



International Interactions

Empirical and Theoretical Research in International Relations

ISSN: 0305-0629 (Print) 1547-7444 (Online) Journal homepage: <http://www.tandfonline.com/loi/gini20>

How International Reputation Matters: Revisiting Alliance Violations in Context

Brad L. LeVeck & Neil Narang

To cite this article: Brad L. LeVeck & Neil Narang (2016): How International Reputation Matters: Revisiting Alliance Violations in Context, *International Interactions*, DOI: [10.1080/03050629.2017.1237818](https://doi.org/10.1080/03050629.2017.1237818)

To link to this article: <http://dx.doi.org/10.1080/03050629.2017.1237818>



View supplementary material [↗](#)



Accepted author version posted online: 20 Sep 2016.
Published online: 20 Sep 2016.



Submit your article to this journal [↗](#)



Article views: 54



View related articles [↗](#)



View Crossmark data [↗](#)

How International Reputation Matters: Revisiting Alliance Violations in Context

Brad L. LeVeck^a and Neil Narang^b

^aUniversity of California, Merced; ^bUniversity of California, Santa Barbara

ABSTRACT

We investigate the role of international reputation in alliance politics by developing a signaling theory linking past alliance violations with the formation of future alliance commitments. In our theory, past violations are useful signals of future alliance reliability conditional on whether they effectively separate reliable from unreliable alliance partners. It follows that states evaluating potential alliance partners will interpret past violations in their context when deciding to enter a new alliance, attaching less weight to violations in “harder times,” when many states are defaulting on their alliance commitments together, and more weight to violations in “easier times,” when fewer states are defaulting on their alliances. We test our theory and find that states are empirically more likely to form new alliances with states that violated in harder times compared to states that violated in easier times. The results have important implications for how scholars understand and estimate the impact of international reputation.


KEYWORDS

Alliances; cooperation; foreign policy; international institutions; international organizations; international reputation; international security; statistics

Scholars of international relations continue to debate whether international reputation matters. In large part, this is because the empirical evidence to date has been largely mixed, both within domains and across domains of international politics. For example, in the realm of international security, most empirical studies of extended deterrence have found little to no relationship between states’ past willingness to follow through on military threats and the likelihood that a state will be challenged in the future (Hopf 1994; Huth and Russett 1984; Mercer 1996, 1997; Press 2004, 2005).¹ On the other hand, at least one study demonstrates that past actions can affect subsequent deterrence outcomes (Weisiger and Yarhi-Milo 2015). Similarly, in the domain of alliance politics, Gibler (2008), Crecenzi, Kathman, Kleinberg, and Wood (2012) and Miller (2012) demonstrate how states that violated

CONTACT Neil Narang  narangn@gmail.com  University of California, Santa Barbara, 3710 Ellison Hall, Santa Barbara, CA 93106, USA.

Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/gini

 Supplemental data for this article can be accessed on the [publisher's website](#)

¹For example, Press (2004:169) concludes that, “Credibility of threats and promises does not hinge on establishing a history of resolute actions. As the current calculus theory explains, threats and promises are only credible if—and only if—they are consistent with important interests and are backed by substantial power.”

their alliance commitments in the past appear less likely to find partners in the future, suggesting that international reputation may matter. Meanwhile, in the field of international political economy, an influential study by Tomz (2007) demonstrates a mechanism through which reputational concerns can affect investors' decision to lend, debtors' decision to repay, and the structure of sovereign loans.

In this article, we aim to further explore the role of international reputation in the context of bilateral security alliances. Bilateral alliances are an interesting forum to test the effect of international reputation for several reasons. First, there is less independent monitoring or enforcement of states' alliance commitments compared to other, more heavily institutionalized domains of international politics (for example, trade, nuclear cooperation, etc.). This suggests that reputation should serve as an important mechanism to ensure compliance in the absence of formal institutions (Hafner-Burton, LeVeck, and Victor [Forthcoming](#)). Second, most studies of reputation in international security have studied it in the context of ultimatum bargaining (Hopf 1994; Huth and Russett 1984; Mercer 1996, 1997; Press 2004, 2005), where states have numerous opportunities to send other costly signals to demonstrate resolve (audience costs, troop mobilizations, etc.) (Fearon 1997). These other signals might mask the effect of reputation and account for why analysts have found mixed support for a reputation effect. Alliance negotiations, by contrast, often afford states fewer opportunities to signal the credibility of their commitments save for past actions and treaty design (Morrow 2000). Finally, alliances, unlike deterrent threats, are explicitly written documents recording the nature of commitments over time. This allows analysts to more clearly determine when states fail to honor their commitments.

Following Tomz (2007), this article proposes a signaling theory of alliance politics in which states infer the reliability of potential alliance partners from observable information and then factor this into the decision to offer a new alliance. We effectively treat security alliances as international contracts where states have private information about their own willingness and ability to honor commitments in the future (reliability). Because states seeking alliances in the international system have private information about their reliability, one way they demonstrate the credibility of their security commitments to (potential) partners is by upholding their alliance commitments. Potential alliance partners, in turn, learn from this signal only when they believe it to be correlated with reliability.

Treating alliances in this way makes a simple but important contribution to our understanding of reputation in alliance politics. Importantly, we demonstrate that alliance violations do not always matter in the exact same way because individual violations do not necessarily signal the same level of unreliability. Specifically, we show that when there are systemic shocks in the international system (such as WWII in Europe or decolonization in Africa), it becomes too

costly for many different states to uphold their alliance obligations. Because these shocks lead so many states to abrogate their alliances together, violations under these conditions convey less information about individual states' reliability, and thus they harm a state's ability to form future alliances much less. Conversely, states that violate their alliances in regions and times where few other states are doing so appear much less likely to find new alliance partners in the future. In this way, our innovation is to show that past violations are useful signals of future alliance reliability conditional on whether they effectively separate reliable from unreliable alliance partners.

Alliance bonds and signaling reliability under uncertainty

International alliance agreements function like bonds in at least one important way: Just as bonds are issued based on the *expectation* that a recipient will eventually repay its debt, states in the international system must enter alliance contracts based only on the *expectation* that a security commitment will eventually be honored (Lake 1999; Leeds and Anac 2005; Morrow 2000).² The key time-inconsistency problem for both bond issuers and states issuing alliance commitments is that neither can directly observe a recipient's intentions to uphold its commitment in the future. Because potential alliance partners always have private information about their willingness and ability to honor their commitment in the future, states issuing alliance contracts always run the risk that potential alliance partners will default on their commitment when it is no longer in their interest to comply.³

Alliance politics thus resembles other contractual relationships in which there is asymmetric information with respect to quality (Akerloff 1970), where—in this case—quality can be understood as future alliance reliability. One potential consequence of the interaction between quality heterogeneity and incomplete information may be disappearance of Pareto-improving agreements altogether where guarantees are indefinite. With respect to international alliance agreements, this may produce suboptimal levels of alliance formation in cases where two or more states may otherwise benefit from cooperating.⁴

A well-known solution to the asymmetric information problem, originally proposed by Spence (1973), is for actors with private information to credibly

²There are, of course, many practical differences between lending and alliance contracts. For example, transactions in bond markets can be much more fluid than cooperation through international alliance contracts. Nevertheless, here we aim to highlight a time-inconsistency problem common to both, where Pareto-improving transactions may end because one side has asymmetric information about its own reliability.

³While stated in terms of issuers and recipients, uncertainty over the likelihood that an alliance will be honored in the future is reciprocal.

⁴One might suppose that a key difference between bonds and the exchange of alliance contracts is a power asymmetry between a borrower and lender that does not exist in the formation of alliances. This is not necessarily true, as cooperating on an alliance must be Pareto improving for cooperation to occur (Lake 1999; Morrow 2000). Just as a borrower has bargaining power in negotiating terms with a lender, potential alliance partners generally have bargaining power vis-à-vis another state seeking an ally for the joint production of security.

signal their type by taking costly actions that poor-quality candidates cannot efficiently mimic. With respect to international alliance agreements, where potential alliance partners have private information about their willingness and ability to honor alliance commitments in the future, states can calculate the likelihood that a potential partner is unreliable using a variety of observable indicators (costly signals) they believe to be correlated with states' underlying reliability—the latent parameter of interest. By combining factors from a potential alliance partner's past alliance behavior and current characteristics, states can update their beliefs about the reliability of potential partners within the current systemic conditions to determine whether or not to issue an alliance contract and with what terms. We treat past decisions to honor an alliance as a costly signal of “type” in much the same way that Spence originally conceived of taking a costly action, like investing in education, *prior* to entering the labor market in anticipation of bargaining over wages. This may seem different from costly signals taken *during* the bargaining process to convey information. However, the mechanism is theoretically identical: States are taking a costly action (in our case, honoring an alliance) at time t in order to impact beliefs about type (in our case reliability) at time $t+n$. The difference appears to be only a practical one related to the actual amount of time between signaling and offers/counteroffers.

When understood this way, much of the existing research on alliances formation appears to fit within these broad analytic categories. *Orthodox* theories of alliance formation have generally focused on systemic conditions that affect the overall demand for alliances. In this view, alliances are understood as attempts by states to “balance” against the capabilities (Waltz 1959) or threats (Walt 1987) of other coalitions by aggregating military capabilities among members. Still other analysts emphasize the role of “bandwagoning” in alliance formation, where nations ally with more capable states in order to enhance their security (Schweller 1998). Within this broader calculus, Altfeld (1984) and Morrow (1991) attempt to explain which states ultimately ally by characterizing a trade-off between security and autonomy. In this view, militarily strong states provide greater security to weaker partners in return for greater policy concessions. Thus, it is diversity in capabilities that often drives alliance formation and helps explain the prevalence and durability (Conybeare 1992, 1994) of asymmetric alliances.⁵

Taken together, much of the traditional research on alliance formation can be usefully understood as characterizing various structural conditions that influence the overall demand for alliances in the international system. Just as borrowers increase their demand for loans in periods of greater financial insecurity, states in the international system appear to increase their demand for alliances in periods of greater national insecurity when the value of

⁵See Gibler and Rider (2004) on how capabilities and interests interact in the autonomy-security trade-off.

securing an alliance increases. And, similar to how the frequency of lending is largely determined by the ratio of capital-rich to capital-poor actors in a system, the rate of alliance formation can be partly driven by the ratio of militarily powerful states to less military powerful states in the population, the latter of which appear to trade autonomy (similar to the logic of comparative advantage).

In contrast to traditional approaches, *contemporary* alliance research has tended to focus more on unit-level factors to explain alliance formation and durability. A primary focus has been on studying the relationship between regime type and alliances. To date, the empirical results have been mixed. With respect to alliance formation, several studies find that similar regime types—specifically democracies—are more likely to ally (Lai and Reiter 2000; Siverson and Emmons 1996), while at least one study finds that differing regimes are more likely to ally (Simon and Gartzke 1996).⁶ With respect to alliance durability, some studies find that democracies make more reliable allies (Leeds 2003a; Leeds, Mattes, and Vogel 2009) while others find that democracies make less-reliable allies (Gartzke and Gleditsch 2004; Leeds and Gigliotti-Labay 2003; Smith 1996). The first set of results is more consistent with the finding that democracies tend to honor their international obligations and are perceived to be more credible by international actors (Leeds 2003a, Choi 2003; Mansfield, Milner, and Rosendorff 2002; Schultz 1999).

On the whole, these results suggest that certain characteristics of states can function as signals or indices of alliance reliability that are difficult to manipulate. If bilateral security alliances can be usefully understood as international contracts where states have private information about their willingness and ability to honor commitments in the future, then just as certain social, political, or economic indicators can signal the reliability of a borrower to a lender, an important mechanism that allows alliance commitments to be exchanged where guarantees are indefinite is for states to take costly actions to signal reliability under uncertainty.

In addition to system- and unit-level factors, states seeking an alliance partner in the international system may also screen potential alliance partners based on past actions. In fact, this is often necessary because important attributes like credibility are not directly observable (Guzmán 2008). To assess these qualities, potential partners observe past choices that correlate with the ability to endure or with weakness in the face of adversity.

Given the importance of reputation in other types of contractual relationships, it is odd that more attention has not been paid to reputation and how it matters in alliance politics. Two notable exceptions are Gibler (2008) and Crescenzi et al. (2012), who test whether violating the terms of a bilateral alliance decreases the

⁶Gartzke and Weiseger (2013) argue that the relationship between regime type and alliance formation is mediated by the system-level prevalence of a particular regime.

chance that a dyad will form a future alliance. However, our view of reputation differs in a number of important ways. Most important is the theoretical definition of reputation itself. In our view, reputation is not the actions taken by states, but the *beliefs* that those actions create. This is important because it makes clear that one cannot capture the concept of reputation by simply modeling a state's violation history; one must also model the context in which those violations occur. As we explain in the next section, our innovation is to draw on signaling theory to show how past alliance violations are not uniformly informative as signals of future alliance reliability. Rather, context has important effects on how and when an action will cause others to update their beliefs about an actor and hence how others will treat that action in the future.

Theory: Signaling alliance reliability with past actions

In this section we describe how the decision to honor or violate an alliance is a costly signal that provides information about a state's underlying propensity to honor its agreements. We do this to show how the context of an alliance violation matters and to derive a testable hypothesis about the conditions under which violating an alliance affects the probability of gaining a future alliance. Our logic is similar to that developed by Tomz (2007), who showed that defaults on sovereign debt during periods of high regionwide default did little to harm countries' reputation for repayment because the widespread violations made it difficult for lenders to distinguish normally reliable states from normally unreliable states based on this signal. Here, we apply this theory to study reputation in military alliances for the first time. In doing so, we demonstrate that a signaling theory of international reputation also provides a clear and testable hypothesis about which alliance violations are actually likely to harm the reputation of states and which violations will be relatively inconsequential.

Consider two states that are contemplating entering into an alliance contract. Presumably, each state stands to benefit from the alliance as long as each member fulfills its duties, either through deterring a potential challenger (Smith 1995) or through burden sharing in the event of an actual conflict (Lake 1999). However, because alliance commitments are ultimately costly to carry out, only some states will be willing to actually honor the agreement and eventually fight on an ally's behalf, even if doing so gives them better access to future alliances. Assuming states prefer to form an alliance with a more-reliable partner to less-reliable partner, the problem becomes how one state might discern the other's latent propensity to abide by their agreement.⁷ One possible solution is to rely on observable indicators

⁷We do not explore the actual determinants of a state's reliability. Like previous literature on reputation, we assume that part of what causes states to honor or violate their commitments is persistent across time and partially expressed through past behavior. However, in the results, we also explore the fact that some elements behind a state's tendency to honor commitments can change over time or by leader, causing the informational value of violations to decay. Doing so does not impact our results.

that may be correlated with a state's underlying reliability, like domestic political institutions or culture, as a visible—albeit imperfect—signal of quality. In Jervis' terms, these would be indices rather than signals (Jervis 1989).⁸

In this context, a particularly good indicator of future alliance behavior may be past behavior. If a state violated its agreements in the past, it seems intuitive that it may be more likely to do so in the future. However, Spence (1973) famously showed that past behavior is not always equally informative and that whether past behavior distinguishes one type from another depends crucially on the behavior's cost. If, for instance, honoring an alliance becomes so difficult that all states are forced to violate their commitments together, then a violation conveys little information about how reliable one state is relative to another. Beyond this extreme example, the general insight is that alliance violations do more to signal that a state is relatively unreliable when many other states appear to be willing and able to honor the same agreement.

Of course, whether other states would honor a particular agreement under similar conditions is often difficult to observe (Narang 2014; Narang and Mehta 2015), as each alliance has elements that are somewhat unique. However, there may be times and regions where system-level shocks cause a large number of countries to simultaneously violate alliance commitments together. This may provide relatively clear evidence to a potential partner that the costs of honoring a previous alliance were so great that even reliable states that would normally honor their commitment were unable to do so.

This discussion has important implications for empirically studying how violating an alliance affects a state's reputation. It is likely that the cost of maintaining an alliance varies significantly by region and time and that one can identify shocks across these dimensions. Figure 1, which plots the percentage of states violating their bilateral security alliance in each region and year based on Leeds et al. (2009), supports this supposition.

In periods of extreme war (such as World War II in Europe), or in periods where regional politics are in great upheaval (such as the period of rapid decolonization in post-1950s Africa), the regional violation rate among states is much higher.⁹ For the purposes of our theory, these regionally and temporally

⁸Jervis (1989:18) draws a distinction between signals that “both the sender and perceiver realize ... can be as easily issued by the deceiver as by an honest actor” and indices, which “are beyond the ability of the actor to control for the purpose of projecting a misleading image.” However, in signaling theory, *costly* signals are by definition statements or actions where costs are asymmetric across type and thus positively correlated with the latent parameter of interest such that they can separate one type from another (Banks 1991). In other words, the concept of “costly signals,” used broadly across the social and life sciences, is the closer theoretical analog to Jervis' concept of indices rather than his concept of signals.

⁹The wave of defaults in Europe occurs due to World War II: 8.2% of European states defaulted in 1939, rising to 10.9% in 1940, and peaking at 11.1% at the height of the War in 1941. Similarly, the wave of defaults in Africa is the product of the decolonization in the 1960s, where once relatively well-ordered and peaceful territories transitioned into violent and inefficient independent states. Alliance violations in Africa peak at 16.6% in 1960—at the height of decolonization, when 17 states declared independence—and again in 1962. Meanwhile, the wave of alliance defaults in the Middle East is well documented by Walt (1987), who surveys the rapidly “changing internal and external circumstances” that caused these violations, and influenced subsequent alliance formation in the post-war period.




Figure 1. Variation in bilateral violation rate by region over time.

clustered moments represent system-level shocks along a set of dimensions where even fairly reliable states can be expected to default on their alliances. Because violations in these periods will do less to distinguish a state as unreliable (as both reliable and unreliable states appear to pool on violating), one should expect the reputational consequences to be lower. We therefore exploit structural shocks in the regional and temporal violation rates to assess the main hypothesis.

H1: States are less likely to enter new alliances with states that violated in periods of low regional violation (when the violation more clearly separates unreliable allies from reliable allies) than they are with states that violated in periods of high regional violation.

To be clear, we do not assert through the hypothesis that a state's direct experience with a potential alliance partner does not matter, nor do we assert that the only "correct" cohort in which violations are interpreted is the group of states in its region. Rather, we propose to exploit clear regional and temporal patterns of alliance violations as one valid proxy for the varying informational value that a single violation can have for signaling reliability to future partners—our key theoretical innovation on previous studies. As discussed in the next section, there need not be anything essential to regions for them to be empirically useful in estimating the reliability of a potential ally. Rather, the underlying events that cause these patterns (spatially correlated wars, economic crises, famines, state failure, etc.) need only provide an observable moment in which some states clearly

separate from others. One could potentially exploit other behavioral patterns that emerge from additional dimensions (regime type, economic, etc.), but doing so is not crucial to testing our hypotheses.

There is, however, compelling qualitative evidence that states have interpreted past violations by potential partners within their regional context. For example, consider Russia's commitment to the Triple Entente between the Bosnian Crisis and the Agadir Crisis, and Austria's commitment to Germany during the first Moroccan Crisis as two separate cases that illustrate the hypothesized mechanism. According to Mercer (1996), the tepid Russian support for the Triple Entente during the Bosnian Crisis did not lead its British and French allies to later perceive it as unreliable during the 1909 Agadir Crisis because its past actions were explained away in situational terms (Mercer 1996:179). In particular, and consistent with the theory, Mercer quotes Lloyd George's explanation for discounting prior instances in which Russia was irresolute by specifically referencing the forgivable strain placed on Russia due to states in the region having collectively experienced "as much war as they can stand" (Mercer 1996:179).

Similarly, during the First Moroccan Crisis in 1906, the Germans knew their Austrian ally would prove unreliable, but they were inclined to interpret a violation in context and attribute it to situational factors rather than to Austria-Hungary's underlying disposition (type). As the German ambassador to Austria-Hungary reported, "The fact that the Dual Monarchy is not inclined or able to act in a military way ... is due to her sorry domestic situation and her reduced circumstances," which included the regionally destabilizing formation of the Balkan League among Greece, Bulgaria, Serbia, and Montenegro against the Ottoman Empire (Mercer 1996:98). Consistent with the theory, the German ambassador's assessment of the credibility of Austria-Hungary appears to have been influenced by its situation in the region, with past behavior discounted in the context of alliance commitments being violated at higher levels through the region and through the decade as new alliances formed around the Balkan League.

Before proceeding, we should also note that denying an unreliable state a new alliance is not the only strategy that potential partners may take. States may still form an alliance with unreliable partners but account for the higher risk of abandonment in the terms of the contract itself (Leeds and Anac 2005; Mattes 2012; Narang and LeVeck 2011; Poast 2012a; 2012b). Our claim is only that withholding an alliance is one strategy that states may pursue, and this strategy becomes more likely when a potential partner is more unreliable.

Research design and method

In demonstrating that the reputational consequences of violating international agreements are mediated by the degree to which past violations effectively

separates states, our purpose in this article is to provide empirical evidence of a specific mechanism through which past actions affect future behavior. This article does not seek to advance the most complete predictive model of alliance formation. As a result, the empirical models used to provide evidence for the hypothesized mechanism do not control for every single variable posited to affect alliance formation in the previous literature. Rather, we control for potential confounds that could affect both the decision to violate an alliance contract in the past and the decision to form a new alliance at any moment.

We test our main hypothesis using data on bilateral alliance violations from 1919 to 2001 (Leeds, Mattes, and Vogel 2009).¹⁰ Our unit of analysis is state-years. In total, the sample includes 234 alliances, 74 of which end in violation of their terms, and 4,997 state-year observations. The key independent variable of our theory is a state's past violation history in a given year with each violation discounted by the regional violation rate, or *violations-in-context*. To test our hypothesis, we measure the impact of this signal on the probability that a state gains a new bilateral alliance in the Alliance Treat Obligations and Provisions Dataset (Leeds, Ritter, Mitchell, and Long 2002).

We limit our analysis to years in which states are involved in at least one alliance, as opposed to including all state-years. We do this because states often choose not to participate in alliance politics entirely for reasons that have little to do with their past alliance reliability. For example, the United States exits the data from 1931–1942 during the period of American isolationism that began in the wake of WWI. Likewise, Switzerland only enters our data in 1956 when it signs a nonaggression pact with the Philippines—an exceedingly rare act by a country with a long history of avoiding foreign commitments. In both cases, the absence of alliances reflected the foreign policy preferences rather than past alliance behavior of each state. This is similar to how studies of war focus on “politically relevant dyads.” Just as certain states will never reasonably engage in conflict, certain states will never reasonably ally regardless of their reputation. Nevertheless, in online Appendix Table A7, we show that the empirical results are robust and even more consistent with our hypotheses when we include all states-years from 1919–1989 in our analysis. However, this almost certainly overestimates the impact of previous violations as they steadily accumulate over time, while states exit the alliance market for other reasons.¹¹

¹⁰We restrict our analysis to the period between 1919 and 1989 because virtually none of the alliances after 1989 has terminated, raising the possibility that there is something fundamentally different about the post-Cold War era. However, extending our analysis to 2001 does not change any of our findings. Furthermore, we restrict the analysis specifically to bilateral alliances because it is much more difficult for both states and analysts to attribute violations in a multilateral context (Leeds and Mattes 2007).

¹¹To be clear, we never disregard violations, which would bias the measurement of the independent variable. Rather, limiting the analysis to states involved in at least one alliance affects the population of cases under observation. For reasons noted, we believe this is both reasonable and preferable to observing all states in the international system because it reduces unobserved heterogeneity from selection into an alliance. As noted, online Appendix Table A7 shows that the results support our argument even more strongly when we rerun the analyses in the full population of state-years from 1911–1989.

Measuring reputation: Data and empirical model

To calculate our independent variable, violation history in context, we use three steps. First, we calculate states' violation history using the Leeds et al. (2009) data. These data code which state in an alliance was actually responsible for violating, allowing us to test the directional implications of our theory.¹² To calculate a state's violation history in a given year, we sum all previous bilateral alliance violations by the state in previous years.

$$Violation\ History_{it_k} = \sum_{t_0}^{t_k} Violation_{it}$$

The second component of our independent variable is the regional violation rate in a given year. We use this as an indicator of how easy or hard the times in which the violation took place actually were. States may find it necessary to abrogate their alliances for a variety of reasons like famines, regional civil wars, economic shocks, natural disasters, and other crises. However, in terms of our theory, the actual measures of these indicators are less important than the observable alliance behavior they produce. In many cases, states choosing to violate alliances may have an incentive to misrepresent how hard times really are in order to abrogate commitments. Given the incentive to bluff, states evaluating potential partners rely on the actual behavior of similar states. This is very similar to how banks issuing mortgages rely on the behavior of an applicant's peer group to discern whether past loan terms were violated because the applicant was simply an unreliable type or if other exogenous conditions caused the violation. At the heart of this strategy is the problem that talk is cheap, and thus verbal justifications for a violation often lack credibility compared to observable behavior.

We utilize the cohort of states in the same region and time as only one among many possible ways that states in the international system may evaluate the past violations of potential allies. Previous research has shown that many of the conditions that could put pressure on an alliances, and constitute harder times, are regionally clustered rather than conserved across other possible cohorts, such as regime type (Huth 1997). However, to be clear, we acknowledge that there are many dimensions along (or cohorts within) which states can (and probably do) evaluate the behavior of potential allies, including similar regime type, similarity between the present and past context, or the degree to which a potential ally is similar to a state's previous partner. There is no a priori reason to suspect that evaluating the behavior of a potential ally within any of these

¹²Leeds et al. (2009:469–470) adopted the following rules for coding which alliance member was responsible for terminating an alliance: (1) If one state obviously violates key terms of the alliance, for instance by attacking the ally or failing to come to the ally's assistance; (2) Barring a clear violation of the treaty, if one state unilaterally declares the alliance over or breaks diplomatic relations; (3) In all other cases, Leeds et al. make judgments based on who most experts judge to have ended an alliance.

cohorts is any more or less informative. Indeed, it may require an equally strong assumption to define the cohorts another way, like regime type, as the existing evidence that democracies and autocracies exhibit similar patterns in alliance violations has been somewhat inconclusive. For our part, we demonstrated in Figure 1 that region is highly correlated with observable alliance behavior in such a way that would be relatively easy for outside observers to identify and interpret. It bears reemphasizing, though, that we are not claiming to have identified the single correct cohort, but rather we have identified *a* correct cohort in which to test our theory.¹³ Which cohort is *most* informative is a question that we purposely leave aside, as our more modest purpose is to simply show how context (whether that be regime type, region, or anything else) matters for how a particular action influences beliefs.¹⁴

To calculate the regional default rate, we sum the total number of bilateral alliance violations in each Correlates of War region in a given year. We then divide this number by the total number of alliances in the region for that year to get the average rate of commitment violation in a given year per region.

$$Regional\ Violation\ Rate_{R_g t} = \frac{\sum_{i \in R_g} Violation_{it}}{\sum_{i \in R_g} InAlliance_{it}}$$

Finally, we perform a third step to create our independent variable. In the theory we previously outlined, not all violations are equally informative. When states violate alliances in periods of high default, the violation reveals less information about how reliable a state is. Likewise, when states honor an alliance in a year in which most other states also honor their alliances, the choice to honor the alliance signals very little to a potential ally. It is only when states violate alliances in easier times or honor alliances in harder times that the behavior effectively separates the less-reliable types from the more reliable. To capture this, for each state, i , in year, t , we divide each previous violation by the regional violation rate in the year the violation occurred. We sum this value across all previous years for an aggregate measure of actual unreliability.

$$Unreliability = \sum_{t_0}^{t_k} \left(\frac{Violation_{it}}{Regional\ Violation\ Rate_{R_g t}} \right)$$

¹³Furthermore, it would be difficult to explain the empirical results toward Hypothesis 1 if region and time were not informative peer groups within which to interpret a violation.

¹⁴We acknowledge that the signal sent by an alliance violation does not uniformly influence beliefs. Different partners may interpret violations as a more or less informative of reliability (Morrow 1994). As a reviewer suggests, democratic states may be inclined to forgive violations caused by democratic constraints. Alternatively, violations with African states may suggest greater unreliability to states in Africa but greater reliability to states in Europe. Finally, certain alliance types (like defensive pacts) may be punished more than others (like consultation pacts). These are all reasonable suppositions that we do not test. The goal here is more modest: to advance a simple but important correction to existing research on alliance violations and reputation by arguing that context matters. In doing so, we report an average treatment effect across the full population over time.

$$Violation_{it} = \begin{cases} 1 & \text{if end alliance by violating terms} \\ 0 & \text{if continue alliance or end for other reason} \end{cases}$$

The variable captures the essential feature of the theory needed to test our hypothesis.¹⁵ Each alliance violation is modified—or effectively deflated—by the regional violation at the time the violation occurred, and the overall unreliability score is the running summation over time.¹⁶ In other words, as the number of violations that occurred in periods of higher regional default increases, a state's violation history is discounted in context, and the overall measure decreases.¹⁷ Alternatively, if a state has only a few violations, but these violations occurred in periods of low regional default, the informational value of each violation is relatively inflated toward perceptions of Unreliability. Finally, the measure *rewards states that honor their alliances*, particularly in the hardest of times, because fewer violations equate to lower perceptions of Unreliability (with the most Reliable states that honored all alliances bounded at 0).¹⁸

In the following analyses, we control for the possibility that some regions are simply more active in alliance politics by modeling regional heterogeneity in alliance activity using regional fixed effects and by controlling for the general frequency of alliance relationships in the region as a whole in online Appendix Table A8. We show that our results are robust to the inclusion of both controls. We also control for the possibility that, in many cases of alliance violations, there may only be a single violation in the year in question—in which case, higher levels of unreliability might simply reflect the existence of fewer bilateral alliances in the region in the year of the violation. In online Appendix Table A4, we show that our results are robust to dropping regions in which only one state tends to violate its alliance by restricting our analysis to Europe, where alliance violations are frequent across different states in a given year. In online Appendix Table A5, we also show that our results are robust to the inclusion of a wider range of 10 years when calculating the regional violation rate, which further ensures that our results are not driven by a single alliance violation constituting the majority of the regional violation rate. Finally, in online Appendix Table A6, we show that our results are robust to both solutions implemented simultaneously, with the sample

¹⁵As a reviewer notes, regional context is certainly not the only variable that shapes the information conveyed by an alliance violation. If this article sought to construct the most complete and accurate empirical model of reputation, we would certainly want to include many additional variables to more accurately predict the likely reputational consequence of a violation. However, this is beyond the scope of the current article.

¹⁶This is effectively an interactive model where a state's violation history is modified by the regional violation rate.

¹⁷Online Appendix Table A1 summarizes the *Unreliability* scores generated for each violation in the data set.

¹⁸We only consider the behavior of states that are in at least one alliance. States that never form alliances are likely to be fundamentally different for many reasons. If we were to include all state-years, our measure would consider a state that has never had an alliance to be more desirable than states that have honored 99% of their obligations. For this reason, we do not include states outside an alliance. Moreover, our test requires that we compare our measure against established measures in the literature, which only look at the effect of past alliance violations on future alliance formation, to determine which theoretical understanding of reputation best explains the data.

limited to Europe only, while including a 10-year range of years in calculating the regional violation rate.

Two additional points are worth mentioning before proceeding to the results. First, our model treats reputation as a characteristic attached to states rather than a characteristic that follows individual leaders (Gibler 2008; Guisinger and Smith 2002; Sartori 2002; Wolford 2007). It is not obvious to us that international reputation completely resets once a new leader comes to power, nor does the empirical evidence in the literature prove this supposition (Leeds et al. 2009; Renshon, Dafoe, and Huth 2016; Weisiger and Yarhi-Milo 2015).¹⁹ In many cases, the preferences of the electorate, compositions of domestic legislature, and domestic political institutions remain constant from one leader to the next. This is not to say that leadership changes do not effect international reputation but rather that some enduring component of that reputation is attached to the state and not the leader. Nevertheless, to address this possibility, we constructed a new measure, which is exactly the same as our current measure of *Unreliability*, except that it measures the unreliability of each leader within a state, resetting with each leadership change. The results using this measure are shown in online Appendix Table A3 column 4. Importantly, we show that this measure of unreliability does not change the results. Second, as a robustness check, we constructed a second measure of alliance reliability using a spatial measure of region. Our results do not change.

Results: International reputation and future alliance gains

If our hypothesis is correct, then we expect past violations to decrease the probability a state will gain an alliance in any given year but that this effect will be much stronger when past violations occurred in region-years where few other states were violating their alliances and much weaker in region-years where many states were violating. As described in the previous section, our measure, *Unreliability*, captures this distinction by penalizing states for violations that occur in region-years with low violation rates, when a violation more clearly distinguishes a state as truly unreliable. Thus, holding other factors constant, we expect to find a negative relationship between an increasing *Unreliability* score and the probability that a state gains a new alliance in any given year.

Additionally, if the context in which a violation occurs actually matters for how a violation is ultimately treated (as our signaling theory uniquely predicts), then our measure of *Unreliability* should show a stronger and more statistically significant effect than the conventional measure of alliance reputation in the existing literature (Crescenzi et al. 2012; Gibler 2008), which

¹⁹Leeds et al. (2009) shows that alliance behavior does not change with leadership turnover in democratic regimes, nor is it deterministic in autocratic regimes. Renshon et al. (2016) provide support for treating reputation as state-centric. Weisiger and Yarhi-Milo (2015) find evidence consistent with reputation affixing to states.

treats all violations the same by summing them over time. This is because in lumping together violations that will have very little effect on a state's ability to gain new alliances with violations that should have a strong effect, the conventional measure will water down the average effect of each violation and increase the variability in the estimated effect.

We demonstrate both effects empirically by competing two different logit models.

Model A :

$$\Pr(y_{it} = 1) = \alpha + \beta_1 \text{Unreliability} + \beta_2 X + u$$

Model B :

$$\Pr(y_{it} = 1) = \alpha + \beta_1 \text{Violation History} + \beta_2 X + u$$

In both of these models, the dependent variable, y_{it} , is a dichotomous variable that captures whether state i gains a new bilateral alliance in year t . β_2 is a vector of covariate parameters, X' is a vector of covariates, and u is the error term.

Note that the two models differ only in how they assess a state's reputation. Model A uses our measure, *Unreliability*, which accounts for the context under which a violation occurred by deflating the value of violations in harder times (high region-year violation rates) and inflating the value of violations in easier times (low region-year violation rates). Model B uses the more conventional measure, which is simply a running count of all the previous violations. We call the existing measure *ViolationHistory* and note that it makes no distinction between the expected effect of any two violations. Again, we expect that the estimated effect for *Unreliability* and *ViolationHistory* will both be negative but that the estimated effect for our measure, *Unreliability*, will be larger and less variable than that of *Violation History*.

Covariates

As recommended by Achen (2005), we keep our model as simple as possible, limiting the covariates to variables that consistently appear in related studies and that present an obvious threat of selection bias.²⁰

Major power status

Major Power Status is a dummy variable coding major power status according to the Correlates of War data. Empirically, major powers are more active in the alliance politics. This is probably because states perceive them to be more attractive partners by virtue of their capabilities. However, this affords

²⁰We exclude variables measuring threat because they do not present an obvious source of selection bias, and we do not intend to present a complete theory of alliance formation.

them more opportunities to violate alliances without loss in demand. Thus, not accounting for major power status risks omitted variable bias.

Democracy

Democracy is a dummy variable coding whether a country's score is greater than 0 and (alternatively) greater than 6 on the Polity IV scale.²¹ Gartzke and Gleditsch (2004) argue that democracies are less-reliable allies, which suggests a potential confound, whereby the negative effect of violating an alliance is actually picking up the effect of a country being a democracy.

Lagged number of alliances

Lagged Number of Alliances is the total number of alliances held by a country in the previous year. This proxies for how active a country is in alliance politics, with higher numbers signaling a higher demand for alliance formation. If countries with varying propensities to form alliances are more or less likely to break commitments, this could confound our result.

Lagged number of alliances²

Lagged Number of Alliances² is the total number of alliances in the previous year squared. This term is included to account for the fact that after gaining a certain number of alliances, a country's demand for alliances will decrease, as it already has gained enough alliances to meet its security needs.

Region

We code a fixed effect for each Correlates of War region, with Europe as the omitted category. Because our independent variable, *Unreliability*, is constructed by adding information about the regional violation rate to each violation, we want to ensure we are not capturing some constant effect of the region. A regional fixed effect controls for any constant mean effect that the region imparts.

Time since alliance gain

We include a cubic time polynomial, which codes the years, years squared, and years cubed, since a country last gained an alliance. We explain this further in the following discussion of the models.

Decay

In some models we also include a cubic time polynomial, coding the years, years squared, and years cubed, since a country's last violation. We explain this further in the following discussion of the models.

²¹Coding a country as a democracy if it is greater than 6 on the Polity scale, or using Polity as a continuous measure, does not substantively affect our results.

In the online appendix, we include additional controls to show that our results are robust to changes in the external security environment and leadership changes (Table A3), as well as the total number of alliances in region in a given year (Table A4).

Model specifications

To test our main hypothesis, we estimate three different specifications for each of the two previous competing models. The first specification is a baseline model designed, while the second two specifications (Models 2 and 3) check if our finding is robust to different assumptions about temporal dependence. The results for each specification can be seen in Figure 2, beginning with the simplest specification to the left and moving to the most complicated specification on the right (explained in the following).

Here we plotted the point estimate for each of our logit coefficients as a dot, with a vertical line representing the 95% confidence interval. For each specification, black dots correspond to Model A, which uses our *Unreliability* variable, while white dots correspond to Model B, which uses the competing *ViolationHistory*. For clarity of presentation, we do not plot the coefficients of the regional fixed effects or cubic time polynomials described in the following. Readers may find the estimated coefficients for all variables in online Appendix Table A2.

To further aid the reader in comparing coefficients, in Figure 2 we transformed the regression inputs for scalar variables (*Unreliability*, *ViolationHistory*, and *Lagged Number of Alliances*) by subtracting the mean and dividing by two standard deviations. This places scalar variables on the same scale as each other and approximately the same scale as the binary variables (Gelman 2008). Each coefficient in Figure 2 can now be interpreted as moving from one standard deviation below to one standard deviation above the mean of the variable. This is particularly helpful when comparing the coefficients for *Unreliability* and *ViolationHistory*, which have very different ranges.²²

Model specification 1: Baseline model

Our main baseline model specification includes all of our main covariates and regional fixed effects but excludes *Time Since Alliance Gain* and *Decay*. To deal with our observations not being independent across time, here we estimate our model using cluster-robust standard errors to correct for serial correlation by allowing the errors to be correlated within each state (Liang and Zeger 1986).²³ This method has been shown to effectively correct

²²Coefficients estimates that have not been rescaled are in online Appendix Table A2.

²³For example, the weariness with alliances that characterized American isolationism (1931–1942) followed, not by coincidence, two decades of intense alliances that proved extremely costly for the United States. Likewise, the subsequent spike in US alliances starting in 1942 was largely driven by the lack of alliances in the prior period.

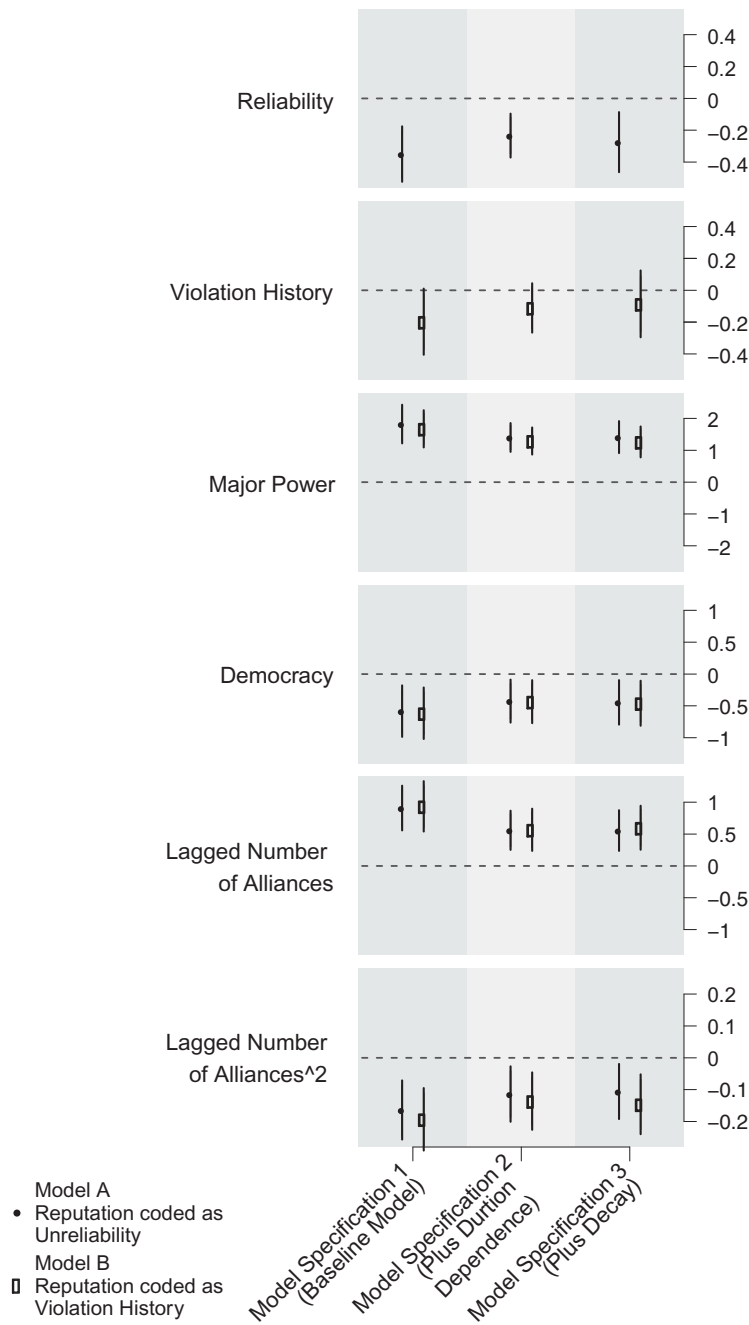


Figure 2. Unreliability versus Violation History and the probability of gaining a new alliance.

standard error estimates when the number of clusters is large (Angrist and Pischke 2009, Bertrand, Duflo and Mullainathan 2004).

The results clearly support our hypothesis. The higher a state's *Unreliability* score, the lower the probability a state will gain a new bilateral alliance.

Furthermore, as expected, the coefficient on *Unreliability* is slightly larger and less variable than the estimated effect for *ViolationHistory*. *Unreliability* is significant at the 1% level, while *ViolationHistory* is only significant at the 10% level.

Also, to more formally show that our measure of reputation better explains the data, we ran the distribution free test suggested by Clarke (2007) for nonnested model selection. This tests whether the median log-likelihood for each observation is higher under Model A (our measure) than Model B. Model A is preferred to Model B ($p < 4e-13$).

In Figure 3, we analyze the substantive effect of violation context by showing the estimated effect of an alliance violation moving from the hardest to easiest of times. Following King, Tomz, and Wittenberg (2000), we simulate draws from the variance-covariance matrix of our model and then estimate the probability of gaining an alliance given one violation at every *Unreliability* score between 6 (the lowest score for one violation) and 81 (the highest score for one violation), holding all other variables at their median.

Recall that a single violation can generate extremely different *Unreliability* scores depending on the context (regional violation rate) in which it occurred. Figure 3 shows these estimates, with a state's *Unreliability* score on the x -axis and the probability of gaining a new alliance in a given year on the y -axis. The gray bands show the estimated uncertainty.²⁴ On average, the probability of




Figure 3. Probability of gaining a new alliance as a function of unreliability.

²⁴The confidence intervals shown in Figure 3 only overlap because of the range considered. The significance of context becomes much greater over the full range of the variable (for example, second to third violation, third to fourth violation), as indicated by the coefficient estimates.

gaining an alliance with one violation and the lowest *Unreliability* score (6) is 1.8 times higher than a state with one violation but the highest *Unreliability* score (81). In other words, violations appear to affect perceptions of alliance unreliability much more when they occur in regions and times where few other states are violating.

The estimated effects of other covariates are also sensible. Major powers are more likely to gain alliances. Additionally, democracies have a harder time gaining alliances, consistent with Gartzke and Gleditsch (2004). Also, as expected, states with more alliances are more likely to gain future alliances. Finally, the coefficient on alliances² is negative, suggesting that states stop demanding more alliances after a certain point.

Model specification 2: Baseline model plus duration dependence

In the baseline Model 1, we attempted to solve the problems of serial correlation using cluster-robust standard errors by state. However, while this solves the problem of underestimating standard errors, this solution does not directly address a second problem of inefficient estimates (Beck and Katz 1997). Substantively, it is possible that gaining (or losing) an alliance may reduce (or increase) a state's demand for alliances in subsequent years. Beck, Katz, and Turner (1998) propose solving this problem by directly modeling temporal dependence, and Carter and Signorino (2010) show that this can be accomplished by including a time polynomial for the years since an event. Following this recommendation, Model 2 and Model 3 listed in online Appendix Table A1 and plotted in Column 2 of Figure 2 includes the number of years, years², and years³ since a country last gained an alliance (listed as Time Since Alliance gain in Table A1).²⁵

Crucially, including the time polynomial does not change the coefficient on our main explanatory variable, *Unreliability*.²⁶ As shown in column 2 of Figure 2, increasing unreliability remains associated with a lower likelihood of gaining an alliance. On the other hand, *Violation History* becomes statistically insignificant at conventional levels.

Model 3: Baseline model plus duration dependence and decay

There is another possible source of temporal dependence in the data. Other work on reputation has suggested that the effect of violating an alliance may decay because changes in a state over time (leadership, institutions, or population) might make past actions less diagnostic of future actions

²⁵Online Appendix Figure A1 uses the time polynomial to show how the predicted probability of gaining a new alliance changes after a recent alliance gain (all other variables set at their median). Immediately following an alliance gain, countries are much less likely to gain an alliance.

²⁶Allowing for a fourth and fifth order term does not change our results.

²⁷There are many ways to control for this possibility. Gibler experiments with limiting reputation to the leader-tenure and with models that limit violations to finite times. Another approach involves directly modeling the decay using a distributed lag model, like the Koyck model. Both approaches, and many others, are reasonable.

(Crescenzi et al. 2012; Gibler 2008).²⁷ If there is decay in the effect of reputation, this implies that a state's baseline probability of gaining an alliance rises after it incurs a violation.²⁸ This risks biasing the result if states with high *Unreliability* systematically violate their alliances later than states with low *Unreliability*, because the former would have less time for the effect of their violation to decay. This might cause us to overestimate the effect of their *Unreliability*.

We address this problem in a manner similar to model specification 2, by including a time polynomial that counts the years, years², and years³ since a state's last violation. This directly controls for any difference in a state's probability in gaining an alliance that is due to the time that has elapsed since its last violation. Column 3 in Figure 2 demonstrates that including the time polynomial in our model does not substantially impact our estimate of *Unreliability*. Notice again that the competing measure, *Violation History*, becomes statistically insignificant.²⁹

Conclusion

In this article, we proposed a signaling theory of alliance politics that structures the existing literature on alliances while generating new testable implications. Bilateral security alliances are simply international contracts where states have private information about their willingness and ability to honor commitments in the future. The mechanism that allows alliance commitments to be exchanged is taking costly actions to signal reliability under uncertainty. In this article, we highlighted the role of reputation as a signal in alliance politics.

Drawing on insights from signaling theory, we proposed a simple model of international reputation that describes how states interpret past alliance behavior when issuing future alliance contracts. This yielded a key hypothesis we test with one possible operationalization of a signal's strength. Specifically, we found that in some contexts, when many countries in a region are violating their alliances (harder times), violating an alliance does not clearly distinguish a state. As the theory predicts, violating an alliance in these contexts affects a state's ability to gain future alliances significantly less than violations in easier times, on average.

We believe the application of costly signaling theory and attribution theory to the study of reputation in alliance politics makes an important contribution to prior work (Crescenzi et al. 2012; Gibler 2008). Previous models tacitly assumed that state leaders were subject to the "fundamental

²⁸Online Appendix Figure A2 plots the predicted probability of gaining a new alliance as a function of the years since a country's last violation. Consistent with the literature, we find evidence that the reputational cost incurred by a violation decays over time.

²⁹Regression results listed in columns 5 and 6 of online Appendix Table A1.

attribution error” and always took alliance violations as a signal of states’ intrinsic unreliability. By contrast, our findings suggest that leaders do take situational factors into account when interpreting the past violations of potential allies. We demonstrate this by showing that models of international reputation that overemphasize dispositional attribution while discounting “situational” attribution underperform those that take the latter seriously.

The theory also identifies several avenues for future research. First, as noted in the theory section, it is likely that states employ other tools besides withholding alliances to offset the risk posed by unreliable partners. For example, just as banks issue variable mortgage rates or credit card companies offer different interest rates, the terms of alliance agreements may be adjusted to compensate states for the higher risk of default posed by unreliable allies. Second, states may diversify their overall alliance portfolio if they believe an ally is unlikely to honor its commitment. This could motivate partners of unreliable states to sign a greater number of alliances in order to hedge against the risk of the ally defaulting. Third, states that are least likely to honor their alliance may also be asked to post some form of collateral in order to reduce the probability they will default in the future. Together, evidence toward these expectations would offer a novel explanation for variations in the institutional design of cooperative security agreements. Finally, we believe future research should further investigate the strategic aspects of alliance violations, as even reliable states may default on alliance obligations when they believe it will not significantly affect their reputation.³⁰

Acknowledgments

We thank Michaela Mattes, Brett Ashley Leeds, Paul Poast, Brett Benson, David Lake, Erik Gartzke, Doug Gibler, Branislav Slantchev, Miles Kahler, Mark Crescenzi, Todd Sechser, James Morrow, James Long, Robbie Narang, and the participants in the International Relations Workshop at the University of California, San Diego, and the Center for International Studies at the University of Southern California. We are especially grateful to Michaela Mattes and Ashley Leeds for helpful discussion and comments on earlier drafts. A version of this paper was presented at the Midwest Political Science Association Meeting in March 2009 and the Annual Meeting of the American Political Science Association on September 3–6, 2009.

References

- Achen, Christopher H. (2005) Let’s Put Garbage-can Regressions and Garbage-can Probits Where They Belong. *Conflict Management and Peace Science* 22(4):327–339.
- Akerlof, George A. (1970) The Market for “Lemons”: Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics* 84(3):488–500.

³⁰On how reputational concerns influence bargaining strategies, see Sechser (2010).

- Angrist, Joshua D., and Jorn-Steffen Pischke. (2008) *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton, NJ: Princeton University Press.
- Altfeld, Michael F. (1984) The Decision to Ally: A Theory and Test. *Western Political Quarterly* 37(4):523–544.
- Banks, Jeffrey S. (1991) *Signaling Games in Political Science*. New York: Routledge Press.
- Beck, Nathaniel, and Jonathann Katz. (1997) The Analysis of Binary Time-Series Cross-Section Data and the Democratic Peace. Paper presented at the annual meeting of the Political Methodology Group, Columbus, OH, July 24–27.
- Beck, Nathaniel, Jonathan Katz, and Richard Tucker (1998) Taking Time Seriously: Time-Series-Cross-Section Analysis with a Binary Dependent Variable. *American Journal of Political Science* 42(4):1260–1288.
- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan. (2004) How Much Should We Trust Differences-In-Differences Estimates? *Quarterly Journal of Economics* 119(1):249–275.
- Carter, David B., and Curtis S. Signorino. (2010) Back to the Future: Modeling Time Dependence in Binary Data. *Political Analysis* 18(3):271–292.
- Choi, Ajin. (2003) The Power of Democratic Cooperation. *International Security* 28(1):142–153.
- Clarke, Kevin A. (2007) A Simple Distribution-Free Test for Nonnested Model Selection. *Political Analysis* 15(3):347–363.
- Conybeare, John A. C. (1992) A Portfolio Diversification Model of Alliances: The Triple Alliance and Triple Entente, 1879–1914. *Journal of Conflict Resolution* 36(1):53–85.
- Conybeare, John A. C. (1994) A Portfolio Benefits of Free Riding in Military Alliances. *International Studies Quarterly* 38(3):405–419.
- Crescenzi, Mark J. C., Jacob D. Kathman, Katja B. Kleinberg, and Reed M. Wood. (2012) Reliability, Reputation, and Alliance Formation. *International Studies Quarterly* 56(2):259–274.
- Fearon, James D. (1997) Signaling Foreign Policy Interests Tying Hands Versus Sinking Costs. *Journal of Conflict Resolution* 41(1):68–90.
- Gartzke, Erik, and Kristian Skrede Gleditsch. (2004) Why Democracies May Actually Be Less Reliable Allies. *American Journal of Political Science* 48(4):775–795.
- Gartzke, Erik, and Alex Weisiger. (2013) Fading Friendships: Alliances, Affinities and the Activation of International Identities. *British Journal of Political Science* 43(1):25–52.
- Gelman, Andrew. (2008) Scaling Regression Inputs by Dividing by Two Standard Deviations. *Statistics in Medicine* 27(15):2865–2873.
- Gibler, Douglas, and Toby Rider. (2004) Prior Commitments: Compatible Interests versus Capabilities in Alliance Behavior. *International Interactions* 30(4):309–329.
- Gibler, Douglas M. (2008) The Costs of Reneging: Reputation and Alliance Formation. *Journal of Conflict Resolution* 52(3):426–454.
- Guisinger, Alexandra, and Alastair Smith. (2002) Honest Threats: The Interaction of Reputation and Political Institutions in International Crises. *Journal of Conflict Resolution* 46(2):175–200.
- Guzmán, Andrew T. (2008). *How International Law Works: A Rational Choice Theory*. Oxford: Oxford University Press.
- Hafner-Burton, Emilie, Brad L. LeVeck, and David G. Victor. (Forthcoming) No False Promises: How The Prospect of Non-Compliance Affects Elite Preferences for International Cooperation. *International Studies Quarterly*.
- Hopf, Ted. (1994) *Peripheral Visions: Deterrence Theory and American Foreign Policy in the Third World, 1965–1990*. Ann Arbor: University of Michigan Press.
- Huth, Paul. 1997. Reputations and Deterrence. *Security Studies* 7(1):72–99.
- Huth, Paul, and Bruce Russett. 1984. What Makes Deterrence Work? Cases from 1900 to 1980. *World Politics* 36(4):496–526.

- Jervis, Robert. (1989) *The Logic of Images in International Relations*. New York: Columbia University Press.
- King, Gary, Michael Tomz, and Jason Wittenberg. (2000) Making the Most of Statistical Analyses: Improving Interpretation and Presentation. *American Journal of Political Science* 44:341–355.
- Lai, Brian, and Dan Reiter. (2000) Democracy, Political Similarity, and International Alliances, 1816–1992. *Journal of Conflict Resolution* 44(2):203–227.
- Lake, David A. (1999) *Entangling Relations*. Princeton, NJ: Princeton University Press.
- Leeds, Brett Ashley. (2003a) Alliance Reliability in Times of War: Explaining State Decisions to Violate Treaties. *International Organization* 57:801–827.
- Leeds, Brett Ashley. (2003b) Do Alliances Deter Aggression?: The Influence of Military Alliances on the Initiation of Militarized Interstate Disputes. *American Journal of Political Science* 47(3):427–439.
- Leeds, Brett Ashley, and Sezi Anac. (2005) Alliance Institutionalization and Alliance Performance. *International Interactions* 31(3):183–202.
- Leeds, Brett Ashley, and Jennifer Gigliotti-Labay. (2003) You Can Count on Me? Democracy and Alliance Reliability. Presented at the annual meeting of the American Political Science Association, Philadelphia, PA, August 27.
- Leeds, Brett Ashley, and Michaela Mattes. (2007) Alliance Politics During the Cold War: Aberration, New World Order, or Continuation of History? *Conflict Management and Peace Science* 24(3):183–199.
- Leeds, Brett Ashley, Michaela Mattes, and Jeremy S. Vogel. (2009) Interests, Institutions, and the Reliability of International Commitments. *American Journal of Political Science* 53(2):461–476.
- Leeds, Brett Ashley, Jeffrey M. Ritter, Sara McLaughlin Mitchell, and Andrew G. Long. (2002) Alliance Treaty Obligations and Provisions, 1815–1944. *International Interactions* 28(3):237–260.
- Liang, Kee-Yung, and Scott L. Zeger. (1986) Longitudinal Data Analysis Using Generalized Linear Models. *Biometrika* 73(1):13–22.
- Mansfield, Edward D., Helen V. Milner, and B. Peter Rosendorff. (2002) Why Democracies Cooperate More: Electoral Control and International Trade Agreements. *International Organization* 56(3):477–513.
- Mattes, Michaela. (2012) Reputation, Symmetry, and Alliance Design. *International Organization* 66(4):679–707.
- Mercer, Jonathan. (1996) *Reputation and International Politics*. Ithaca, NY: Cornell University Press.
- Mercer, Jonathan. (1997) Reputation and Rational Deterrence Theory. *Security Studies* 7 (1):100–113.
- Miller, Gregory D. (2012) *The Shadow of the Past: Reputation and Military Alliances before the First World War*. Ithaca, NY: Cornell University Press.
- Morrow, James D. (1991) Alliances and Asymmetry: An Alternative to the Capability Aggregation Model of Alliances. *American Journal of Political Science* 35(4):904–933.
- Morrow, James. D. (1994). Alliances, Credibility, and Peacetime Costs. *Journal of Conflict Resolution* 38(2):270–297.
- Morrow, James D. (2000) Alliances: Why Write Them Down? *Annual Review of Political Science* 3:63–83.
- Narang, Neil. (2014) When Nuclear Umbrellas Work: Signaling Credibility in Security Commitments through Alliance Design. Presented at the annual meeting of the American Political Science Association, Washington, DC, August 28–31.
- Narang, Neil, and Brad LeVeck. (2011) Securitizing International Security: How Unreliable Allies Affect the Structure of Alliance Portfolios. Presented at the annual meeting of the American Political Science Association, Seattle, WA, September 2.

- Narang, Neil, and Rupal Mehta. (2015) The Unforeseen Consequences of Extended Deterrence: Moral Hazard in a Nuclear Client. Presented at the annual meeting of American Political Science Association, San Francisco, September 3–6.
- Poast, Paul. (2012a) Can Issue Linkage Improve Treaty Credibility? Buffer State Alliances as a “Hard Case.” *Journal of Conflict Resolution* 57(5):739–764.
- Poast, Paul. (2012b) Does Issue Linkage Work? Evidence from European Alliance Negotiations, 1860 to 1945. *International Organization* 66(2):277–310.
- Press, Daryl G. (2004) The Credibility of Power. *International Security* 29(3):135–169.
- Press, Daryl G. (2005) *Calculating Credibility: How Leaders Assess Military Threats*. Ithaca, NY: Cornell University Press.
- Renshon, Jonathan, Alan Dafoe, and Paul Huth. (2016) To Whom Do Reputations Adhere: Experimental Evidence on Influence-Specific Reputation. Available at <https://protect-us.mimecast.com/s/4QneBAFQoMo8ig?domain=07e85df4-a-62cb3a1a-s-sites.googlegroups.com>.
- Sartori, Anne E. (2002) The Might of the Pen: A Reputational Theory of Communication in International Disputes. *International Organization* 56(1):123–151.
- Schultz, Kenneth A. (1999) Do Domestic Institutions Constrain or Inform?: Contrasting Two Institutional Perspectives on Democracy and War. *International Organization* 53(2):233–266.
- Schweller, Randall. (1998) *Deadly Imbalances: Tripolarity and Hitler’s Strategy of World Conquest*. New York: Columbia University Press.
- Sechser, Todd S. (2010) Goliath’s Curse: Coercive Threats and Asymmetric Power. *International Organization* 64(4):627–660.
- Simon, Michael W., and Erik Gartzke. (1996) Political System Similarity and the Choice of Allies: Do Democracies Flock Together or Do Opposites Attract? *Journal of Conflict Resolution* 40(4):617–635.
- Siverson, Randolph M., and Juliann Emmons. (1996) Birds of a Feather: Democratic Political Systems and Alliance Choices in the Twentieth Century. *Journal of Conflict Resolution* 35(2):285–306.
- Smith, Alastair. (1995) Alliance Formation and War. *International Studies Quarterly* 39(4):405–425.
- Smith, Alastair. (1996) To Intervene or Not to Intervene: A Biased Decision. *Journal of Conflict Resolution* 40(1):16–40.
- Spence, Michael. (1973) Job Market Signaling. *The Quarterly Journal of Economics* 87(3):355–374.
- Tomz, Michael. (2007) *Reputation and International Cooperation: Sovereign Debt across Three Centuries*. Princeton, NJ: Princeton University Press.
- Walt, Stephen M. (1987) *The Origins of Alliances*. Ithaca, NY: Cornell University Press.
- Waltz, Kenneth N. (1959) *Man, the State, and War*. New York: Columbia University Press.
- Weisiger, Alex, and Keren Yarhi-Milo. (2015) Revisiting Reputation: How Past Actions Matter in International Politics. *International Organization* 69(2):473–495.
- Wolford, Scott. (2007) The Turnover Trap: New Leaders, Reputation, and International Conflict. *American Journal of Political Science* 51(4):772–788.