Journal of Peace Research

Journal of Peace Research
1–16

DOI: 10.1177/0022343318808844
journals.sagepub.com/home/jpr
SAGE

International reputation and alliance portfolios: How unreliability affects the structure and composition of alliance treaties

Neil Narang

Department of Political Science, University of California, Santa Barbara

Brad L LeVeck

Department of Political Science, University of California, Merced

Abstract

Why do states ever form military alliances with unreliable partners? States sign offensive and defensive military alliances to increase their fighting capabilities in the event of war and as a signal to deter potential aggressors from initiating a crisis. Yet, signing an alliance with an unreliable partner is at odds with both of these rationales. This should be particularly concerning for peace scholars and policymakers, since the uncertainty generated by unreliable partners may increase system-wide conflict. This article provides an answer to this puzzle by arguing that states continue to form alliances with unreliable partners because they can adopt rational portfolio-diversification strategies. Drawing on well-developed models from portfolio theory, we present evidence that states design their overall alliance portfolios to minimize the risks posed by allies with a reputation for being unreliable. Specifically, we show that unreliable allies are more likely to be pooled into multilateral alliances that dilute risk rather than bilateral alliances, and that states allied with unreliable partners form a greater number of alliances to hedge against the added risk of default. Together, our results demonstrate why unreliable partners may not lead to increased conflict initiation, while also providing a novel explanation for previously unexplained variation in the structure of alliance portfolios. The article contributes to the literatures on international reputation and the rational design of international institutions by demonstrating how international reputation matters in subtle and often overlooked ways.

Keywords

alliances, collective security, international reputation, international security, rational design of international institutions, signaling

Introduction

Why do states ever form military alliances with unreliable partners? States largely sign offensive and defensive military alliances to increase their fighting capabilities in the event of a conflict (Lake, 1999). In the case of defensive alliances, states may also hope that this increased fighting capacity will deter aggressors from initiating disputes in the first place (Johnson & Leeds, 2011; Leeds, 2003; Mehta & Narang, 2018). Yet, signing an alliance with an unreliable ally is at odds with both of these rationales. First, unreliable alliance partners are, by definition, demonstrably less likely to honor their commitment in the event of a war. This risk of abandonment may nullify the benefit of an ally's additional fighting capabilities (Snyder, 1984). Second, alliances with unreliable partners may even leave a state more vulnerable than it would have been without an alliance. This is because states often rely on the added capabilities of allies to reduce their own indigenous military expenditures

Corresponding author:
narangn@gmail.com

(Lake, 1999; Morrow, 1993; Narang, 2018a). Such burden-sharing can be highly advantageous if a partner actually honors the alliance, but disastrous if it does not (Lake, 2009). Third, unreliable allies expose a state to all of the risk of entrapment from 'entangling relations', but reduce the potential benefits due to a higher risk of abandonment (Beckley, 2015; Lake, 1999; Narang & Mehta, forthcoming). Finally, allying with an unreliable partner may increase uncertainty about whether the unreliable partner will honor its commitment. Such uncertainty can lead to conflicting expectations about the strength of an alliance and increase the chance that deterrence will fail (Sagan, 1986).

These risks combine to make alliances with unreliable alliance partners particularly concerning for peace scholars and international policymakers alike. Beyond the consequences for individual states, unreliable allies serve to increase uncertainty about which acts of aggression will trigger war (Sagan, 1986). This uncertainty, in turn, may increase system-wide conflict (Fearon, 1995; LeVeck & Narang, 2017a; Narang, 2014, 2015a). And yet, despite the risks from allying with unreliable partners, previous research has demonstrated that states continue to form alliances with such states (Crescenzi et al., 2012; Gibler, 2008; LeVeck & Narang, 2017a; Mattes, 2012). Therefore, the question remains: given the potentially disastrous consequences that states risk, why do they continue to form alliance agreements with demonstrably unreliable partners?

In this article, we argue that states continue to form alliances with some unreliable allies because they can adopt rational portfolio-diversification strategies. These strategies allow states to benefit from the military capabilities of unreliable states while minimizing the risk of abandonment. Our primary theoretical innovation is to leverage well-developed models from portfolio theory (Markowitz, 1959). Using the framework of portfolio analysis, we generate testable propositions about how states structure their alliance portfolios. Specifically, we present three pieces of evidence that states design their overall alliance portfolios to minimize the risks posed by allies with a reputation for being unreliable. First, we show that unreliable allies are more likely to be pooled into multilateral alliances that distribute risk better than bilateral alliances. Second, we further show that states internally construct these multilateral alliances to reduce the trade-off between additional capabilities and additional risk of abandonment. Third, we demonstrate that states in bilateral alliances with unreliable partners also hedge against the risk of abandonment by signing a greater number of bilateral alliances.

Taken together, our results serve to demonstrate how remarkably general the Nobel prize-winning logic of portfolio theory is across domains, while also contributing to the literatures on international reputation and the rational design of international institutions (Koremenos, Lipson & Snidal, 2001; Mattes, 2012). More importantly, our theory and evidence provide an explanation for why unreliable partners may (happily) not lead to a large increase in dispute initiation throughout the international system. Since the structure of alliance portfolios is visible to other states in the international system, portfolio-diversification strategies that reduce the risk of abandonment should also increase their deterrent effect.

In addition to explaining the puzzle of unreliable allies, a second important motivation and contribution of this project is to provide a novel explanation for the structure of alliance portfolios. To date, variation in alliance portfolios has generally been used as an explanatory independent variable, rather than the phenomenon to be explained (Bennett & Stam, 2000, 2003). For decades, scholars of international relations have exploited variation across alliance portfolios as a conceptual proxy for the overall foreign policy similarity. Since Bueno de Mesquita (1975), researchers have used either Kendall's τ_b score or Signorino & Ritter's (1999) S-score to identify alliance clusters and to measure the extent to which those clusters overlapped or were discrete (Bueno de Mesquita, 1978, 1981, 1988; Ostrom & Aldrich, 1978; Stoll, 1984; Stoll & Champion, 1985). Alliance portfolios were thus treated as revealed preferences over security issues, as states with greater commonality in their alliance portfolios were generally assumed to have more similar security interests (Altfeld & Bueno de Mesquita, 1979). Although useful, here we are the first to demonstrate that variation in alliance portfolios is not purely driven by foreign policy similarity, but that it also reflects attempts to mitigate the risk of abandonment in the event of war.

Literature: The role of reputation in alliances

Previous work has argued that alliance agreements can be usefully analogized to bonds (Conybeare, 1992; LeVeck & Narang, 2017a; Mattes, 2012). Similar to how loan applicants seek loans from potential lenders in order to increase their financial security, states in the international system seek to negotiate alliance contracts in order to increase their national security. And just as lenders must ultimately issue loans based only on the *expectation* that a recipient will eventually repay their debt, states in the international system must enter alliance contracts

based only on the *expectation* that a security commitment will eventually be fulfilled when it is invoked. These expectations matter because neither lenders nor states seeking alliances can directly observe applicants' intentions to uphold their commitment in the future. Since potential alliance partners always have private information about their willingness and ability to honor their commitment in the future, states issuing alliances always run the risk that potential alliance partners will default on their commitment when it is no longer in their interest to comply (Altfeld & Bueno de Mesquita, 1979; Hafner-Burton, LeVeck & Victor, 2017a; Leeds, Long & Mitchell, 2000; Siverson & King, 1980; Smith 1995).

Understood this way, the exchange of security commitments thus resembles other contractual relationships in which there is incomplete information with respect to quality, where – in this case – quality can be understood as future alliance reliability (Akerlof, 1970). One potential consequence of the interaction between quality heterogeneity and incomplete information may be disappearance of Pareto-improving agreements altogether where credible commitments are a problem. With respect to international alliances, this may produce suboptimal levels of alliance formation in cases where two or more states may otherwise benefit from cooperating.

A well-known solution to the incomplete information problem, originally proposed by Spence (1973), is for actors with private information to credibly signal their type by taking costly actions that poor-quality candidates cannot efficiently mimic. In international alliance politics, 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 alliance 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.[1] By combining factors from a potential alliance partner's past alliance behavior and current national profile, states can update their beliefs about the reliability of potential alliance partners based on costly signals they believe to be correlated with reliability to determine whether or not to issue an alliance and with what terms.

Consistent with the logic of costly signaling under incomplete information, a recent wave of alliance research has examined how states can screen potential alliance partners using past actions (or reputation) as a costly signal of reliability, or credible commitments. For example, Gibler (2008) tests whether previously violating the terms of a bilateral alliance decreases the chance that a dyad will form a future alliance. He finds that states that have honored their commitments in the past are more likely to find alliance partners in the future. LeVeck & Narang (2017a) also investigate the impact of alliance violations on future alliance activity with a more nuanced model of reputation formation that emphasizes the role of states' beliefs. They show that past alliance violations are only useful signals of future alliance reliability conditional on whether they effectively separate reliable from unreliable allies. Mattes (2012) shows that alliance design is also motivated by reliability concerns, as previously unreliable allies lead future partners to design subsequent contracts with greater precision, issue-linkages, and military institutionalization. Finally, Crescenzi et al. (2012) also show that past alliance violations affect the probability of future alliance formation, paying particular attention to the dynamics of reputation by modeling how the effects of past violations decay over time.

Interestingly, however, while the previous literature has shown that reputation matters for the quantity and quality of future alliance commitments, researchers have yet to fully investigate how states' perceptions of alliance reliability might affect their alliance choices at the portfolio level. Yet, we know that there is considerable variation in how countries choose to structure their overall alliance portfolios across states and over time. For example, in the period from 1919 to 2001, the average number of offensive and defensive bilateral alliances held by states in an alliance was 0.62 per year.[2] However, this number varied considerably, reaching a high of 12 separate offensive and defensive bilateral alliances held in a single year by the most active state during this period (the Soviet Union in 1967 and again in 1970) to a low of 0 bilateral alliances in a single year by the least active states.[3] Figure 1 illustrates this variation graphically, by plotting the distribution in the number of bilateral offensive and defensive alliances in all member-years from 1919 to 2001. Notice that states hold more than one bilateral alliance in their portfolio simultaneously only about a quarter of the time (27%).

[1] See Fearon (1997). For empirical investigations of costly signaling in alliance politics see Fuhrmann & Sechser (2014).

[2] According to the ATOP member-level data, there were a total of 467 bilateral alliances in the period 1919–2001.

[3] Out of the 149 states involved in alliances, 98 have at least one year in which they are involved in only multilateral alliances and no bilateral alliances.

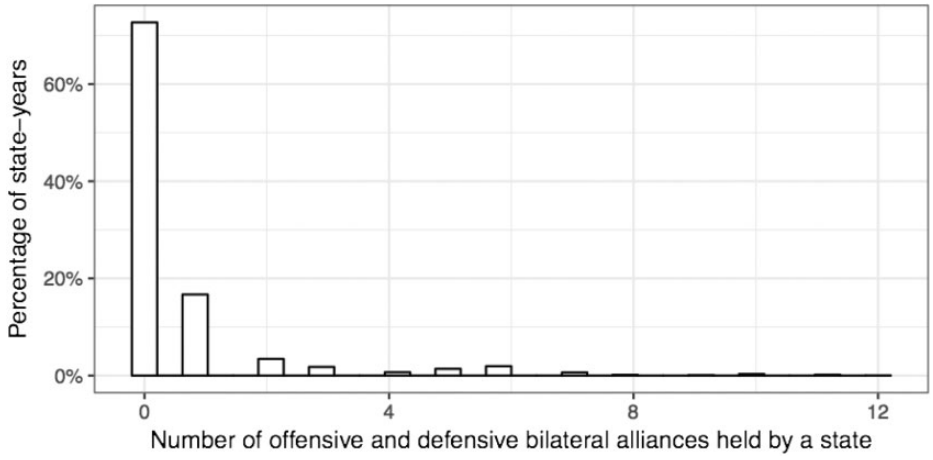

Figure 1. Histogram plotting the distribution in the total number of bilateral alliances held by active alliance members in the period 1919–2001

Additionally – as a second source of variation in alliance portfolios – states appear to differ considerably in the likelihood they will enter bilateral versus broader multilateral alliance agreements. For example, among the 535 unique military alliances active in the period 1919–2001, 467 (or 87.3%) were bilateral alliances negotiated by only two states, while 69 (or 12.7%) were multilateral alliance agreements negotiated by three or more states. However, the data suggest that a disproportionate number of the bilateral alliances were held by only a few states. Among all states actively involved in an alliance, bilateral alliances constituted, on average, only 24.5% of states' overall alliance portfolios. This means that, for some reason, many states are much more likely to find themselves packaged in broader multilateral alliances compared to others.

Equally interesting is that the overall size of these multilateral alliance agreements – when chosen over bilateral commitments – varies dramatically. Figure 2 illustrates this variation by plotting the distribution in the number of members across multilateral alliances from 1919 to 2001. Within the 69 multilateral alliances that were either offensive or defensive in nature, the average number of members was roughly seven, but this number varied dramatically across multilateral alliances. The modal number of alliance members in this period was the minimum three, while the broadest offensive or defensive agreement – signed by the Allies during World War II from 1 January 1942 to 2 September 1945 – contained 43 separate members.

In the next section, we argue that some of this variation can be explained by the risk posed by different alliance partners. In some cases, states may judge a potential ally to be too risky and decide not to ally with these potentially unreliable states. In other cases, however, states may still ally with previously unreliable states. When this occurs, states can still mitigate the risk of allying with a potentially unreliable partner by altering the structure of their alliance commitments. Therefore, variation in the structure of alliances and alliance portfolios is partially explained by the fact that some of these structures are better suited to mitigating the risk posed by potentially unreliable allies.

Theory: The role of reputation in building alliance portfolios

In this section, we follow Conybeare (1992) by suggesting that the broader collection of alliance commitments that compose states' overall alliance portfolios can be usefully analogized to the collection of financial investments in investment portfolios.[4] To begin, we assume that states invest in forming alliances because they have some expectation of potential benefits, or returns, from holding these contracts. However, we further assume that investments in alliances are always risky in that there is always some chance that a state will default on its promise to come to its partner's assistance in the event of a war (Hafner-Burton, LeVeck & Victor, 2017a; Leeds, Long & Mitchell, 2000, 2003; Leeds, Mattes & Vogel, 2009; Leeds & Savun, 2007).

[4] Conybeare (1992) details how each of the key concepts of portfolio theory formally elaborated by Markowitz, Miller, and Sharp are satisfied in the domain of military alliances, including the concepts of risk, return, correlation of returns, portfolio weights and asset, or portfolio dominance.

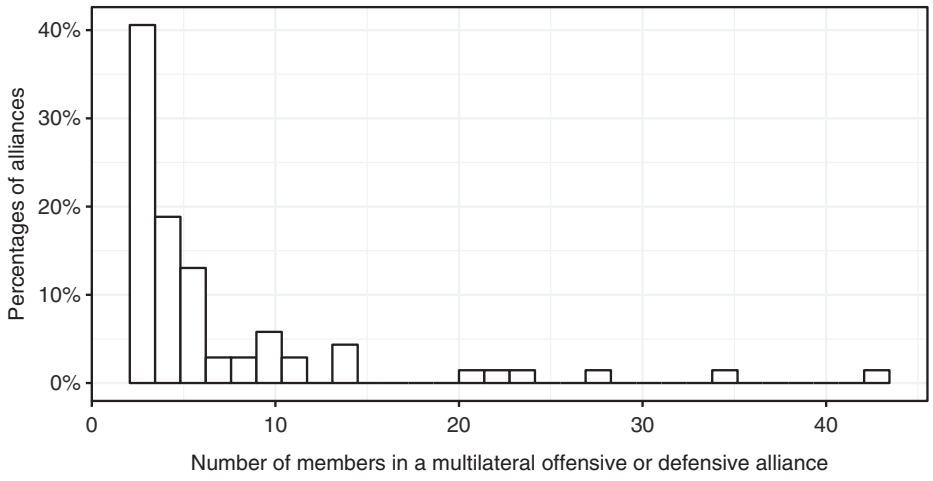

Figure 2. Histogram plotting the distribution in the total number of alliance members across multilateral alliances signed in the period 1919–2001

Treating alliance portfolios as simply collections of individual assets that are each characterized by some level of risk and return lends itself well to conventional methods of portfolio analyses. In finance, a common means for reducing risk is diversifying a portfolio by investing resources broadly in a variety of assets like stocks, bonds, and cash (Markowitz, 1959). Assuming investors do not know ex ante which assets will perform better than others, portfolio diversification limits the down-side risk of investing solely in the asset that turns out worst, though at the opportunity cost of investing completely in the one asset that turns out best. Conversely, a non-diversified portfolio invested in a single asset is fully exposed to movements in just that asset.

In terms of alliance investments, an undiversified alliance portfolio in which a state holds only one alliance commitment is far riskier with respect to the likelihood of abandonment than a diversified alliance portfolio in which a state holds multiple alliance commitments simultaneously. While the former has the advantage of limiting the transaction and coordination costs associated with alliance negotiations, it is also riskier in the sense that a single state may easily renege on its alliance commitment in the event of a war. In this case, the expected value of the alliance portfolio depends fully on the reliability of a single alliance partner. A diversified alliance portfolio like NATO, by contrast, requires greater transaction and coordination costs upfront, but it has much lower exposure to the risk of abandonment should an alliance partner fail to provide assistance in the event of a war.

While states can never fully know in advance exactly which alliances commitments will prove reliable and which will not, they may form prior beliefs about the likelihood that an alliance partner will honor its commitment based on observable indicators like those reviewed in the previous section. For instance, states may update their beliefs about a partner's reliability based on costly signals, such as whether a state honored its alliance commitments in the past (Crescenzi et al., 2012; Gibler, 2008; LeVeck & Narang, 2017a; Mattes, 2012). In cases where a state has proven itself to be unreliable by defaulting on past alliance commitments, it is reasonable to expect that other countries forming subsequent alliances with this state will take this into account in structuring their alliance portfolios. They may be less willing to let unreliable states join them in bilateral alliances, where the impact of a single state defaulting is high. However, as discussed above, they might be relatively more willing to let these states join them in multilateral alliances where other members dilute the risk of abandonment from a single default.[5] Unreliable states should thus be much more likely to end up in multilateral agreements that limit the risk of abandonment through diversification versus bilateral agreements that

[5] One might wonder if the applicability portfolio theory is limited if, in finance, an investor can freely choose across different investments, whereas in alliance formation, potential alliance partners need to accept the state as a partner. However, this distinction is illusory: market actors assembling financial portfolios are similarly constrained by the willingness of counter-parties to accept their investment. Market actors routinely reject offers of capital if the terms are unattractive, and they also exercise agency by rejecting offers of capital by not selling below a minimum price.

leave their partners fully exposed to this risk. This leads to the first hypothesis we test:

Hypothesis 1: More unreliable states will be more likely to enter a multilateral (rather than bilateral) alliance.

A good example of this hypothesized mechanism is the case of NATO expansion. NATO has undergone six rounds of enlargement since 1949 to reach its current membership of 28 states. The order in which new members were added to NATO was clearly not random. Instead, as Kamp (1995) summarizes, the sequence was a rational cost–benefit calculation by existing members with respect to whether NATO could function effectively as a military alliance once it expanded to include Eastern European countries in the fourth round. The specific concern among the original NATO members was that Eastern European states would prove to be unreliable allies both in their capabilities and their commitment to NATO's mission. This created concern among founding members that NATO would dissolve into an ineffective decisionmaking body (Waltz, 2000). Ultimately – and only after successive rounds of admitting decidedly reliable Western European states like the United Kingdom in the founding round and Germany and Spain in the second and third rounds – NATO was willing to include riskier members that were potentially less reliable in the fourth round.

Note, however, that while simply adding greater numbers of alliance partners to a given agreement should – on average – decrease the risk posed by any one particularly risky security asset (Conybeare, 1992), it is not the optimal way to manage risk, particularly if the additional allies are subject to the same regional and temporal shocks that cause states to default on their obligations. Ideally, more reliable states would be included in an alliance to offset the risk posed by particularly unreliable states in a way such that the underlying risks are less correlated. The net effect is a portfolio diversified with respect to the risk of default. Much as financial investors may try to limit the downside of more variable assets like stocks by also investing in safer assets like bonds, very unreliable states may be admitted into alliances where their risk of default (abandonment) can be mitigated by other, particularly reliable states (Markowitz, 1959).

This suggests a second, corollary hypothesis about the internal composition of the multilateral alliance portfolios that would provide a sharper test of our theory. If the structure of alliance portfolios partly reflects a desire by states to reduce their overall risk of abandonment, then we would expect states to construct multilateral alliances in such a way to reduce the correlation in the risk of default among the members.

Hypothesis 2a: Compared to bilateral alliances, the risk of default by member-states will be less correlated in multilateral alliances.

Another corollary hypothesis about the internal composition of the multilateral alliance portfolios follows from the logic of rational portfolio diversification. As in the domain of financial portfolios, it is reasonable to expect that when states form alliances with more unreliable states, these same unreliable states will – on average – provide greater military capabilities to the alliance in order to compensate for the added risk of default through expected returns.

Hypothesis 2b: The more unreliable an alliance partner, the greater its military contribution will be in future alliances.

This hypothesis is important in understanding why states do not simply avoid forming alliances with unreliable partners altogether. Similar to how the increased risk of junk bonds is offset by potentially higher returns, riskier allies should return a higher yield to states holding their commitments. This argument is also consistent with earlier work, which argued that states sometimes prefer unreliable allies because those allies provide additional capabilities above and beyond what other, more reliable allies might provide (Cha, 2010; Miller, 2012).[6]

Finally, we test a third related hypothesis in order to not underestimate the overall logic of portfolio theory. In addition to packaging unreliable security assets with reliable assets into a single multilateral alliance in order to mitigate the risk of an individual default, states may also reduce the risk of abandonment from an ally defaulting by signing multiple, discreet bilateral alliances. In theory, this is functionally equivalent to pooling risky assets in with less risky assets into a single multilateral alliance, as it also offsets potential losses that might be incurred from any single bilateral alliance. However, as a matter of practice, states may have a preference for signing bilateral alliances over multilateral alliances for a variety of

[6] Miller argues that reliable states have an easier time finding other reliable allies and have greater autonomy within their alliances. He provides evidence from the Anglo-Japanese Treaty to show that agreements between reliable states contain lower levels of obligation, precision, and delegation (2012: 90).

reasons, including lower contracting costs (Hemmer & Katzenstein, 2002).

Hypothesis 3: States with more unreliable partners will sign a greater number of offensive and defensive bilateral alliances.

The above hypotheses notwithstanding, it is also important to briefly theorize about the costs to states from adding more allies to an alliance portfolio. This is important because – in focusing on the incentive for states to reduce the risk of default – the logic thus far might be taken to imply that states should sign the broadest possible alliance commitments with as many partners as possible. However, we know that this is not the case empirically. This is because states incur at least two types of costs that may increase simultaneously with the benefits of portfolio diversification. The first, mentioned above, comes in the form of greater transaction and coordination costs associated with adding more allies. The second is the risk of entrapment. At the same time that portfolio diversification can hedge against the added risk of an unreliable alliance partner abandoning its commitment, broader and more diverse portfolios may perversely increase the separate risk of entrapment simultaneously. This trade-off is increasingly well established in the literature (Benson, 2012; Cha, 2000; Fang, Johnson & Leeds, 2014; Morrow, 2000). Both types of costs suggest that portfolio diversification is more appropriately understood as an optimization problem rather than a maximization problem. In this way, it is important to understand the hypotheses above as comparative static claims that hold, ceteris paribus.

Research design: Measuring reputation

We test our hypotheses with a measure of reputation constructed using data on bilateral alliance violations in the period 1919–89, which we call *Unreliability*.[7] Our measure of a state's perceived *Unreliability* is simply the total number of times a state has violated one of its past alliances before time *t*.[8] For example, if a state violated two alliances prior to time *t*, its unreliability score is simply 2.[9] To determine which state was responsible for the violation that terminated an alliance, we rely on the dataset coded by Leeds, Mattes & Vogel (2009).[10] According to this dataset, there were 71 alliance violations in the period between 1919 and 1989.[11] This represents a substantial percentage of the 272 active alliances in this period that ended in violation (26%), and it suggests that each violation contains a significant amount of information, since a violation occurs in only the most unreliable quartile of alliances. Table A1 in the Online appendix shows the full list of alliance violators and year by the number of violations. Figure 3 displays the distribution of alliance violations (the independent variable) in two ways. On the left, we plot temporal variation in alliance violations across COW regions, along with a worldwide total, as shown in LeVeck & Narang (2017a). On the right, we provide a cross-sectional 'heat-map' that shades countries darker by their aggregate number of violations in the period.

Several studies in the previous literature have assumed that a reputation for unreliability is a characteristic attached to *leaders* rather than a characteristic of the *state* itself (Gibler, 2008; Mattes, 2012). These studies measure reputation using a binary variable that codes

[7] We follow Leeds, Mattes & Vogel (2009) in limiting our analysis to estimating the impact of alliance violations before 1989 on alliance formation through 2001 because virtually no alliances have been terminated after 1989.

[8] We also follow Leeds, Mattes & Vogel (2009) in limiting the analysis to bilateral alliance violations. We do so because it is much more difficult for both states and analysts to attribute violations in a multilateral context (Leeds & Savun, 2007). Thus, any measure based on attributing violations in a multilateral context is likely to be equally noisy with measurement error. As a result, multilateral violations may have less impact since other states presumably have the same problem of attribution. Lastly, using bilateral violations to predict both bilateral and multilateral alliance behavior is in line with previously published work, including Gibler (2008), Mattes (2012), Crescenzi et al. (2012), and LeVeck & Narang (2017a).

[9] One risk of our measure is that it may lump together states with zero violations because they followed through on past commitments, with those that never had their commitment invoked. Ideally, we would control for the number of times a state has faced the choice to honor or violate its alliance. To address this, we include controls to proxy for how active a state is in alliance politics, including the number of alliances and measures like major power status and CINC.

[10] Leeds, Long & Mitchell (2000) offer an alternative dataset that more narrowly codes violations depending on whether states joined their allies in fighting when they were obligated to do so. We prefer Leeds, Mattes & Vogel (2009) which also codes a violation even when states unilaterally declare the alliance over or break diplomatic relations with the ally, because these types of violations also affect prospective estimates of whether a state will honor an agreement in the future. Our choice is likely to bias against our hypotheses if states do indeed treat these more forgivingly.

[11] Leeds, Mattes & Vogel (2009) adopted the following rules: (1) if one state obviously violates terms of the alliance; (2) barring a clear violation, if one state unilaterally declares the alliance over or breaks diplomatic relations; (3) if there is no clear violation and no statement ending the alliance, a decision is made based on who most experts judge ended it.

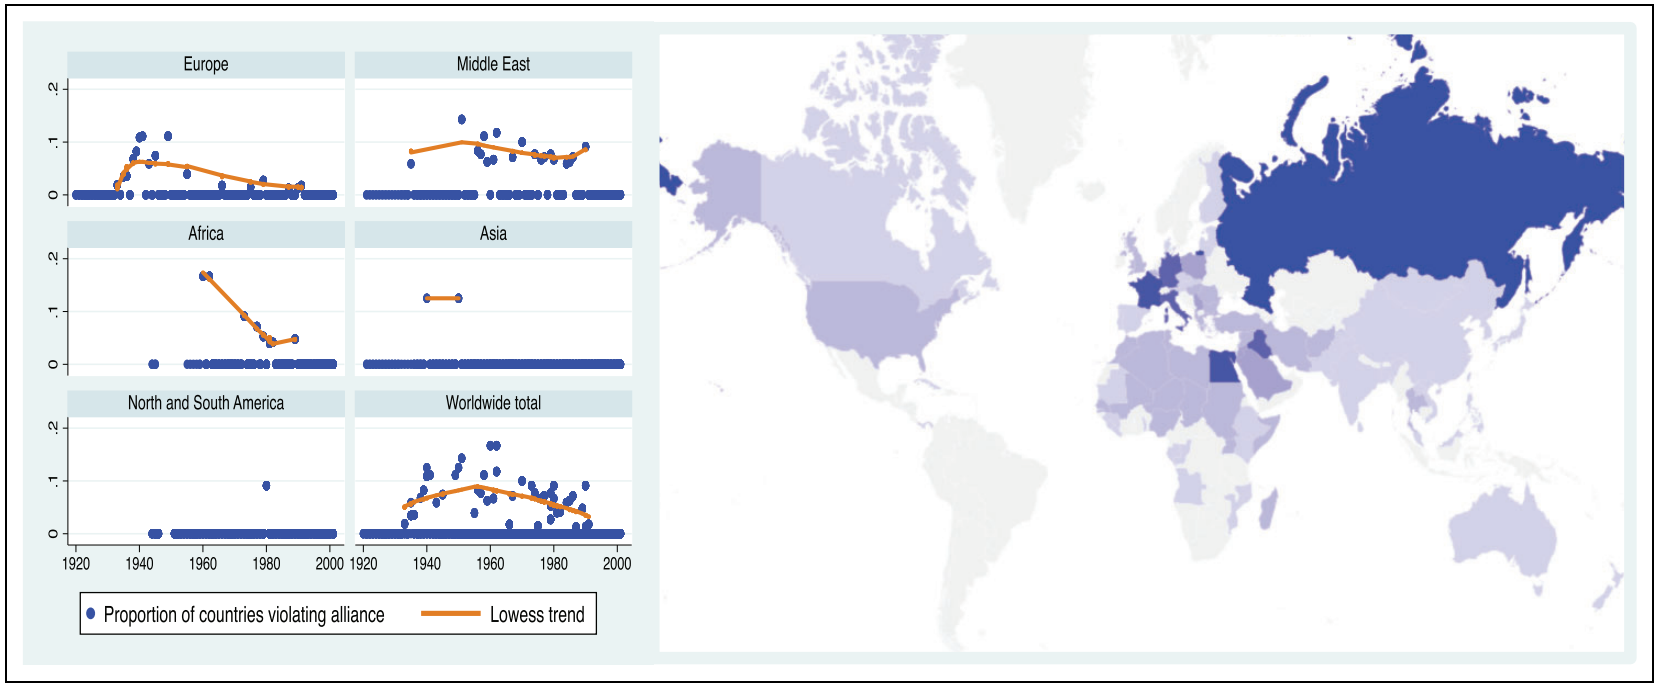

Figure 3. Distribution of alliance violations

Plotted left: temporal variation in violations across COW regions, along with a worldwide total. Plotted right: 'heat-map' shading countries darker by their aggregate number of violations in 1919–89.

whether the current leader of a state previously violated an alliance. Relatedly, other studies assume that reputations exist at the level of the state, but that violations are forgotten after a finite time-span of ten years. Both measurement choices require equally strong assumptions about the dynamics of international reputation.

In contrast to some earlier studies, we prefer to measure reputation as a state-level characteristic for both substantive and technical reasons.[12] First, alliances are 'sticky' in that they generally persist beyond a leader's tenure. Therefore, it is reasonable to assume that agreements are formalized to reflect factors beyond a leader's narrow tenure. Second, leaders often have little time to form reputations before being replaced, especially in democratic systems. In these cases, it would strain plausibility that the reputation of a nation resets entirely every time a leader leaves office.[13] A much more reasonable assumption is that there is at least *some* component of international reputation that endures from one administration to the next that may be embodied in the preferences of its citizens, the nature of its political institutions, or a number of other factors. Our only assumption is that states assess the future *unreliability* of potential alliance partners based – in part – on past alliance behavior, and that some component of these inferences are formulated based on domestic actors and institutions that are separate from the leader.

Finally, similar to Crescenzi et al. (2012) and LeVeck & Narang (2017a), we also model the decay in violations over time. We do this by including a linear term for the number of years since a state last violated an alliance. This is to control for the fact that a past violation may only impact a state's reputation for a finite amount of time. We used a linear function because we do not know the exact functional form, and linear functions are known to be more robust to misspecification error (Dawes, 1979). In particular, they are more likely to correctly estimate the sign and general magnitude of a relationship, even when the 'true' function is not linear (Dawes, 1979). However, we also follow LeVeck & Narang (2017a) by allowing the decay function to be nonlinear and fitting a polynomial function. We report these results in Table A4 in the Online appendix. Allowing the decay function to be non-linear does not change the sign, magnitude, or significance of our main independent variable, *Unreliability*.

Analysis and results

Alliance unreliability and multilateral versus bilateral alliances

Hypothesis 1 expects that unreliable states will find it more difficult to enter into bilateral offensive and defensive alliances with other states than multilateral offensive and defensive alliances.[14] Again, this is because more unreliable states represent riskier security assets that other states should be less willing to directly ally with in bilateral agreements, as they pose a higher risk of default based on past actions (Gibler, 2008; Leveck & Narang, 2017a). As a result, we posited that potential alliance partners should be more willing to admit these riskier assets into multilateral alliances that dilute the overall risk of abandonment in the event that an unreliable state fails to honor its commitment. To test this claim we examined if, each time a state gains an alliance, that alliance is either a multilateral alliance or a bilateral alliance as a function of a state's reputation for unreliability using the measure described above. Specifically, we estimated the effect of reputation on this variable using a logistic regression for the following equation:

$$\Pr(\textit{Multilateral alliance})_{it} = \alpha_i + \beta_1 \textit{Unreliability}_{it} + \beta_x X_{it}$$

where X_{it} is a vector of controls for state i in year t, and α_j is a region fixed effect for state i in region j, which allows us to capture known regional differences in states to form multilateral vs. bilateral alliances (Cha, 2010) (which may correlate with regional differences in state unreliability found by LeVeck & Narang, 2017a). We estimated the model using logistic regression with errors clustered by alliance to account for the fact that each state joining an alliance is not a statistically independent event. In two specifications, we also include a cubic time polynomial (Carter & Signorino, 2010) to account for any temporal changes in the baseline probability of

[12] See Crescenzi et al. (2012: 9–11) for a reasoned justification of modeling reputation as a state-level characteristic.

[13] Crescenzi et al. (2012: 264) provide the following illustration. 'Do states erase their knowledge of prior interactions with changes of administrations? This seems unlikely. For example, consider the current leadership change in the United States. By all accounts, this change appears to signal a dramatic change in foreign policy strategy. Yet the commitments of the United States remain largely unchanged [. . .].'

[14] We restrict our analysis to offensive and defensive alliances because they are the only alliances that obligate members to provide active military support. Because they are riskier, we expect leaders to account for the unreliability of potential allies in structuring alliance portfolios to offset some risk.

joining multilateral vs. bilateral alliances. For example, one might be concerned about the possibility that the frequency of multilateral alliances grows over time along with the cumulative number of violations, in which case a positive spurious relationship might emerge.

We include the following control variables for our baseline model:

- ***Number of alliances:*** Total number of alliances held by a country in the previous year. This proxies for how active a country is in alliance politics. If countries with a higher (lower) propensity to form alliances are also more (less) likely to violate their alliances, then this could confound our result.
- ***Major power:*** A dummy variable indicating whether a country is a major power. Empirically, major powers are more desirable partners, which means they find themselves included in a larger number of alliances despite past violations.
- ***CINC score:*** This further controls for a state's relative capabilities, which is known to be correlated with alliance violations, and may also be correlated with state's relative propensity to join multilateral (rather than bilateral alliances).
- ***Democracy:*** A dummy variable coding whether a country's score is greater than 0 and, alternatively, greater than 6 on the Polity IV scale.[15] Existing literature has shown that alliance reliability is correlated with regime-type (Gartzke & Gleditsch, 2004; Leeds, Long & Mitchell, 2003). Additionally, we know from previous work that democracies and autocracies are more likely to find themselves in alliances with their own regime-type (Lai & Reiter, 2000; Simon & Gartzke, 1996; Siverson & Emmons, 1996). If democracies also have a preference for forming broader, multilateral coalitions, this may confound our results.
- ***MID previous one year and MID previous five years:*** A dummy variable coding whether a country experienced a militarized interstate dispute in the previous one or five years based on the MID 4.0 data (Palmer et al., 2015). Threats to states may drive the structure of alliance formation and previous alliance violations.
- ***Decay:*** A linear term for the number of years since a state last violated an alliance. Violations may only impact reputation for a finite amount of time.

Table I reports the results for our regression in four models. Models 1 and 2 include MIDs over the previous one and five years, respectively, with only regional fixed effects, while Models 3 and 4 include MIDs over the previous one and five years, respectively, with regional fixed effects and time polynomials. Notice that, in all specifications, the coefficient on our main independent variable is significant and in the expected direction.[16] States with a greater reputation for alliance unreliability are significantly more likely to form a multilateral alliance rather than a bilateral alliance conditional on gaining an alliance. Substantively, for every violation of an alliance in the past, a state is approximately 4% more likely to be packaged in a multilateral alliance instead of a bilateral alliance. This suggests that potential alliance partners are more willing to form alliances with unreliable states through multilateral agreements than bilateral agreements, as our theory predicts.[17]

Alliance unreliability, uncorrelated risk, and expected returns

Hypotheses 2a and 2b can be seen as extensions of Hypothesis 1. If states construct multilateral alliances, in part, in an effort to dilute the added risk of default posed by an unreliable security asset, then one would expect the average covariance in the past alliance violations that comprise states' unreliability to be lower for multilateral alliances compared to bilateral alliances. This is because simply adding more allies to an alliance alone may not substantially reduce the risk if many allies are subject to the same regional and temporal shocks that cause states to default on their obligations. Thus, states should construct multilateral alliances to reduce the covariance in the risk of default among the members as a strategy for reducing the overall risk of abandonment. To test this expectation, we compare the average pairwise covariance in the measured unreliability of members

[15] Coding 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.

[16] While we do not aim for prediction, the R^2 for in-sample prediction is 0.30 for the first two models and 0.35 for the second two models. Our unreliability variable increases the R^2 by one percentage point. This is statistically distinguishable from the null hypothesis of 0% according to a Wald test ($p < 0.006$).

[17] In Table A2 in the Online appendix, we show that the results are also robust to an alternate measure of reliability used by LeVeck & Narang (2017a).

Table I. Impact of unreliability on the probability of gaining offensive or defensive alliance, conditional on gaining a new alliance

	DV: Enter multilateral defensive alliance (versus bilateral)			
	(1)	*(2)*	*(3)*	*(4)*
Unreliability	0.136*	0.135*	0.144*	0.145*
	(0.059)	(0.059)	(0.060)	(0.060)
MID previous 1 year	–0.116		–0.193	
	(0.276)		(0.299)	
MID previous 5 year		0.074		0.124
		(0.275)		(0.302)
Decay	0.005	0.005	0.025	0.024
	(0.015)	(0.016)	(0.020)	(0.020)
Major power	0.677	0.588	0.928	0.790
	(0.608)	(0.589)	(0.581)	(0.568)
Democracy	1.058*	1.053*	1.038*	1.035*
	(0.506)	(0.508)	(0.475)	(0.476)
Anocracy	0.842**	0.831*	0.508	0.503
	(0.303)	(0.306)	(0.304)	(0.302)
Number of alliances	–0.127**	–0.127**	–0.153**	–0.154**
	(0.045)	(0.045)	(0.043)	(0.045)
CINC	–11.075**	–11.146**	–12.173**	–12.263**
	(2.607)	(2.612)	(2.780)	(2.791)
Region fixed effects	X	X	X	X
Time polynomial			X	X
Observations	602	602	602	602
Log likelihood	–312.816	–312.895	–287.266	–287.457
Akaike inf. crit.	651.631	651.789	606.531	606.914

Robust standard errors in parentheses. $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\dagger}p < 0.1$.

within multilateral alliances over time and we compare this to the covariance in the unreliability of members in bilateral alliances.

Figure 4 is a box plot showing the distribution of unreliability covariances across alliances. We subset the distribution of covariances for bilateral alliances on the left, and multilateral alliances on the right. As expected, Figure 4 shows that the covariance is much lower on average (as shown by the black line) and also more consistently lower (as shown by the box around the interquartile range for the 75th and 25th percentile).

Hypothesis 2b posits that if there is indeed a risk–return trade-off as implied by our portfolio theory, then we should generally expect a positive relationship between increasing unreliability of an alliance partner and the expected returns from including that alliance member in the portfolio. Specifically, we expect that when states form alliances with more unreliable states, these same unreliable states will – on average – provide greater military capabilities to the alliance in order to compensate for the added risk of default through expected returns.

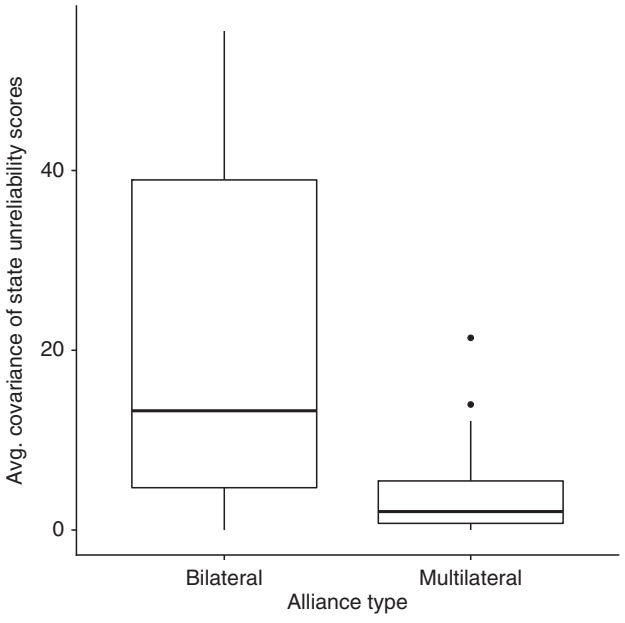

Figure 4. Covariance in the violations that comprise states' unreliability over time in bilateral versus multilateral offensive and defensive alliances

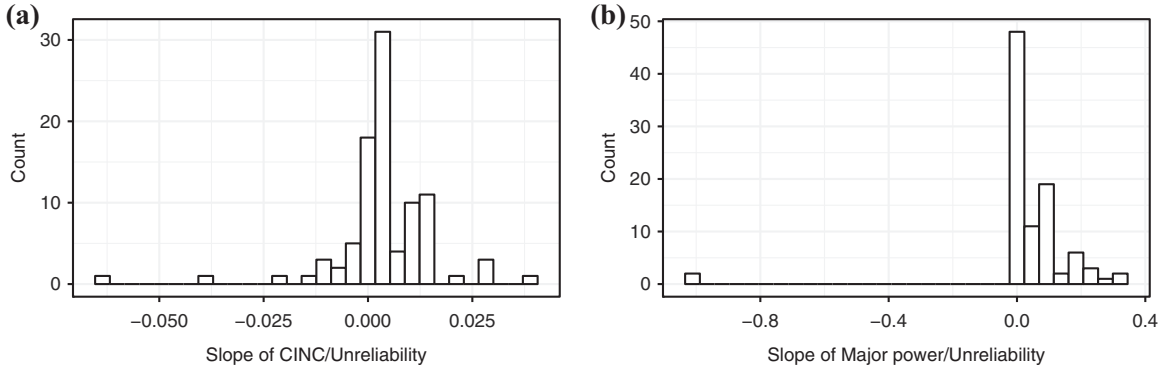

Figure 5. Histograms of beta coefficients for the estimated relationship between unreliability and the return to alliance members, measured using CINC scores and major power status

To test this hypothesis, we fit an OLS regression between unreliability and two different operationalizations of military capabilities: *CINC scores* and *Major power* status. Figure 5 displays a histogram of every beta coefficient from each alliance in our analysis using each measure. As expected, the overwhelming majority of the coefficients are positive. For *CINC score*, we can reject the null hypothesis that the relationship in the median alliance is less than or equal to zero using either a Wilcoxon signed-ranked test ($p = 7.31 \times 10^{-7}$) or a sign test ($p = 5.19 \times 10^{-8}$). For *Major power*, we can reject the null hypothesis that the relationship in the median alliance is less than or equal to zero using either a Wilcoxon signed-ranked test ($p = 6.66 \times 10^{-8}$) or a sign test ($p = 6.53 \times 10^{-11}$). Taken together, we believe this provides strong evidence towards our theory, as risk appears strongly associated with higher expected return in multilateral alliances.

When combined with our finding from Hypothesis 1 that unreliable allies are more likely to enter a multilateral alliance, the results for both aspects of Hypothesis 2 provide evidence that offensive and defensive multilateral alliances are constructed in a way that helps manage and offset the risk posed by allies with worse reputations. Consistent with portfolio theory, the composition of multilateral alliances appears to be structured to reduce the covariance in alliance risk and to offset the risk of default through greater returns in the form of military contributions to the alliance.

Alliance unreliability and the number of bilateral alliances

Hypothesis 3 expects that states in bilateral alliances with unreliable partners will be more likely to hedge against the risk of that partner defaulting by signing a greater number of bilateral alliances with other states. To test this expectation, we estimated a model where the dependent variable is the total number of bilateral alliances a state has in a given year (see Figure 1). Our main independent variable is the *Unreliability* score of a state's most unreliable partner. These partners represent the weakest link in a state's bilateral alliance portfolio. The regression equation is:

$$\#Alliances_{it} = \alpha_i + \beta_1 max[partner\ unreliability\ scores]_{it} + \beta X_{it}$$

We estimated the equation using an OLS regression with White's standard errors, clustered by state. Each observation is a state-year where state i was in at least one bilateral alliance in year t. In addition to controls that we used when testing Hypothesis 1, we added a state's own *Unreliability* score to control for the fact that both the number of alliances held (the dependent variable) and a state's propensity to sign alliances with unreliable states (the independent variable) might each be a function of the state's own unreliability. In two specifications, we also add a state-level fixed effect to control for individual states' propensity to hold a certain number of alliances over time, as well as a cubic time polynomial to control for any temporal changes in states' general propensity to hold a certain number of alliances.

The results are shown in Table II. As predicted, there is a positive association between how unreliable a state's most unreliable partner is, and the number of bilateral alliances signed by a state.[18] Substantively, if a state's least reliable ally has one more violation, on average, it

[18] The R^2 is 0.56 for the first two models and 0.86 for the second two models. Removing our unreliability variable reduces the R^2 by 37 percentage points (p < 2.2 x 10^–16 according to a Wald test).

Table II. Total number of offensive and defensive bilateral alliance commitments held by a state as a function of the unreliability score of its least reliable bilateral alliance partner

	DV: Number of defensive bilateral alliances			
	(1)	*(2)*	*(3)*	*(4)*
Unreliability of worst ally	0.145**	0.148**	0.118**	0.118**
	(0.034)	(0.035)	(0.024)	(0.024)
Own unreliability	0.445**	0.444**	0.213	0.215
	(0.057)	(0.058)	(0.188)	(0.188)
MID in last year	–0.384*		–0.032	
	(0.172)		(0.069)	
MID in last 5 years		–0.123		–0.091
		(0.183)		(0.100)
Major power	–0.274	–0.384	0.927	0.923
	(0.757)	(0.771)	(0.549)	(0.542)
Democracy	–0.135	–0.100	–0.602**	–0.601**
	(0.443)	(0.441)	(0.223)	(0.223)
Anocracy	–0.359	–0.359	–0.164	–0.165
	(0.197)	(0.198)	(0.171)	(0.170)
CINC	10.094*	9.588*	10.838	10.811
	(4.468)	(4.366)	(11.348)	(11.331)
State fixed effects			X	X
Time polynomial			X	X
Observations	1,832	1,832	1,832	1,832
R^2	0.561	0.555	0.867	0.868
Adjusted R^2	0.559	0.554	0.861	0.861
Residual std error	1.430 (df = 1,824)	1.440 (df = 1,824)	0.804 (df = 1,746)	0.803 (df = 1,746)
F statistic	333.163***	325.504***	134.320***	134.534***
	(df = 7; 1,824)	(df = 7; 1,824)	(df = 85; 1,746)	(df = 85; 1,746)

Standard errors in parentheses. $^{**}p < 0.01$, $^{*}p < 0.05$, $\dagger p < 0.1$.

increases the size of the state's bilateral alliance portfolio by about 5%. This evidence suggests that states in bilateral alliances with unreliable partners do appear to hedge against the risk of that partner defaulting by signing a greater number of bilateral alliances.[19] However, we note that although this result is fully consistent with the expectation derived from the theory, it does not fully exclude the alternative possibility that states with many bilateral alliances may be more likely to have at least one alliance partner with a poor reputation. Therefore, this result should be interpreted with some caution.

Conclusion

We conclude by returning to the questions motivating this article. Does international reputation matter? If so, why do states sometimes sign alliances with demonstrably unreliable partners? This is puzzling because a poor reputation should signal several risks outlined at the outset of this article. As a result, one might expect demonstrably unreliable states to be largely precluded from gaining future alliances. Thus, the fact that unreliable states frequently form new alliances might suggest that reputation only has a weak effect on alliance politics.

Yet, despite the arguments above, this article demonstrates that international reputation actually matters in ways that are often overlooked. While previous work has shown how perceptions of alliance reliability affect the quantity and quality of future alliances, we argued that alliance reliability should also affect how future alliance portfolios are structured. Specifically, we demonstrate that more unreliable alliance partners appear to cause states to rationally diversify their overall alliance portfolios in order to hedge against a greater risk of abandonment.

This article is also the first to identify and characterize variation in alliance portfolio-structure as a broader theoretical and empirical puzzle. To date, there does not

[19] In Table A3 in the Online appendix, we show that the results are also robust to an alternate measure of reliability used by LeVeck & Narang (2017a).

exist any theoretically coherent models to explain why states prefer cooperating with allies through multilateral alliance portfolios over bilateral alliance portfolios. Nor has there been a particularly clear explanation for why states sometimes sign a large number of bilateral alliances. As it turns out, states' preferences for multilateral versus bilateral alliances systematically vary with their need to manage the risk posed by unreliable allies. Similarly, the risk posed by unreliable partners may drive states to sign a larger number of bilateral alliances.

Together, these arguments suggest that previous work on alliances was correct to see alliance agreements as analogous to bonds (Conybeare, 1992; LeVeck & Narang, 2017a; Mattes, 2012). What previous work missed, however, is that investors do not simply manage risk and return through the creation and acquisition of single financial instruments. Instead, investors also manage the structure of aggregate portfolios (Markowitz, 1959). Our results show that states use similar techniques when managing the risk and returns of alliance agreements. This article therefore follows previous research by showing how very general theoretical frameworks, developed outside of political science, can help in politics (see Fearon, 1995, 1998; Horowitz & Narang, 2014; Lake, 1999; Narang, 2015b, 2018b).

Replication data

The dataset, codebook, and do-files for the empirical analysis in this article, as well as the Online appendices, can be found at http://www.prio.org/jpr/datasets.

Acknowledgments

We thank Michaela Mattes, Brett Ashley Leeds, Paul Poast, Brett Benson, David Lake, Erik Gartzke, Branislav Slantchev, Miles Kahler, Mark Crescenzi, Todd Sechser, James Morrow, James Long, Robert Trager, Daniel Maliniak, Robbie Narang, and the participants in the International Relations Workshop at the University California, San Diego. 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 Annual Meetings of the American Political Science Association in 2011 and 2018. We also thank the editorial team at the *Journal of Peace Research*, along with many peer reviewers.

References

Akerlof, George A (1970) The market for 'lemons': Quality uncertainty and the market mechanism. *Quarterly Journal of Economics* 84(3): 488–500.

Altfeld, Michael F & Bruce Bueno de Mesquita (1979) Choosing sides in wars. *International Studies Quarterly* 23(1): 87–112.

Beckley, Michael (2015) The myth of entangling alliances: Reassessing the security risks of US defense pacts. *International Security* 39(4): 7–48.

Bennett, D Scott & & Allan Stam (2000) EUGene: A conceptual manual. *International Interactions* 26(2): 179–204.

Bennett, D Scott & & Allan Stam (2003) *The Behavioral Origins of War*. Ann Arbour, MI: University of Michigan Press.

Benson, Brett V (2012) *Constructing International Security: Alliances, Deterrence, and Moral Hazard*. Cambridge: Cambridge University Press.

Bueno de Mesquita, Bruce (1975) Measuring systemic polarity. *Journal of Conflict Resolution* 19(2): 187–216.

Bueno de Mesquita, Bruce (1978) Systemic polarization and the occurrence and duration of war. *Journal of Conflict Resolution* 22(2): 241–267.

Bueno de Mesquita, Bruce (1981) Risk, power distributions, and the likelihood of war. *International Studies Quarterly* 25(4): 541–568.

Bueno de Mesquita, Bruce (1988) The contribution of expected utility theory to the study of international conflict. *Journal of Interdisciplinary History* 18(4): 629–652.

Carter, David B & Curtis S Signorino (2010) Back to the future: Modeling time dependence in binary data. *Political Analysis* 18(3): 271–292.

Cha, Victor D (2000) Abandonment, entrapment, and neoclassical realism in Asia: The United States, Japan, and Korea. *International Studies Quarterly* 44(2): 261–291.

Cha, Victor D (2010) Powerplay: Origins of the US alliance system in Asia. *International Security* 34(3): 158–196.

Conybeare, John A (1992) A portfolio diversification model of alliances: The triple alliance and triple entente, 1879–1914. *Journal of Conflict Resolution* 36(1): 53–85.

Crescenzi, Mark JC; Jacob D Kathman, Katja B Kleinberg & Reed M Wood (2012) Reliability, reputation, and alliance formation. *International Studies Quarterly* 56(2): 259–274.

Dawes, Robyn M (1979) The robust beauty of improper linear models in decision making. *American Psychologist* 34(7): 571–582.

Fang, Songying; Jesse C Johnson & Brett Ashley Leeds (2014) To concede or to resist? The restraining effect of military alliances. *International Organization* 68(4): 775–809.

Fearon, James D (1995) Rationalist explanations for war. *International Organization* 49(3): 379–414.

Fearon, James D (1997) Signaling foreign policy interests: Tying hands versus sinking costs. *Journal of Conflict Resolution* 41(1): 68–90.

Fearon, James D (1998) Bargaining, enforcement, and international cooperation. *International Organization* 52(2): 269–305.

Fuhrmann, Matthew & Todd S Sechser (2014) Signaling alliance commitments: Hand-tying and sunk costs in

extended nuclear deterrence. *American Journal of Political Science* 58(4): 919–935.

Gartzke, Erik & Kristian Skrede Gleditsch (2004) Why democracies may actually be less reliable allies. *American Journal of Political Science* 48(4): 775–795.

Gibler, Douglas M (2008) The costs of reneging: Reputation and alliance formation. *Journal of Conflict Resolution* 52(3): 426–454.

Hafner-Burton, Emilie M; Brad L LeVeck & David G Victor (2017) No false promises: How the prospect of non-compliance affects elite preferences for international cooperation. *International Studies Quarterly* 61(1): 136–149.

Hemmer, Christopher & Peter J Katzenstein (2002) Why is there no NATO in Asia? Collective identity, regionalism, and the origins of multilateralism. *International Organization* 56(3): 575–607.

Horowitz, Michael & Neil Narang (2014) Poor man's atomic bomb? Exploring the relationship between 'weapons of mass destruction'. *Journal of Conflict Resolution* 58(3): 509–535.

Johnson, Jesse C & Brett Ashley Leeds (2011) Defense pacts: A prescription for peace? *Foreign Policy Analysis* 7(1): 45–65.

Kamp, Karl-Heinz (1995) The folly of rapid NATO expansion. *Foreign Policy* 98(Spring): 116–129.

Koremenos, Barbara; Charles Lipson & Duncan Snidal (2001) The rational design of international institutions. *International Organization* 55(4): 761–799.

Lai, Brian & 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: American Foreign Policy in its Century*. Priceton, NJ: Princeton University Press.

Leeds, Brett Ashley & Burcu Savun (2007) Terminating alliances: Why do states abrogate agreements? *Journal of Politics* 69(4): 1118–1132.

Leeds, Brett Ashley; Andrew G Long & Sara McLaughlin Mitchell (2000) Reevaluating alliance reliability: Specific threats, specific promises. *Journal of Conflict Resolution* 44(5): 686–699.

Leeds, Brett Ashley; Andrew G Long & Sara McLaughlin Mitchell (2003) Alliance reliability in times of war: Explaining state decisions to violate treaties. *International Organization* 57(4): 801–827.

Leeds, Brett Ashley; Michaela Mattes & Jeremy S Vogel (2009) Interests, institutions, and the reliability of international commitments. *American Journal of Political Science* 53(2): 461–476.

LeVeck, Brad & Neil Narang (2017a) How international reputation matters: Revisiting alliance violations in context. *International Interactions* 43(5): 792–821

LeVeck, Brad & Neil Narang (2017b) The democratic peace and the wisdom of crowds. *International Studies Quarterly* 61(4): 867–880.

Markowitz, Harry M (1959) *Portfolio Selection: Efficient Diversification of Investments*. New York: John Wiley & Sons.

Mattes, Michaela (2012) Reputation, symmetry, and alliance design. *International Organization* 66(4): 679–707.

Mehta, Rupal & Neil Narang (2018) Do nuclear umbrellas increase the diplomatic influence of client states? A theory and empirical evidence. Paper presented at the American Political Science Association, Boston, September.

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 (1993) Arms versus allies: Trade-offs in the search for security. *International Organization* 47(2): 207–233.

Morrow, James D (2000) Alliances: Why write them down? *Annual Review of Political Science* 3: 63–83.

Narang, Neil (2014) Humanitarian assistance and the duration of peace after civil war. *Journal of Politics* 76(2): 446–460.

Narang, Neil (2015a) Forgotten conflicts? Need versus political priority in the allocation of humanitarian assistance across conflict areas. *International Interactions* 42(1): 189–216.

Narang, Neil (2015b) Assisting uncertainty: How humanitarian aid can inadvertently prolong civil war. *International Studies Quarterly* 59(1): 184–195.

Narang, Neil (2018a) When nuclear umbrellas work: Signaling credibility in security commitments through alliance design. Paper presented at the American Political Science Association, Boston, 30 August–2 September.

Narang, Neil (2018b) Biting the hand that feeds: An organizational theory explaining attacks against aid workers in civil war, 1997–2010. Paper presented at the UCSB Orfalea Center, Santa Barbara, 19 January.

Narang, Neil & Rupal Mehta (forthcoming) The unforeseen consequences of extended Deterrence: Moral hazard in a nuclear client. *Journal of Conflict Resolution*. https://doi.org/10.1177/0022002717729025.

Ostrom, Charles W, Jr, & John H Aldrich (1978) The relationship between size and stability in the major power international system. *American Journal of Political Science* 22(4): 743–771.

Palmer, Glenn; Vito D'Orazio, Michael Kenwick & Matthew Lane (2015) The MID4 data set: Procedures, coding rules, and description. *Conflict Management and Peace Science* 32(2): 222–242.

Sagan, Scott D (1986) 1914 revisited: Allies, offense, and instability. *International Security* 11(2): 151–175.

Signornio, Curtis & Jeffrey Ritter (1999) Tau-b or not tau-b: Measuring the similarity of foreign policy positions. *International Studies Quarterly* 43(1): 115–144.

Simon, Michael W & 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 & 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.

Siverson, Randall M & Joel King (1980) Attributes of national alliance membership and war participation, 1815–1965. *American Journal of Political Science* 24(1): 1–15.

Smith, Alastair (1995) Alliance formation and war. *International Studies Quarterly* 39(4): 405–425.

Snyder, Glen H (1984) The security dilemma in alliance politics. *World Politics* 36(4): 461–495.

Spence, Michael (1973) Job market signaling. *Quarterly Journal of Economics* 87(3): 355–374.

Stoll, Richard J (1984) Bloc concentration and the balance of power: The European major powers, 1824–1914. *Journal of Conflict Resolution* 28(1): 25–50.

Stoll, Richard J & Michael Champion (1985) Capability concentration, alliance bonding, and conflict among the major powers. In: Alan Ned Sabrosky (ed.) *Polarity and War: The Changing Structure of International Conflict*. Boulder, CO: Westview, 67–94.

Waltz, Kenneth N (2000) NATO expansion: A realist's view. *Contemporary Security Policy* 21(2): 23–38.

NEIL NARANG, b. 1982, PhD in Political Science (University of California, San Diego, 2012); Assistant Professor of Political Science at the University of California, Santa Barbara (2012–); Council on Foreign Relations International Affairs Fellow serving as a Senior Advisor in the Office of the Secretary of Defense for Policy (2015–16).

BRAD L LEVECK, b. 1982, PhD in Political Science (University of California, San Diego, 2012); Assistant Professor of Political Science at the University of California, Merced (2014–).