expected magnitude of mean reversion, or the extent of the shift back toward the long run




Figure 2: Comparison of the cutpoints for different types of uncertainty. Horizontal dotted black lines depict cutpoints for Scenario 0; dashed green lines for Scenario 1; and solid red lines for Scenario 2.

trend.¹⁷ Uncertainty about the underlying data generating process increases the probability of peace.

Since Scenario 1 does not allow for the possibility for a structural break, it continues using the whole history to approximate a single distribution and does a poor job capturing the post-break distribution. The cutpoints slowly approach but do not converge on Scenario 0's cutpoints during the 40-period history considered in the figure. Finally, Scenario 2 allows for structural breaks in the approximation of the distribution and captures the new distribution relatively quickly. This shift is not as immediate as Scenario 0, however, as it takes states a few periods to identify a structural break based on past and present realizations.¹⁸

Figure 3 compares the three scenarios in terms of the overall probability of peace.¹⁹ Like above, the black dotted line represents the probability of peace for Scenario 0. The red solid lines belong to Scenario 2. Scenario 1 shares the red solid line with Scenario 2 before the structural break, and is represented by the

¹⁷This dynamic attenuates under some conditions. If there is (1) a sufficiently large number of periods, (2) equal weighting of all realizations across time, and (3) no structural breaks, then each new realization induces less adjustment to the prevailing estimate as the weight of the present period's draw from the distribution converges to zero in estimating the distribution, and there is convergence in the cutpoints across the scenarios.

¹⁸Note that these figures show the average trend based on multiple histories. In a single history, it is always possible that with scenario 2, a false positive structural break is detected in either direction even though it does not exist in the true data generating process, and war could result due to this. In comparing scenarios 1 and 2, the point we would like to emphasize is that in expectation, scenario 2 is more peaceful overall due to its ability to capture shifts when they take place and avoid an unnecessary war between the two states.

¹⁹The y-axis values are meaningful for evaluating the *relative* peacefulness of the scenarios. Their absolute levels shift with alternative parameter values.




Figure 3: Uncertainty and the probability of peace. The dotted line depicts the probability of peace for Scenario 0; the dashed line for Scenario 1; and the solid line for Scenario 2.

green dashed lines after the break. As expected, both Scenarios 1 and 2 are more peaceful than Scenario 0 through the first 14 periods before the break. States use more extreme cutpoints in determining whether to attack when they take into account the present round's favorable realization and update the distributional estimate. Intuitively, Scenario 2 is more peaceful than Scenario 1 after the structural break, as Scenario 2 allows states to recognize the break while Scenario 1 does not. Strikingly, right after the break is a dangerous period for Scenario 2 (and trivially for Scenario 1 as it does not allow for breaks) since states do not recognize the break immediately. The faster states capture the break and recognize the newly favorable trend, the shorter the period in which the risk of war is elevated.²⁰ It is also worth noting that there is a high probability of peace once a state does recognize the break. Because of the short history of realizations under the new distribution, the present period's realization has more weight in approximating the parameters of the distribution, and states cutpoints get more extreme.

The bottom line is that allowing for uncertainty about the underlying distribution for the balance of power and incorporating the possibility of structural breaks in states' estimation calculus make peace more likely overall, with the exception of short periods right after significant structural shifts in the balance of power occur, before states can recognize them. Uncertainty about the underlying distribution produces more peace than setups without such uncertainty. And uncertainty from the possibility of structural breaks produces more peace than a setup which precludes such breaks with certainty. Models overpredict war when assuming leaders are privy to more information about future power balances than they likely have in reality.

²⁰Frequent breaks would produce a higher probability of war if states require multiple periods to recognize them as in the model. In reality, breaks may be more recognizable given an observable disjuncture (e.g., the splintering of the Soviet Union).

Mapping Theory to Data

This section has three objectives. It states testable implications from the theoretical model. It operationalizes concepts from the model such as how actors estimate the distribution determining future power balances, what constitutes a power shock, and when do we observe structural breaks. Finally, it lays out the research design for testing the argument.

Testable Implications

We assess four theoretical implications. These follow from Scenario 2 in which states (1) update expectations about the balance parameter based on past and present realizations and (2) allows for structural breaks in the underlying distribution governing the balance parameter. First, *favorable power shocks increase the probability of war*. War occurs when states have a sufficiently favorable “realization” of their military capabilities because they expect the favorable balance is temporary—that is, they anticipate mean reversion. In the model’s terms, states attack when their realized military capabilities exceed their estimated cutpoint. Because power shocks are indicative of a favorable realization, such shocks increase the probability of war. This expectation is consistent with past models (Fearon, 2004) but is untested in the context of interstate wars.

We can refine the first implication. As the theory emphasizes, not all favorable military realizations increase the likelihood of war. Commitment problems are tractable provided a favorable realization is below the cutpoint. A second implication follows: *only sufficiently large favorable realizations cause war while modestly favorable realizations are unrelated to war onset*.

Incorporating the possibility of structural breaks into states’ estimates of military capabilities offers a third implication. Consistent with Scenario 2, the broader strategic context matters for evaluating the effect of a positive shock (or favorable realization). Stylized scenarios help clarify the idea. If a state has a favorable shock and estimates that this new realization is indicative of sharply more favorable environment—that is, indicative of a structural break—then the cutpoints increase and states avoid war.²¹ This is equivalent to a state in the model quickly recognizing that a break occurred. If a shock of the same magnitude occurs but the state believes this is an aberration as the strategic environment is not shifting in its favor, then this shock increases the probability of war. Shocks that occur when a state is pessimistic about the strategic environment pose an even more severe danger. The commitment problem is especially likely to bind as the state estimates the window of opportunity is rapidly closing. War is most likely in this case. A third testable implication is that *a state’s estimate of the power balance trajectory conditions the effect of a positive realization in a given period*.

²¹Though substantively different, the intuition relates to the contention in Krainin and Wiseman (forthcoming) that any mechanism that allows a newly advantaged state to expect that its advantage will persist—such as through transfers—mitigates commitment problems.

A final implication makes the importance of estimation uncertainty especially clear. A power shift between periods above a given magnitude incentivizes conflict in classic accounts of commitment problems (e.g., Fearon (1995)). This is not the case in our approach. Consider an example where years stand in for periods. The implications of a year-over-year change in the balance parameter of, say, 10% depends on the extent of estimation uncertainty states have over the dyadic balance. A 10% shift is war-promoting if the new balance constitutes a shock versus expectations for that year's balance. In contrast, that same 10% shift is not particularly problematic if the resulting balance falls inside the cutpoints, or within the expected range for the dyadic balance. Greater estimation uncertainty increases the width of this range. A fourth implication is that *shocks, as benchmarked against expectations, increase the probability of war even after conditioning on the absolute year-over-year change in the dyadic power balance.*

Operationalizing Concepts

Consistent with the theory, we specify that states estimate future values of the military balance based on its past realized values. This matches our suggestion that states would not reliably know the true data generating process but instead must estimate it using observed values. The empirical approach builds from Bell and Johnson (2015). Bell and Johnson (2015) begin by fitting models where military capabilities in period $t - 1$, defined as the average of a state's share of global military personnel and global military expenditures as coded by CINC scores (Singer, Bremer and Stuckey, 1972), is the outcome variable and a series of observable predictors in period $t - 2$ are the independent variables. These predictors include measures for economic potential, rivals, the toll of recent military conflicts, regime type, alliances, and the past year's value of the outcome variable. A separate OLS model is run for each period $t - 1$ using all data until that point with country fixed effects and where the unit of analysis is the country year.²² Then, using the coefficient estimates for the predictors, we can plug in the values of the predictors in period $t - 1$ and estimate military capabilities for period t .

Rather than have states only look at past values of the military balance, we model them using the past relationship between predictors and values of the military balance. Does this approach for generating estimates of future capabilities approximate the actual processes that states employ? There is reason to believe it does based on both theory and empirics. Leaders assessing the capability trajectory of other states must rely on observables to help craft their assessments. Past work specifies the type of observables most likely

²²With the fixed-effects specification, intercepts vary by country while coefficients for the predictors do not. The predictors are operationalized as follows. To proxy for a state's economic potential, Bell and Johnson (2015) use the four non-military components of CINC scores—a state's urban population, total population, energy consumption, and iron and steel production. The rivalry indicator is based on the Thompson and Dreyer (2010) coding. The natural log of interstate war battle deaths, coded by Sarkees and Wayman (2010), provides the measure for the recent toll of military conflicts. The regime type predictor is *polity2* scores from the Polity IV project (Marshall and Jaggers, 2002). Finally, the alliance variable indicates whether a state was party to an alliance that necessitated military contributions in that year as coded by the Alliance Treaty Obligations and Provisions project (Leeds et al., 2002).

to influence assessments (Fordham, 2011). These include a state’s economic potential, whether they face security threats, and domestic institutional arrangements. The predictors used in the empirical approach are designed to reflect these factors.

Beyond relying on past theoretical work, we look to actual evaluations of trends in adversary capabilities to assess the approach’s construct validity. A number of US intelligence community assessments, National Intelligence Estimates (NIEs), of trends in rivals’ capabilities are declassified. These NIEs consistently highlight factors related to the observable predictors used in the empirical approach. Consider annual NIEs from the late-1950s on five-year trends in Soviet capabilities. These discuss economic prospects at length. For instance, to gauge future capabilities the 1958 NIE examines current levels of Soviet industrial production, noting “severe” shortages in steel and discussing population dynamics affecting the labor supply.²³ Besides economic potential, the NIE addresses Soviet perceptions of the threat environment, assessing its military objectives and regional interests and challenges. It further discusses trends in relations with allied Communist states. Analysis of economic potential, security threats, allies, and the domestic political landscape is a fixture of all similar NIEs from this period.

Intelligence analysts of course draw on more granular data than raw energy consumption as used in the predictive model. For instance, Soviet pronouncements about industrial plans or force reductions are valuable sources of information that our approach, building from Bell and Johnson (2015), does not systematically capture. Similarly, analysts may anticipate major technological breakthroughs—such as developing nuclear weapons—which our approach cannot. However, this only dampens our results as states may attack following a shock due to a qualitative shift in technology that the measure will not recognize.²⁴ Despite these limitations, there is meaningful overlap in the observable indicators used to form future assessments in the NIEs and in the statistical predictions we employ. In addition to similarities in the estimative process, we also observe similarities in the military capability predictions. NIEs throughout this period anticipate declines in Soviet force sizes and stability in military expenditures.²⁵ Similarly, the predictive model anticipates declines in the Soviet share of military capabilities throughout this time period. The Supporting Files further addresses the issue, describing the similarity of our predictions to those in a RAND study that the Office of the Secretary of Defense funded. In sum, there are substantive similarities, at least anecdotally, between our statistical approach and the approach of the intelligence community in terms of the estimative process and the estimates themselves.

Our setup differs from past studies in two ways. First, whereas Bell and Johnson (2015) require a point estimate of future capabilities, we need to estimate a full distribution. To do so, we use the uncertainty

²³National Intelligence Estimate: Number 11-4-58, “Main Trends in in Soviet Capabilities and Policies, 1958-1963,” 23 December 1958. Accessed July 20, 2016 at https://www.cia.gov/library/readingroom/docs/DOC_0000267651.pdf.

²⁴In a further effort to account for the measure’s inability to account for nuclear weapons, we show that all results hold when restricted to a pre-1945 temporal window.

²⁵Military expenditure projections are rarely explicit in the reports. Instead, they are implicit based on projections for submarine, aircraft, and surface naval vessel construction as well as projected personnel levels.

around the point estimate to capture the broader expectation for where next period's capabilities will fall. For a given dyad we draw 1,000 estimates of each state's capabilities for period t and pair each of the 1,000 draws to construct a measure of the balance parameter using a standard ratio form reflecting an initiator's share of dyadic capabilities. This provides an estimated distribution for the balance of capabilities for each dyad in each year. This is an admittedly imperfect mapping from the theory where actors estimate a Beta distribution that governs the balance parameter. However, this empirical approach using draws from the predictive model to construct an estimated distribution captures the basic intuition where states form expectations and some realizations are seen as consistent with or extreme versus those expectations. A second difference from prior work is that estimates serve a dual purpose in the theory as they provide cutpoints that determine when sufficiently large shocks occur and offer context for the trajectory of a state's capabilities (i.e., whether a structural break occurred). Consequently, we estimate military capabilities for both period t and $t + 1$ where the former helps define power shocks and the latter shapes projections about the broader trajectory.

After estimating a full distribution for a dyad's balance of capabilities in period t , we can specify a cutpoint above which a favored state attacks. We define the cutpoint as the estimate's upper bound of the 95% confidence interval.²⁶ While in part divorced from the analytical definition of the cutpoint provided in the theory, using the upper confidence bound captures the underlying logic. States only attack when a realization in the military balance is highly favorable versus *ex ante* expectations and they consequently anticipate a sizable reversion to the mean power balance. It follows that a power shock is defined as a realization of military capabilities in period t that exceeds the cutpoint for that period. This approach differentiates between favorable realizations that are above the mean estimate (but below the cutpoint) and those that are shocks in that they significantly diverge from expectations (above the cutpoint).

Consider the US-Hungary directed dyad in 1926 as an example. Leaders in the empirical setup examine the historical relationship between observable indicators in the US and Hungary, such as their regime types or the presence of rivals, and their military capabilities up to 1925. Plugging the predictors' values in 1925 into the model of the historical relationship yields estimates for each state's capabilities in 1926. We then take draws from these estimates to construct an estimated directed dyadic balance parameter with a mean of 94.8 and 95% confidence bounds of 93.3 and 96.4 where these values represent the US share of dyadic military personnel and military spending. We then observe the realized military balance value of 1926 is 94.2. Because this falls below the upper confidence bound, this is not coded as a positive shock. In aggregate, positive shocks occur in 15% of the politically relevant directed dyad years for which there are estimates.²⁷ To assess the predictive validity of the estimate, we also define a standardized score as the absolute difference between the estimated mean and the realized value, divided by the estimated mean. In

²⁶All results are robust to using 99% confidence bounds to determine cutpoints.

²⁷Estimates are missing due to either missing values in the CINC data set or when the predicted value for either state's capabilities has an upper confidence bound less than or equal to zero. OLS generates negative predictions in some country years which by definition would be coded as power shocks. We drop these cases.

the US-Hungary example this equals $\frac{94.8-94.2}{94.8}$. The measure's median value is 5%.²⁸

The empirical setup thus far provides measures for the estimated dyadic balance in t and an indicator for power shocks in that period. We still require a measure for the expected future trends in relative capabilities between states to assess how the broader context conditions the effect of power shocks. Bell and Johnson (2015), using the basic approach already outlined, provide this exact measure. Their substantive focus is on projected future capabilities. The measure, which we adopt, reflects the expected future trajectory of the balance of power, taking on negative values when the specified “initiator” state anticipates a favorable shift in $t + 1$ and positive values when the “target” expects a favorable shift.

Expected changes to the relative balance of dyadic capabilities reflect future power trends. Tying back to the theoretical implication, we anticipate that the measure of future trends (estimates for $t + 1$) conditions the effect of positive shocks in t . Very favorable expectations about future relative capabilities are akin to states benefiting from and recognizing a favorable structural break in the model.²⁹ Accordingly, when the “initiating” state anticipates a favorable shift in relative capabilities in $t + 1$, a positive shock in t should have little effect on the probability of war. When the state's future projections are less optimistic, shocks today (t) do increase the probability of war, presenting a window of opportunity. Hence the interactive relationship between future trends and current realizations. As constructed, future trend reflects the one-year outlook for the dyadic balance. This differs from conceptions in which actors anticipate a persistent rate of change in the power balance, as in Krainin (forthcoming). A persistent trend in the power balance is indicative of a non-stationary process, which is outside the scope of our theory. Facing a sufficiently adverse and persistent trend, the state disadvantaged by a power shock may attack rather than suffer unfavorable bargains in the future. In our theory without power balance trends, only the temporarily advantaged state has incentives to attack. That being said, this omission does not bias the results in favor of our theory. In fact, if the disadvantaged state has incentives to attack due to such trends, since those cases are not recorded as shocks in our existing analysis, this would mean that the estimated effects of shocks are biased downwards.

Mid-1930s Germany and France illustrate the broader concept. Based on data from preceding years, the model predicts a directed German-French balance parameter in 1934 of 38.1 ± 1.0 . The actual balance in 1934 was 46.3, thus representing a positive shock. A theory devoid of the broader trend of relative

²⁸Recall that the objective for the estimation process is to use past observations in a way that plausibly approximates the estimation approach state leaders or intelligence agencies might use. The predictive accuracy of the measure is a separate matter. A better predictor would reduce the number of observations coded as positive shocks. However, greater accuracy is only desirable insofar as leaders and their advisers are more accurate. Fitting more complex predictive models—for example, using random forests—is inappropriate unless we believe leaders, including those in 1840 for instance, used this approach. That said, we check and find that more conservative modeling approaches with wider cutpoints, and thus fewer shocks, yield similar results.

²⁹A positive shock coupled with a favorable outlook for the future is admittedly not identical to the structural breaks in the theory. However, it is a reasonable approximation of the intuition we wish to capture. A shock coupled with a favorable future expectation represents a major change in the estimated distribution determining the dyadic power balance. This is substantively similar to the theoretical discussion of structural breaks, which are major changes in the estimated distribution determining the power balance.

capabilities might predict that the advantage granted by the shock is fleeting, leading Germany to quickly attack. However, the future trend is crucial. Our approach estimates that the balance of military capabilities between Germany and France would shift an additional 5.6% points in Germany's favor in 1935, which is in the top quartile of most favorable trend expectations. Consequently, Germany's favorable power realization in 1934 is not indicative of a fleeting window of opportunity but rather of a more lasting change in the balance of power. Given the optimistic projections for 1935, our theory anticipates the favorable shift in 1934 has little effect on the likelihood of Germany initiating a conflict in that year.

Specifications

The unit of analysis for our statistical tests is the directed rivalrous dyad year. A directed dyad approach is necessary as the theory specifies which actor initiates conflict when the military balance exceeds the cutpoint. The specification uses the Thompson (2001) rivalry coding because it does *not* rely on prior militarized conflicts to code rivalry presence. The concern with alternative rivalry codings is that given the empirical relationship between prior conflict and future conflict, prior conflict would affect both sample qualification and the outcome variable. The sample consists of roughly 9,200 observations. Using an alternative unit of analysis, such as the directed politically relevant dyad year, which yields a sample with over 140,000 observations, produces substantively and statistically similar results.³⁰

War initiation is a dichotomous outcome variable equal to one during the first year a state either originates or joins a war. Following Bell and Johnson (2015), joiners are only coded as such if they join the side of the originating initiator of the war. Those joining the side of the originally targeted state are coded as targets. Just over 1% of all rivalrous dyad years fit this criteria. We rely on the interstate war data of Sarkees and Wayman (2010). Given the binary outcome variable, all models in the manuscript use logistic regression with standard errors clustered on the directed dyad. Results are robust to using rare event logit models (King and Zeng, 2001) and employing cluster-robust standard errors (Aronow, Samii and Assenova, 2015).

Power shock is the primary explanatory variable. As operationalized above, it is a binary indicator equal to one when a dyad's realized capabilities in period t exceed the cutpoint. Recall, the cutpoint is set at the upper bound of the 95% confidence interval for that year's estimated balance. To differentiate between highly favorable realizations, as indicated by a power shock, and favorable but less extreme realizations, we also code a binary indicator equal to one when a dyad's balance in t is greater than the mean estimate but below the cutpoint. We do not anticipate that such minor favorable realizations change the probability of war. To test whether broader trends condition the effect of power shocks, we use the previously described future trend measure which captures the expected change in the balance of capabilities between t and $t + 1$. The variable takes on negative values when trends favor the initiating state. We interact the future trend and

³⁰Dyads are politically relevant when they include at least one major power or the constituent states share a border.

power shock measures to test the conditional relationship between favorable realizations in a given period and the future outlook.

The models control for an array of variables that correlate with war initiation according to prior work. These include the dyadic balance of capabilities, reflecting the initiator's share of dyadic capabilities where capabilities are measured using all six components of CINC scores. Past studies show that a more balanced dyadic distribution of power increases the risk of conflict (Reed, 2003). Given that CINC score components are also used to construct the power shock variable, we exclude the relative capabilities control variable from robustness test specifications and find similar results. War is more likely when states are contiguous (Bennett and Stam, 2004; Starr and Thomas, 2005). Accordingly, contiguity indicates whether states share a common land border. We control for joint dyadic democracy as democratic states are less likely to fight each other. Jointly democratic states each have *polity2* scores of six or greater (Marshall and Jaggers, 2002). To account for similarity of foreign policy interests, we include the Signorino and Ritter (1999) S-score. The measure reflects the similarity of two states' alliance portfolios with values close to one indicating strong similarity and values close to negative one indicating divergence. Finally, we include a measure of the duration of peace since the prior dyadic dispute and this term's squared and cubic values to account for temporal dependence in the time series cross sectional data (Carter and Signorino, 2010).

Results: Shocks, War, and Restraint

This section reports findings assessing each of the four empirical implications. We first show that power shocks increase the probability of war. Next, the results illustrate that not all favorable realizations increase the probability of war, only those of a sufficient scale to constitute shocks. Third, we demonstrate that expectations regarding the future trend in the dyadic balance of power condition the effect of power shocks. Finally, power shocks increase war likelihood even conditional on the size of the year-over-year shift in the power balance.

Model 1 of Table 1 reports results indicating that power shocks increase the risk of war. As shown, there is a positive and statistically significant relationship. A highly favorable realization of military capabilities increases the probability that state initiates or joins the initiating side in a war. Even within a sample of rivalrous dyads which have a higher baseline rate of war, greater than expected military capabilities substantially increase the risk of war. This is consistent with theories of windows of opportunity in which states are incentivized to attack before the capability balance reverts to a less favorable level. To facilitate interpretation, we simulate substantive quantities of interest using *Clarify* (King, Tomz and Wittenberg, 2000). With all covariates set to their median values, a power shock increases the probability of war from 0.5% by 0.5% points, $\pm 0.4%$ at the 95% confidence bounds. That is, a shock doubles the probability of war. An illustrative example below provides further intuition for the result.

Turning to Model 2 and the second testable implication, the results show that not all favorable power

Table 1: Shocks, Trends, and War: Rivalrous Dyads

	(1)	(2)	(3)	(4)	(5)
Power Shock	0.63** (0.28)		0.54* (0.29)	0.21 (0.62)	0.79*** (0.26)
Small Positive Realization		0.48 (0.29)			
Future Trend			-2.22 (1.56)		
Power Shock*Trend			5.44** (2.21)		
Annual Power Shift					-2.35* (1.32)
Relative Capabilities	-0.67 (0.48)	-1.00 (0.63)	-0.62 (0.49)	0.39 (1.17)	-0.65 (0.49)
Contiguity	-0.14 (0.30)	-0.13 (0.46)	-0.10 (0.31)	-0.68 (0.63)	-0.11 (0.30)
Foreign Policy Similarity	-0.53 (0.56)	-0.37 (0.72)	-0.58 (0.55)	1.13 (1.10)	-0.53 (0.57)
Peace Years	-0.35*** (0.07)	-0.53*** (0.12)	-0.36*** (0.08)	-0.38 (0.26)	-0.35*** (0.07)
Constant	-3.00*** (0.51)	-3.06*** (0.56)	-3.02*** (0.52)	-4.47*** (0.92)	-3.08*** (0.52)
N	9,321	7,708	9,290	2,160	9,321

* $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$

Notes: Logistic regression with directed rivalry dyad year as unit of analysis; standard errors clustered on the directed dyad. Models 1,3, and 5 include the full sample of observations. Model 2 drops observations with power shocks. Model 4 includes only those observations with *Future Trend* values in the most favorable quartile. Higher order terms for *Peace Years* are not shown. *Joint Democracy* drops out of the models due to collinearity with the outcome.

realizations increase the probability of war. This specification excludes observations experiencing power shocks and analyzes the effect of a potential initiator having a favorable power realization of a modest magnitude. That is, the dichotomous small positive realization is coded one if the initiator's share of dyadic capabilities in year t exceeds the mean predicted value but is less than the cutpoint. As shown, a small positive realization has little effect on war likelihood. Being slightly stronger than expected does not associate with more war initiation than being weaker than expected. This is consistent with the theoretical model, and with existing models of commitment problems, in that state's only initiate conflict when they anticipate a sizable unfavorable change in the distribution of capabilities. An expected minor reversion back from a small positive realization does not preclude successful bargains.

Model 3 of Table 1 incorporates estimated future trends into the analysis to test the third empirical implication. In the model, structural breaks condition the effect of power shocks. States that recognize a break occurred and thus anticipate remaining strong—that is, have favorable estimated future trends—do not view a power shock as a window of opportunity in which to attack but rather as indicative of the capability balance that will prevail for some time. Results from the interactive model strongly support the hypothesis. A negative outlook for future trends (positive values in the future trend variable) coupled with a positive shock is especially prone to war. In contrast, when future trends are favorable to the initiator, a positive shock has little effect on war likelihood. If the state favored by the power shock does not anticipate mean reversion in the power balance the next year, then the shock does not generate a strong incentive for attack.

To fix this point, Model 4 subsets the sample to only include those observations with the most favorable future trends expectations. Only those in the top quartile of favorable outlooks are included.³¹ Results show that under such conditions, a power shock has a substantively limited and statistically insignificant effect on war initiation.



Figure 4: Power shocks and the probability of war.

Consider the Iraq-Iran directed dyad from 1975 through 1980, depicted in Figure 4, to illustrate the logic. Through 1977, before sizable anti-Shah protests in Iran, there are no shocks in the dyadic balance as compared to expectations. Accordingly, there is a relatively low predicted probability of war (1.8% in 1977 based on Model 1, as shown in the right panel). Starting in 1978, there are shocks each year as Iran’s capabilities dropped, relative to Iraqi capabilities, amidst domestic turmoil and revolution. From the Iraqi perspective, its realized dyadic capabilities exceeded the upper bound of the estimated dyadic balance

³¹The quartile cut off is -2% for the future trend variable.

as shown in the left panel. Iraq experienced positive shocks in 1978 and 1979. However, a favorable future outlook—that is, a reduced probability of quick mean reversion in the dyadic balance—tempered the incentive to attack. The dynamic changed in 1980. Iran’s capabilities again fell short versus expectations but the future outlook was no longer as favorable for Iraq. The positive shock absent a favorable outlook created a window of opportunity in which to attack. As shown in the right panel of Figure 4, the predicted probability of war increases to 3.1% when all variables are set to their 1980 values.³² A power shock absent a favorable outlook increased the probability of war by over 70% relative to a year with no power shock. This elevated risk of war accords with the historical record as Iraq initiated the Iran-Iraq War seeking to exploit Iran’s temporary weakness (Hiro, 1991; Weisiger, 2013; Shirkey, 2016). Shocks to the dyadic power balance in this period were largely exogenous to the states’ rivalry. Iranian military spending declined roughly 40% in 1979 and another 40% in 1980 due to domestic instability.

To further ease interpretation, the left panel of Figure 5 plots the marginal effect of a positive shock on the probability of war across a range of salient values for the future trends variable. When the future appears bleak, as depicted on the right side of the figure, power shocks present an enticing window of opportunity to initiate conflict. Specifically, going from no power shock to a power shock increases the probability of war by 1.4% points, versus a baseline rate of 1.2%. A power shock when future trends appear pessimistic more than doubles the probability of war. As the future, from the initiator’s view, becomes relatively more attractive, power shocks have a diminishing effect on the probability of war initiation. When there is a neutral outlook, power shocks have a substantively smaller but still statistically significant effect. Finally, when the initiator anticipates a favorable distribution of capabilities will extend into the future, depicted by the left side of the figure, a fortuitous realization today does not induce an intractable commitment problem. The effect is further diminished and no longer precludes the null hypothesis.

Model 5 addresses the fourth implication. Unlike in conventional models of shifting power, the absolute magnitude of a power shift versus the prior year’s balance is not the pivotal element for determining whether war occurs. Rather, it is the value of the new balance versus states’ expectations for and their uncertainty around the power balance. To assess this contention we construct a variable for the annual power shift which captures the year-over-year change in the power balance. For instance, the variable equals 0.04 when the power balance shifts from the potential initiator holding 55% of dyadic capabilities in $t - 1$ to holding 59% of capabilities in t . Larger annual shifts are associated with positive shocks but there are substantively important differences between the two (correlation coefficient of 0.29). When conditioning for the annual shift, power shocks are still strongly associated with more war. In addition to the results in Model 5, we further illustrate this in the right panel of Figure 5. This plots favorable year-over-year changes in the power balance, broken down into 10% point increments. Within each increment, the plot shows the probability of war depending on whether the shift was a power shock—it exceeded the upper bound of expectations—or not. For all increments, excluding the most extreme one with limited observations, a power shock makes a

³²To make for a fairer comparison, we keep *Peace Years* set to the 1977 level. Predicted probability is based on Model 1.

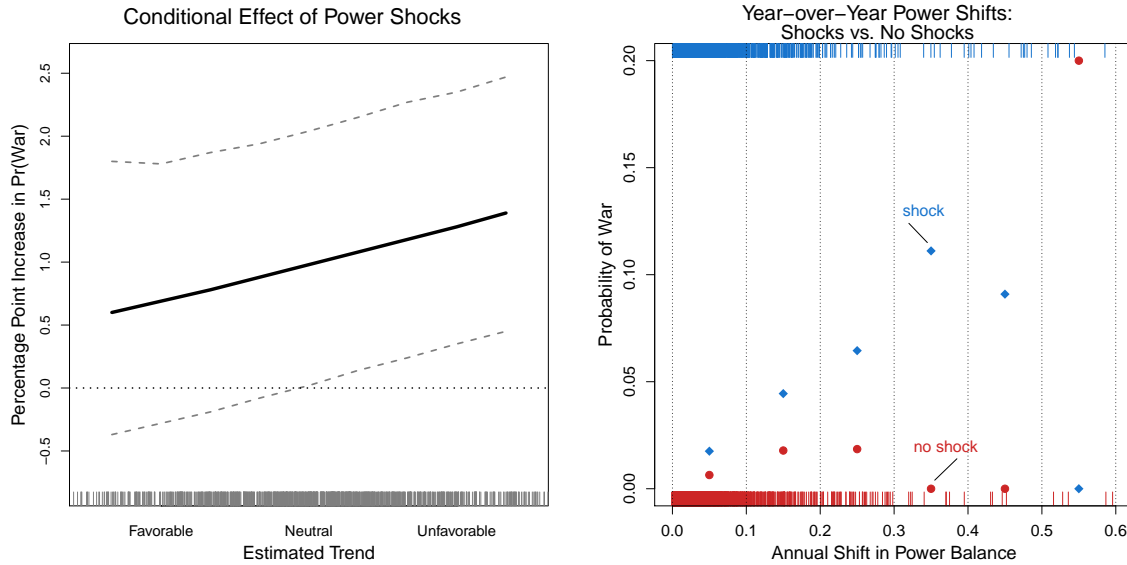


Figure 5: Left panel: Estimated future trends condition the marginal effect of power shocks on the probability of war. Rug plot shows distribution of the data on the moderating variable. Right panel: Year-over-year shift in the dyadic power balance, split by whether the shift constituted a power shock versus expectations. For a given size of power shift, those that exceed expectations (shock observations) have a higher probability of war than those that do not (no shock observations). Rug plots show distribution of annual shift for shocks (top) and no shocks (bottom).

shift of a given size more likely to generate war. Expectations and uncertainty affect whether a shift versus the prior year, even of quite substantial sizes, result in war.

In sum, the empirical results accord with the theoretical expectations. As in the theory, we have states estimate the future capability balance based on past realizations. When the current period's balance is highly favorable and unlikely to prevail in the future, commitment problems bind and war becomes more likely. This result attenuates and disappears though when the current period's balance is only slightly favorable—that is, below the cutpoint—and when the favorable balance is expected to endure in the future. In aggregate, power shocks endanger peace. However, there is important heterogeneity as not all power shocks do so. Shocks can be indicative of either a fleeting window of opportunity or a more structural change in the future distribution of capabilities. Uncertainty induced by the latter possibility fosters peace.

Robustness

To ensure that our sampling limitation to rivalrous dyads is not driving our results, we rerun all specifications using the directed politically relevant dyad year as the unit of analysis. As evident in Table 2, we find similar results across all models. Power shocks increase the probability of war (Model 6); small favorable power

realizations do not (Model 7); the effect of power shocks is conditional on the future balance trajectory (Models 8 and 9); and power shocks increase war likelihood given a fixed size shift in power from the prior year (Model 10).³³

Table 2: Shocks, Trends, and War: Politically Relevant Dyads

	(6)	(7)	(8)	(9)	(10)
Power Shock	0.85*** (0.18)		0.71*** (0.19)	-0.08 (0.54)	0.97*** (0.16)
Small Positive Realization		0.01 (0.18)			
Future Trend			0.73 (2.07)		
Power Shock*Trend			3.38 (2.15)		
Annual Power Shift					-3.81** (1.58)
Relative Capabilities	0.63*** (0.18)	0.48** (0.22)	0.61*** (0.18)	-0.94 (0.69)	0.65*** (0.18)
Contiguity	0.58*** (0.17)	0.62*** (0.19)	0.56*** (0.18)	-0.64 (0.46)	0.57*** (0.17)
Joint Democracy	-3.45*** (1.00)	-3.10*** (1.00)	-3.44*** (1.00)	—	-3.46*** (1.00)
Foreign Policy Similarity	-0.57** (0.29)	-0.52 (0.32)	-0.62** (0.29)	-0.23 (0.75)	-0.58** (0.29)
Peace Years	-0.12*** (0.02)	-0.12*** (0.02)	-0.12*** (0.02)	-0.20*** (0.06)	-0.12*** (0.02)
Constant	-5.54*** (0.29)	-5.47*** (0.32)	-5.49*** (0.29)	-3.39*** (0.62)	-5.58*** (0.29)
N	142,400	121,469	141,975	9,931	142,398

* $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$

Notes: Logistic regression with directed politically relevant dyad year as unit of analysis; standard errors clustered on the directed dyad. Models 6, 8, and 10 include the full sample of observations. Model 7 drops observations with power shocks. Model 9 includes only those observations with *Future Trend* values in the most favorable decile. Higher order terms for *Peace Years* are not shown. *Joint Democracy* drops out of Model 9 due to collinearity with the outcome.

The Supporting Files contain descriptive statistics and additional robustness tests. One robustness check excludes the relative capabilities control variable as its value depends on CINC scores which are also used

³³Though the interaction term in Model 8 falls short of conventional levels of statistical significance ($p = 0.12$), the coefficient points in the expected direction and the effect of substantive interest is further shown in Model 9.

to construct the explanatory variable. Due to the relative rarity of war initiation, we also check and find that results are substantively and statistically similar when using penalized, or rare events, logistic regression (Firth, 1993). Specifications without control variables produce similar results. To further corroborate that the future trend conditions the effect of power shocks, we subset the sample into three buckets based on the moderating variable's (future trend) value. Power shocks occur in each bucket (9% at a minimum and 23% at a maximum). Importantly, the association between power shocks and war becomes substantively and statistically stronger across the buckets as the future trend becomes more pessimistic. This suggests the interactive effect shown in the left panel of Figure 5 is not an artifact of model extrapolation. Additionally, all results are robust to accounting for complex dyadic dependencies with cluster-robust standard errors (Aronow, Samii and Assenova, 2015).

Confounding, or endogeneity, is a potential concern in the research design. In the theory, the balance parameter is determined via an exogenous stochastic process. A state having a favorable realization in a given period is unrelated to other elements of the strategic situation. This is plausible for many cases due to a variety of factors that affect military capacity but are partly exogenous to strategic considerations at the international level. These include the rate of economic growth, domestic instability of many forms, and technological breakthroughs. The drop in Iran's capacity before the Iran-Iraq War due to domestic turmoil is one example. Even with extreme exogenous events, leaders still retain some control over military capacity. We acknowledge that there is no automatic relationship between, say, economic decline and military capability. However, these exogenous events have implications for budgetary and feasibility constraints that are difficult for leaders to ignore. Reducing arms spending and making one's state temporarily vulnerable runs the risk of inviting attack. Yet that risk may be either expedient or necessary due to risks of domestic turmoil if the leader did not cut allocations to guns in an effort to continue providing butter. Thus, military spending may be driven by endogenous decisions but those decisions can be subject to powerful exogenous constraints.

We view endogeneity as a more pressing empirical problem, rather than a theoretical problem, and address it at length. The empirical challenge is that in some instances favorable power realizations generating power shocks may arise due to states' arming in preparation for attacking. If so, then the intention to attack, which is unobserved, is a common cause of both power shock and war initiation.

Noting the potential problem, we conduct three additional empirical tests and offer two theoretical considerations that mitigate the concern. The first robustness test accounts for endogenous arming. A power shock, as operationalized in specifications thus far, can occur because the initiator's capabilities exceed its upper bound expectation or the target's are below its lower bound expectation or some combination of the two. The first form is the most problematic from an endogeneity perspective. An initiator planning to attack can increase military capabilities, inducing a shock and a war. To address this concern we alter the power shock measure and make a different comparison. In this analysis the power shock variable only takes on a value of one when the target is weaker than expected. Endogeneity concerns are less apparent in this case

because a state is unlikely to *reduce* military capabilities when expecting an attack. We then compare the conflict proclivity of these observations (target is weaker than expected) against those where the target's capabilities are above the lower bound. For a fair comparison set the analysis excludes observations where the would-be initiator's capabilities exceed expectations—that is, those observations most subject to problematic endogenous arming.³⁴ The effect of a target state being weaker than expected remains substantively and statistically significant, both for rivalrous dyads and all politically relevant dyads. As shown in the Supporting Files, the results remain similar when using a power shock indicator that is less susceptible to inferential problems related to endogenous arming.

This approach of restricting the operationalization of power shocks to instances where the target's strength is below the lower bound expectation helps rule out another alternative explanation for the result. A typical asymmetric information account from bargaining models—e.g., Fearon 1995—could suggest that positive shocks are actually private information for the initiator rather than common knowledge as we maintain. If private information, then war could follow from the standard risk-reward tradeoff if the target believes they face a weak type when in fact the initiator is a strong type. However, this possibility no longer holds when restricting power shocks to cases where the target is weaker than expected. If this weakness is actually private information, then the target knows it is a weak type and accepts all pooling or screening offers from the initiator, which reduces war likelihood. However, the results are exactly the opposite. The probability of war increases when the target is weaker than expected. This is inconsistent with an asymmetric information account and consistent with our account focusing on stochastic power shifts and expectations of mean reversion.

A second concern is that the predictive model causes some states to be consistently coded as experiencing shocks. The Bell and Johnson (2015) fixed-effects approach, which we adopt, forces all states to have common coefficients on the predictors in a given year. If the relationship between an important predictor—e.g., economic potential—and future capabilities fundamentally differs for a state, the model will (over)underestimate its future capabilities. Consistent underestimation would cause us to indicate this state is experiencing positive shocks which could bias results. To guard against this possibility, we construct a control variable measuring the percentage of years in which a state's capabilities exceed the upper bound of its estimated capabilities (95% confidence interval). If the model severely underestimates a state's capability in 60% of years, this variable equals 0.6. All main results are substantively and statistically similar when controlling for this variable (see Supporting Files). The same is true when excluding states that experience shocks in a high percentage of years. Thus, states with military capabilities that frequently exceed the predicted levels—that is, are coded as having shocks that would not be shocking to more informed observers—are not driving the results.

A third robustness check analyzes the relationship between natural disasters, which are perhaps more

³⁴More precisely, the test drops observations if the initiator's capabilities exceed the upper bound expectation *and* the target's capabilities are not below the lower bound expectation.

analogous to true “shocks,” and conflict onset. Among rivalrous dyads since 1900, which is when the disaster data begins, there is a positive relationship between the number of individuals killed in a natural disaster and war onset. See the Supporting Files for results. Japan’s 1931 invasion of Manchuria provides an illustrative example. Flooding of the Yangtze created one of the deadliest natural disasters in history, with up to four million reported fatalities. Moreover, it suppressed Chinese economic output, stifled troop movements, and impaired national finances due to flood relief provisions (Jordan, 1987, p. 49). Among other reasons, Japanese field officers seized upon this window of vulnerability and initiated the incursion into Manchuria. Though supporting our contentions, we urge caution in interpreting the statistical result. Few natural disasters are of a sufficient scale to meaningfully impair a state’s military capacity. Regional crises are not equivalent to national military shocks. Additionally, there is no intuitive measure of whether a state expects its rival, struck by a disaster, to remain weak for an extended period. This prevents testing a number of our hypotheses.³⁵ With the proper caveats noted, natural disasters provide further evidence of a link between exogenous power shocks and conflict onset.

An additional empirical point concerning results from the interactive models incorporating future power trends further tempers endogeneity concerns. To the extent that power shocks are endogenous, those that occur alongside a favorable future trend are similarly endogenous in that the beneficiary opts to increase military capabilities. And yet we find no relationship between shocks and war initiation when coupled with a favorable future outlook. Consequently, these models are less susceptible to the confounding critique.

Two theoretical considerations further mitigate the endogeneity concerns in the main specifications. First, theoretically a state may preventively attack if it recognizes its adversary is endogenously investing in military capabilities. In such cases, the attacking state’s window of opportunity is fleeting but there is no power shock. This dynamic would only weaken our results.

A second theoretical consideration may further assuage endogeneity concerns. Imagine a dyad with escalating tensions that ultimately goes to war. The initiating state might build arms in anticipation of attacking, as noted. However, target states are rarely oblivious to this threat and so, in turn, build their own arms. Because our main measure of a power shock is dyadic—that is, it accounts for the arming choices of the initiator and the target—the competitive arming will frequently cancel out and thus produce no power shock by our measure. In this case, an intention to attack is still associated with war but is no longer also a cause of power shocks. Thus, if this dynamic exists, it would likely attenuate the effect of power shocks that we estimate. Intention to attack will only be a common cause of both war and power shocks if the target fails to anticipate an attack is coming.

Arms races provide the classic example of highly endogenous arming choices coupled with an elevated base rate of war. It would be problematic if the power shock indicator were merely capturing arms races.

³⁵Limitations in the disaster data further encourage caution. The existing data does not provide dates, opening the possibility that some wars began before the disaster. However, it is difficult (though not impossible) to envision wars causing deadly natural disasters.

However, this is not the case, likely for the strategic reasons noted above. We find a slightly *negative* correlation (-0.07) between the power shock measure and an indicator for arms races (Rider, Findley and Diehl, 2011). Confounding concerns dissipate if targets recognize initiators' threats and arm themselves in preparation. Such cases generate arms races, but not power shocks. However, if arming is a lumpy process where states' arming choices do not inherently cancel out one another, then arms races could create power shocks and windows of opportunity in which to attack (Morrow, 1989*b*).

Conclusion

Commitment problems due to shifting power are a primary reason states fight costly wars. In this paper, we qualified prior theoretical results and showed that sudden power shifts can be less problematic than previously theorized, and not all power shifts should be equally dangerous. To better approximate the estimative challenges leaders confront we relaxed the common assumption in the existing literature that states know with certainty the magnitude and direction of power shifts, or the processes that generate these shifts. Instead, we allowed states to approximate future power balances and the distribution generating them by observing past and present power trends. Incorporating this type of fundamental uncertainty in a model of conflict reveals that power shifts are less likely to create commitment problems and war than previously argued.

The empirical literature on power shocks and commitment problems is limited. Our findings establish that shocks do indeed increase the risk of interstate war. However, there are important qualifications to this result which are consistent with our theory. Uncertainty about the future power balance can reduce the likelihood of conflict. In particular, when states project a favorable future trend in the power balance, akin to a structural break, a power shock has little effect on war likelihood. Moreover, not all large year-over-year changes in the balance are equally problematic. Those that occur alongside greater uncertainty are more prone to peace because the advantaged state is less concerned that there will be rapid mean reversion in the power balance.

Subsequent studies could take multiple directions to build upon this piece. Future work could profitably disentangle the various sources and forms of power shocks. This paper focuses on realizations of the military power balance that diverge from expectations. Others may delve into the sources of these surprising realizations. For instance, shocks due to domestic political turmoil versus economic crises versus rapid armament programs could have heterogeneous effects. Additionally, the substantive metric used to establish power shocks might have important implications. We focus on realized military capabilities. Shifts in latent capabilities, for instance due to trends in population demographics or long run economic growth, might be less prone to conflict if states can credibly commit not to translate latent capabilities into realized capabilities through arms agreements. As shown in this study, power shocks are a source of conflict in aggregate but there is important heterogeneity as to whether and when shocks generate intractable commitment problems.

References

- Acemoglu, Daron and James A. Robinson. 2000. "Why Did the West Extend the Franchise? Democracy, Inequality, and Growth in Historical Perspective." *The Quarterly Journal of Economics* 115(4):1167–1199.
- Aronow, Peter M., Cyrus Samii and Valentina A. Assenova. 2015. "Cluster-Robust Variance Estimation for Dyadic Data." *Political Analysis* 23(4):564–577.
- Bai, Jushan. 1994. "Least squares estimation of a shift in linear processes." *Journal of Time Series Analysis* 15(5):453–472.
- Bas, Muhammet A. and Andrew J. Coe. 2012. "Arms Diffusion and War." *Journal of Conflict Resolution* 56(4):651–674.
- Bas, Muhammet A and Andrew J Coe. 2016. "A Dynamic Theory of Nuclear Proliferation and Preventive War." *International Organization* 70(4):655.
- Bas, Muhammet and Robert Schub. 2016. "Mutual Optimism as a Cause of Conflict: Secret Alliances and Conflict Onset." *International Studies Quarterly* 60(3):552–564.
- Bell, Sam R. and Jesse C. Johnson. 2015. "Shifting Power, Commitment Problems, and Preventive War." *International Studies Quarterly* 59(1):124–132.
- Bennett, D. Scott and Allan C. Stam. 2004. *The Behavioral Origins of War*. Ann Arbor, MI: University of Michigan Press.
- Blainey, Geoffrey. 1988. *The Causes of War*. London, Macmillan.
- Carter, David B. and Curtis S. Signorino. 2010. "Back to the Future: Modeling Time Dependence in Binary Data." *Political Analysis* 18:271–292.
- Chadefaux, Thomas. 2011. "Bargaining over Power: When Do Shifts in Power Lead to War?" *International Theory* 3(2).
- Chapman, Terrence L., Patrick J. McDonald and Scott Moser. 2015. "The Domestic Politics of Strategic Retrenchment, Power Shifts, and Preventive War." *International Studies Quarterly* 59(2):133–144.
- Clark, David H., Benjamin O. Fordham and Timothy Nordstrom. 2011. "Preying on the Misfortune of Others: When Do States Exploit Their Opponents' Domestic Troubles?" *Journal of Politics* 73(1):248–264.
- Debs, Alexandre and Nuno P. Monteiro. 2014. "Known Unknowns: Power Shifts, Uncertainty, and War." *International Organization* 68(1).
- Fearon, James D. 1995. "Rationalist Explanations for War." *International Organization* 49(3):379–414.
- Fearon, James D. 2004. "Why Do Some Civil Wars Last So Much Longer than Others?" *Journal of Peace Research* 41(3):275–301.
- Firth, David. 1993. "Bias Reduction of Maximum Likelihood Estimates." *Biometrika* 80(1):27–38.

- Fordham, Benjamin O. 2011. "Who Wants to be a Major Power? Explaining the Expansion of Foreign Policy Ambition." *Journal of Peace Research* 48(5):587–603.
- Fudenberg, Drew and Jean Tirole. 1991. *Game Theory*. Cambridge, MA: The MIT Press.
- Garfinkel, Michelle R. 1994. "Domestic Politics and International Conflict." *American Economic Review* 84(5):1294–1309.
- Gilpin, Robert. 1981. *War and Change in World Politics*. Cambridge: Cambridge University Press.
- Gleditsch, Kristian S. and Andrea Ruggeri. 2010. "Political Opportunity Structures, Democracy, and Civil War." *Journal of Peace Research* 47(3):299–3010.
- Goddard, Stacie E. 2015. "The Rhetoric of Appeasement: Hitler's Legitimation and British Foreign Policy, 1938-1939." *Security Studies* 24(1):95–130.
- Hiro, Dilip. 1991. *The Longest War: The Iran-Iraq Military Conflict*. New York: Routledge.
- Houweling, Henk and Jan G. Siccama. 1988. "Power Transitions as a Cause of War." *Journal of Conflict Resolution* 32(1):87–102.
- Jordan, Donald A. 1987. "The Place of Chinese Disunity in Japanese Army Strategy during 1931." *The China Quarterly* (109):42–63.
- Kim, Nam Kyu. 2016. "Revisiting Economic Shocks and Coups." *Journal of Conflict Resolution* 60(1):3–31.
- Kim, Woosang and James D. Morrow. 1992. "When Do Power Shifts Lead to War?" *American Journal of Political Science* 36(4):896–922.
- King, Gary and Langche Zeng. 2001. "Explaining Rare Events in International Relations." *International Organization* 55(3):693–715.
- King, Gary, Michael Tomz and Jason Wittenberg. 2000. "Making the Most of Statistical Analyses: Improving Interpretation and Presentation." *American Journal of Political Science* 44(2):347–61.
- Krainin, Colin. forthcoming. "Preventive War as a Result of Long Term Shifts in Power." *Political Science Research and Methods* .
- Krainin, Colin and Thomas Wiseman. forthcoming. "War and Stability in Dynamic International Systems." *Journal of Politics* .
- Leeds, Brett Ashley, Jeffrey M. Ritter, Sara McLaughlin Mitchell and Andrew G. Long. 2002. "Alliance Treaty Obligations and Provisions." *International Interactions* 28(2):237–260.
- Levy, Jack S. 1987. "Declining Power and the Preventive Motivation for War." *World Politics* 40(1):82–107.
- Marshall, Monty G. and Keith Jagers. 2002. *Polity IV Dataset [Computer file]*. College Park, MD: University of Maryland.
- Morrow, James D. 1989a. "Capabilities, Uncertainty, and Resolve: A Limited Information Model of Crisis Bargaining." *American Journal of Political Science* 33(4):941–972.

- Morrow, James D. 1989b. "A Twist of Truth: A Reexamination of the Effects of Arms Races on the Occurrence of War." *Journal of Conflict Resolution* 33(3):500–529.
- Nielsen, Richard A., Michael G. Findlay, Zachary S. Davis, Tara Candland and Daniel L. Nielsen. 2011. "Foreign Aid Shocks as a Cause of Violent Armed Conflict." *American Journal of Political Science* 55(2):219–232.
- Organski, A. 1967. *World Politics*. 2 ed. New York: Knopf.
- Organski, A. and Jacek Kugler. 1980. *The War Ledger*. Chicago, IL: University of Chicago Press.
- Powell, Robert. 1999. *In the Shadow of Power: States and Strategies in International Politics*. Princeton, N.J.: Princeton University Press.
- Powell, Robert. 2004. "The Inefficient Use of Power: Costly Conflict with Complete Information." *American Political Science Review* 98(2):231–241.
- Powell, Robert. 2006. "War as a Commitment Problem." *International Organization* 60(1):169–203.
- Reed, William. 2003. "Information, Power, and War." *American Political Science Review* 97(4):633–41.
- Rider, Toby J., Michael G. Findley and Paul F. Diehl. 2011. "Just Part of the Game? Arms Races, Rivalry, and War." *Journal of Peace Research* 48(1):85–100.
- Sarkees, Meredith R. and Frank Wayman. 2010. *Resort to War: 1816-2007*. Washington DC: Congressional Quarterly Press.
- Shannon, Megan, Daniel Morey and Frederick J. Boehmke. 2010. "The Influence of International Organizations on Militarized Dispute Initiation and Duration." *International Studies Quarterly* 54(4):1123–1141.
- Shirkey, Zachary C. 2016. "Uncertainty and War Duration." *International Studies Review* .
- Signorino, Curtis S. and Jeffrey M. Ritter. 1999. "Tau-b or not Tau-b: Measuring the Similarity of Foreign Policy Positions." *International Studies Quarterly* 43(1):115–144.
- Singer, J. David, Stuart Bremer and John Stuckey. 1972. Capability Distribution, Uncertainty, and Major Power War, 1820-1965. In *Peace, War, and Numbers*, ed. Bruce Russett. Sage: Beverly Hills.
- Slantchev, Branislav L. 2011. *Military Threats: The Costs of Coercion and the Price of Peace*. Cambridge: Cambridge University Press.
- Slantchev, Branislav L. and Ahmer Tarar. 2011. "Mutual Optimism as a Rationalist Explanation for War." *American Journal of Political Science* 55(1):135–148.
- Starr, Harvey and G. Dale Thomas. 2005. "The Nature of Borders and International Conflict: Revisiting Hypotheses on Territory." *International Studies Quarterly* 49(1):123–139.
- Thompson, William and David Dreyer. 2010. *Handbook of International Rivalries*. Washington DC: Congressional Quarterly Press.
- Thompson, William R. 2001. "Identifying Rivals and Rivalries in World Politics." *International Studies Quarterly* 45(4):557–86.

- Van Evera, Stephen. 1999. *Causes of War: Power and the Roots of Conflict*. Ithaca, NY: Cornell University Press.
- Wagner, R. Harrison. 1994. "Peace, War, and the Balance of Power." *The American Political Science Review* 88(3):593–607.
- Walter, Barbara F. 1997. "The Critical Barrier to Civil War Settlement." *International Organization* 51(3):335–364.
- Weisiger, Alex. 2013. *Logics of War: Sources of Limited and of Unlimited Conflicts*. Cornell University Press.
- Wolford, Scott, Dan Reiter and Clifford J Carrubba. 2011. "Information, commitment, and war." *Journal of Conflict Resolution* 55(4):556–579.