

Testing Tripwire Theory Using Survey Experiments

Paul Musgrave

Assistant Professor of Political Science

University of Massachusetts Amherst

musgrave@umass.edu

Steven Ward

Lecturer in the Department of Politics and

International Studies

University of Cambridge

sw986@cam.ac.uk

September 2021

Abstract: Many scholars and policymakers assume that attacks on forward deployed U.S. troops—“tripwires”—are likely to prompt strong domestic political support for escalation against the attacker. This conjecture informs policy and has deep theoretical roots, yet it is undertheorized and largely untested. We identify and develop two theoretical mechanisms – reputation and revenge – capable of explaining why attacks on forward deployed troops might prompt support for escalation, even though prior research shows that casualties suffered during a conflict reduce support for intervention. We then use two survey experiments to examine whether and how attacks on forward deployed U.S. troops influence Americans’ support for intervention. We find that hypothetical attacks on contingents of troops deployed overseas increase support for escalation only modestly and in ways that better reflect demands for revenge rather than concerns about reputation. Our findings imply that confident assessments that forward deployed troops serve as strong pre-commitment devices need to be tempered, pending further theoretical and empirical analysis.

Introduction

For decades, the United States has stationed troops abroad. Today, American soldiers are deployed in 800 locations and more than 70 countries that host forces ranging in size from a handful to tens of thousands of personnel (Cooley and Nexon 2013). Former President Donald Trump's skepticism of stationing troops abroad prompted sporadic efforts to withdraw U.S. troops from countries such as South Korea and Germany. His attempts occasioned a new round of a longstanding debate about whether the presence of U.S. troops abroad, and the prospect of American casualties resulting from foreign attacks, increases the credibility of U.S. pledges to defend partners (Bandow 2018; Gordon 2020).

This ongoing policy debate implicates a largely unexamined conventional scholarly wisdom. Many theorists have described forward deployments – contingents of troops stationed in foreign countries – as potentially serving as hands-tying signals (Fearon 1997). Such deployments may help leaders establish deterrent credibility because of an anticipated “tripwire” effect through which an attack on those troops would raise the likelihood of escalation. One important pathway through which the tripwire effect could function is through the prospect of domestic political pressure on leaders to respond after an attack. This pressure, according to theorists and policy analysts, is supposed to constrain leaders to follow through on promises to act on behalf of the threatened country. Many analysts have long assumed that this effect is strong, particularly in the case of reactions in the United States to attacks on U.S. forward deployed troops—potentially strong enough to even make credible leaders’ threats to escalate conventional conflicts to the nuclear level. At the height of the Cold War, for example, Glenn Snyder conjectured that if Soviet forces attacked U.S. and allied troops in Europe, “strong emotions favoring a nuclear retaliatory response will be generated in Western public opinion” because “an attack on such forces would implicate the honor, prestige, and other emotional values of” countries deploying troops abroad “to a greater degree than a mere treaty commitment” (G. H. Snyder 2015, 130–31). Many contemporary analysts and policymakers continue to agree that the prospect of U.S. escalation in case of an attack on U.S. troops simultaneously reassures allies and deters adversaries, even if the number of U.S. troops deployed is insufficient to shift the balance of battlefield capabilities. Even opponents of current U.S. foreign policy embrace the same assessment of the strength of public opinion and the potential tripwire effect. Instead of viewing such a mechanism as a means of carefully calibrating deterrence, however, they fear an enraged public pushing Washington to become entangled in unnecessary conflict and therefore counsel against such deployments.

There are reasons to question this conventional wisdom. Recent work has yielded mixed evidence about the ability of similar mechanisms, like audience costs generated by public threats, to tie leaders' hands (Chaudoin 2014; Davies and Johns 2013; Kertzer and Brutger 2016; Levendusky and Horowitz 2012; Nomikos and Sambanis 2019; J. Snyder and Borghard 2011; Tomz 2007; Trager and Vavreck 2011). Research showing that casualties incurred during conflicts weaken public support for war (Kriner and Shen 2012; Mueller 1973) also raises doubts: if public casualty sensitivity during a conflict reduces support for war, then the contention that pre-conflict attacks produce demands for intervention may be dubious.

Both theorists and policymakers would therefore benefit from a careful assessment of the relationship between U.S. forward deployment and U.S. public support for escalation. For theorists, understanding whether attacks on forward-deployed troops produce demands for escalation could clarify how different kinds of pre-commitment devices work—in particular, whether Snyder's conjecture that blood is thicker than parchment is empirically accurate. Policymakers and the public, both in the United States and in countries reliant on U.S. security commitments, would thus be well-served to know more about whether and how forward deployment functions as a pre-commitment device, and how strong the effect it produces through the public opinion channel is (Quek 2021). Attention to the role of public opinion in producing a tripwire effect is all the more critical given recent experimental work suggesting that public opinion can sway elected and unelected officials regarding the use of force (Lin-Greenberg 2021; Tomz, Weeks, and Yarhi-Milo 2020).

This paper advances the field's understanding of these issues in two ways. First, we establish that conventional theories about the production of tripwire effects often invoke, explicitly or implicitly, a crucial role for domestic political pressure arising from an attack. We argue that, although these explanations play a central role in many accounts of tripwire effects, they are nonetheless underspecified. We identify two mechanisms alluded to in the literature that might plausibly connect attacks on troops deployed abroad to public support for intervention: revenge and concern for reputation. Both mechanisms potentially explain the central proposition of tripwire theories that rely on public opinion: why pre-conflict casualties might *increase* support for intervention, even though studies consistently show that casualties incurred *during* conflict erode support. Accordingly, we clarify how these mechanisms could operate and derive expectations about how we would know if one, both, or neither were operating. Second, we employ two original survey experiments to investigate whether (and, if so, how and under what conditions) attacks on forward-deployed U.S. troops influence Americans' attitudes toward escalatory intervention and, crucially, to

assess whether the revenge and reputation mechanisms drive responses to attacks on troops deployed abroad.

Our findings suggest that pre-conflict casualties *can* increase public support for escalation, and that they do so in ways that better reflect a desire for revenge than a concern about the state's reputation for credibility. We also find, however, that the ability of forward deployment to produce public support for intervention is limited. First, the effect of attacks on forward deployed forces appears much smaller than other aspects of a situation. Second, the effect is substantial only when very large numbers of casualties are incurred. Third, such attacks most strongly increase support for only moderately escalatory responses, rather than highly escalatory ones.

In sum, while we provide the first experimental evidence showing that attacks on forward deployed troops can increase domestic support for intervention, our findings suggest caution for analysts and policymakers who imagine that those troops are capable of prompting a strong tripwire effect that virtually guarantees further intervention. Forward deployed troops are not obviously better able to produce Americans' public support for escalation than less risky and costly pre-commitment devices (like formal alliances). Indeed, the relatively small increases in support for escalation that come from even very large numbers of casualties imply that both proponents and opponents of forward postures need to temper their assessments of how well troop deployments tie leaders' hands.

We are careful to note that our findings do not necessarily imply that forward deployment (of even small contingents of troops) does not deter.¹ Our analysis aims at clarifying and testing the relationship between forward deployment and public support for escalation in the United States because of the importance of claims about the capacity of forward deployed troops to serve as tripwires and the prominence of U.S. forward deployments explicitly or implicitly meant to serve in that role. Although our findings suggest that American public opinion likely cannot support a strong tripwire effect, a different mechanism still could produce such an effect. We leave the identification and empirical analysis of these alternatives to future work.

Forward Deployment and the Tripwire Effect

Forward deployment refers to the stationing of military personnel overseas in areas of special strategic interest, a practice not limited to but most visible in post-1945 U.S. foreign policy (Schmidt

¹ Although others have recently made this claim; see Reiter and Poast (2021).

2020; Vanaga and Rostoks 2018). Forward deployment takes many forms and can serve many purposes. Large numbers of troops may be stationed abroad to alter the local balance of power between an ally and a potential adversary, as with U.S. troops in South Korea. Most deployments, however, are much smaller. These may be intended to bolster security and stability within the host country, support other U.S. deployments elsewhere, increase the professionalism and capabilities of host state militaries, or improve battlefield performance through joint training and exercises.

As Reiter and Poast (2021) observe, forward deployments are often meant to deter potential adversaries. This is the case for both “tripwire” deployments and for larger deployments that substantially shift the local balance of military capabilities. We focus on troops that can produce what we term the “tripwire effect:” the ability of a contingent of forward deployed troops to prompt dynamics that increase the probability of further intervention by the deploying state if the forward deployed contingent or its host country is attacked. Importantly, this effect might operate regardless of the size (or even the purpose) of the deployment. Depending on how the tripwire effect is conceptualized, it could be triggered *whenever* American troops come under attack, no matter their numbers or their reason for being overseas. Indeed, the main purported advantage of what Reiter and Poast call “tripwire” deployments is that they hold out the possibility of deterring (by producing strong expectations of escalation in case of an attack) without generating security dilemma dynamics (because the deployment can be small). There is no reason that such a tripwire effect could be produced only by contingents explicitly designated as tripwire deployments. Attacks on much larger deployments could also plausibly produce pressure for further escalation. (During the Cold War, for example, large U.S. deployments in Berlin and South Korea were often described as serving tripwire functions.) What is distinctive about tripwire deployments is that they *depend solely* on the tripwire effect for their deterrent force. Larger deployments might be able to deter *both* by changing beliefs about the balance of military capabilities and by changing beliefs about the likelihood of escalation. Because tripwire deployments are not large enough to deter by altering potential adversaries’ beliefs about the relative military balance, they can only deter by altering beliefs about the likelihood of escalation in case of an attack.

Forward Deployment, Signaling, and Extended Deterrence

The tripwire effect relates to two concepts that have received substantial attention from IR theorists: extended deterrence and signaling. Extended deterrence refers to the use of deterrent threats to prevent a potential adversary from harming a third party. Extended deterrent threats are typically

understood as more difficult to make credible than primary deterrent threats (preventing harm to one's own country). During the Cold War, observers frequently assumed that a Soviet nuclear attack on American soil would have prompted U.S. nuclear retaliation. Doubts arose about whether an attack on a U.S. ally would trigger the same response, given that Soviet retribution could subsequently threaten U.S. cities. Fears that deterrence could fail in this manner percolated at the highest levels. In 1961, U.S. President John F. Kennedy expressed frustration to French President Charles de Gaulle that de Gaulle himself "had asked whether we [Americans] would be ready to trade New York for Paris. If the General himself, who has worked together with the United States for so long, could question American firmness, Mr. Khrushchev can question it also."²

One solution to this problem is to implement policies that credibly signal commitment. Signaling refers to the transmission of previously private or uncertain information about a state's preferences, intentions, or strategies. Exploring how to generate credible commitments via signaling stimulated an influential generation of scholarship during the Cold War (Schelling 1960, 1966; G. H. Snyder 2015; Trachtenberg 1989). A later generation employs formal modeling to articulate different signaling mechanisms, various kinds of empirical methods to assess the effectiveness of different kinds of signals, or both (Fearon 1994, 1997; Quek 2021).

This scholarship has established a useful distinction between two types of signals: sunk-cost signals and tying-hands signals. Sunk-cost signaling occurs when a committed state visibly invests resources that an uncommitted state would not. This reveals how much the state is willing to pay to defend the commitment, but does not change that state's calculations if the commitment is challenged. By contrast, tying-hands signals involve arrangements that may be established relatively cheaply, but that increase the costs of backing down from a commitment if challenged. In theory, tying-hands signals produce "audience costs," the political cost that a state's leadership would pay – imposed primarily by a domestic audience – if it failed to redeem its promise. These kinds of signals thus make it more difficult for a leader to renege on a commitment when challenged. Fearon highlights three examples of hands-tying signals: a formal alliance, a public threat, and the deployment of troops in a "threatened area" (Fearon 1994, 1997, 70). In theory, all three enhance

² Memorandum of Conversation, Paris, May 31, 1961, *Foreign Relations of the United States, 1961-1963*, Volume XIV, Berlin Crisis, 1961-1962, Document 30, <https://history.state.gov/historicaldocuments/frus1961-63v14/d30>.

credibility because failing to honor the commitment would lead to domestic public disapprobation, jeopardizing the leader's hold on power.

Much research over the past two decades has explored audience costs, focusing primarily on testing the claim that public threats generate domestic pressure to stand firm. Some researchers report results broadly supportive of the thesis. Tomz (2007) finds experimental evidence that leaders pay a domestic political price after failing to follow through on threats. Trager and Vavreck (2011) similarly use experiments to uncover evidence that presidents can manipulate audience costs before hostilities to make their threats more credible (see also Kertzer and Brutger 2016; Levy et al. 2015). Davies and Johns (2013) find in their experiments that contextual audience costs exist (especially among the politically engaged) in the United Kingdom, but that prime ministers have limited freedom in manipulating those costs.

Some analysts dissent. Qualitative research has been skeptical that audience costs influence decisionmaking (Burns and Stravers 2020; Reiter and Poast 2021; J. Snyder and Borghard 2011). Other work has also raised objections. Chaudoin (2014) finds that respondents' preferences over policy matter more than perceptions of inconsistency in driving assessments of leaders who deviate from prior commitments. Nomikos and Sambanis (2019) challenge experiments that follow Tomz's (2007) lead on theoretical and methodological grounds. They argue that such work has conflated audience costs with "incompetence costs" caused by audiences' perception that leaders failed to reach the optimal outcome, and that audience costs may therefore be exaggerated in those experiments. Along similar lines, Casler and Clark (Forthcoming) suggest that – at least in the context of trade disputes – leaders pay domestic costs for threatening tariffs, but not for backing down from those threats. Gelpi and Grieco (2015) analyze U.S. presidential legislative efforts to argue that competency costs matter more than traditional audience costs. Others contend that, despite earlier evidence that audience costs are nonpartisan (Levendusky and Horowitz 2012), increasing polarization and audience fragmentation in a social media age raises questions about whether they continue to matter today (Baum and Potter 2019).

The Tripwire Effect and Public Opinion

While public threats have so far dominated research into hands-tying signals, Fearon argued that forward deployment was also a means of generating domestic audience costs.³ This categorization echoes a common and longstanding assumption about the role of forward-deployed troops that is closely associated with Thomas Schelling. Schelling sought to resolve the dilemmas of Cold War deterrence—especially fears that the United States would “decouple” itself from European security—by asserting that attacks on American troops stationed in Europe could trigger a U.S. response even in the face of potential nuclear retaliation. He illustrated his theory in a classic passage about the Berlin Brigade, a small contingent of U.S. forces deployed in West Berlin:

The garrison in Berlin is as fine a collection of soldiers as has ever been assembled, but excruciatingly small. What can 7,000 American troops do, or 12,000 Allied troops? Bluntly, they can die. They can die heroically, dramatically, and in a manner that guarantees that the action cannot stop there. (Schelling 1966, 47)

Like the modest deployments at the heart of Fearon’s hands-tying mechanism, this “excruciatingly small” force did *not* materially change the military balance, but nevertheless *could* deter the Soviets through the anticipation of its nearly automatic impact on further escalation. More recent analysis makes similar claims about the function of troops deployed abroad (see Brooks and Wohlforth 2016, 96).

Analysts do not always specify exactly *how* the tripwire effect is supposed to function. Those who do frequently posit a link between attacks on troops stationed overseas and public support for escalation. This is evident in Snyder’s (2015, 130–31) claim about “Western public opinion” demanding escalation in case of an attack on allied troops in Europe. Rovner and Talmadge (2014, 554) note that “light presence” deployments “often deliberately create a ‘tripwire’ ensuring that regional aggression will necessarily entail early engagement with the hegemon” and illustrate their argument with reference to Cold War-era U.S. deployments in Germany, which worked by “guaranteeing public support for European allies” in case of a Soviet attack. Fuhrmann and Sechser (2014, 923) argue that forward deployed nuclear forces may be “militarily superfluous” but nevertheless serve an important role because their destruction during an initial attack on the partner might produce “domestic political pressure for the patron to enter the war wholeheartedly.” Slack

³ Tomz and Weeks (Forthcoming) have recently published the first experimental study on alliances (another potential hands-tying signal) and support for war.

(2018, 25, 31) argues that “the public death of fellow nationals serving the flag ignites and engages the national chauvinist portion of the domestic audience who then demand that the state uphold its honor,” and that “the public loss of military forces generates a large political cost for backing down”. A RAND study asserts that forward deployed troops “deter potential adversaries” because an attack on them might “engage the U.S. public, which could add additional pressure on policymakers to respond with a larger U.S. force” (O’Mahony et al. 2018, 24). And Stephen Saideman recently explained to a national newspaper audience that if Canadian troops deployed in the Baltics died, “their deaths would lead to an automatic response – politicians can’t just let their troops die without a response – which then would lead to an escalation and on and on, so that [NATO] deaths serve to deter the Russians” (MacKinnon 2017).

The notion that attacks on troops deployed abroad prompt public support – and even demands – for escalation thus forms a common theoretical justification for claims about the ability of even small contingents of forward deployments to serve as strong pre-commitment devices. Surprisingly, though, this contention has not received direct, systematic empirical investigation despite good reasons to question its validity. Evidence from related debates about casualties and public support for war appears to contradict the idea that attacks on troops deployed abroad should lead to demands for escalation. Scholars have consistently found that the U.S. public responds to casualties during a conflict by becoming less supportive of continuing the intervention (Fazal 2021; Gartner 2008; Gartner and Segura 1998; Kriner and Shen 2012; Mueller 1973). In other words, during an *ongoing* conflict, mounting U.S. troop deaths lead audiences to demand withdrawal, and harm the domestic popularity of the incumbent president (Geys 2010)—just the opposite of how the tripwire effect is supposed to influence domestic public opinion. While the pre-conflict context may differ in important ways from that of an ongoing conflict, these differences remain untheorized. It is thus not obvious that the same mechanism implicated in an ongoing conflict would not function prospectively in the pre-conflict context. If it did, the result might be a reverse tripwire effect – attacks on troops deployed overseas could actually generate opposition to escalation among domestic audiences.⁴

⁴ Indeed, even if one believes a tripwire effect is produced entirely through other mechanisms, such as international reputation, such concerns would warrant investigating the domestic public opinion channel to better understand the magnitude of public support for escalation.

Moreover, there have been many instances in which states hosting third-party troops – or the troops themselves – have been attacked without sparking irresistible public demands for escalation. These include North Korea’s seizing of the *U.S.S. Pueblo* in 1968 (1 sailor killed, 82 taken prisoner), downing of an EC-121 spy plane in 1969 (killing 31), and attack on soldiers clearing trees in the Demilitarized Zone in 1976 (killing two U.S. officers); the 1983 Beirut bombing of a Marine Corps barracks (killing 241 U.S. servicemembers); the 1994 killings of peacekeepers in Rwanda; and the 2020 Iranian attacks on U.S. troops in Iraq. Together, these examples suggest that attacks on troops deployed abroad might not always prompt the kind of strong public responses that the conventional wisdom assumes. For instance, the 1976 North Korean attack occurred during a U.S. presidential election, when one might expect the pressures of public opinion to be most intense, and yet President Gerald Ford responded with a show of force and a demand for a North Korean apology, not armed retaliation (Sander 2017).

In sum, this discussion suggests the need for greater theoretical development and empirical investigation of how attacks on forward deployed troops affect public attitudes toward intervention. Accordingly, in the remainder of this section, we identify and develop two theoretical mechanisms capable of explaining why attacks on forward deployed troops might generate demands for escalation, even though casualties incurred *during* a conflict reduce support for war. In the following section, we empirically assess those mechanisms using the same kinds of survey experimental methods that have been used productively to examine the strength of other kinds of pre-commitment devices.

Why Should Attacks on Forward Deployed Troops Produce Support for Escalation?

Though analysts commonly assert that attacks on forward deployed troops raise the probability of further intervention because they are likely to induce demands for escalation among domestic audiences, this represents only a partial theoretical logic. The unanswered question involves the relationship between attacks and public opinion. Why should the public be expected to react to attacks on troops deployed abroad by demanding escalation, especially when research overwhelmingly shows that it reacts to casualties incurred during conflicts by supporting de-escalation?

We identify and develop two answers to this question that are implicit in claims about the pre-commitment function of forward deployment. These take the form of two distinct mechanisms linking attacks on soldiers deployed overseas to changes in individual attitudes toward conflict with

the perpetrator of the attack. Importantly, there are theoretical reasons to expect both of these mechanisms to operate most strongly prior to the start of open conflict with the perpetrator of the attack, and more weakly (if at all) in the context of an ongoing conflict.

Outrage and revenge

The first mechanism is the activation of the emotion of outrage, leading to a desire for revenge. Certain kinds of provocations – potentially including unexpected and shocking attacks on troops deployed overseas – may produce in audiences “an individually felt emotional experience of anger” (Hall 2017, 5). Anger may alter preferences, at least in the short-run, in ways that change attitudes toward favoring military escalation (Shandler et al. 2021). In particular, Hall (6-7) notes that outrage can make actors less sensitive to risk. It can also make them demand immediate action to rectify or respond to the perceived wrong. Overall, according to Hall (8), outrage produces “an inclination toward rash, aggressive action.” Stein (2015) notes that the demand for revenge is a potentially important factor in attitudes toward conflict. She defines revenge as “the belief that wrongdoers *deserve* to be repaid for their crimes” (2015, 558), and empirically links “vengefulness” (conceptualized as a country-level population characteristic) to belligerence in foreign policy. By extension, it could be that attacks on forward deployed troops provoke outrage and demands for policies that retaliate against and punish the perpetrator. This mechanism appears implicitly in some prominent articulations of how forward deployment works as a pre-commitment device. Schelling emphasized the importance of the spectacle of the deaths of U.S. soldiers in Berlin (they had to die “heroically, dramatically”), while Snyder emphasized the “emotions” that would be activated by attacks on U.S. troops deployed in Europe. Though neither author so clearly specified the mechanism laid out here, the notions that an attack might serve as a public spectacle and lead to an emotional reaction make sense in the context of an explanation centering on public outrage and demands for revenge.

This mechanism is also capable of explaining why casualties suffered before a conflict begins should have a different effect than those incurred during a conflict. Prior to the start of open conflict, an attack – especially one resulting in casualties – may seem like a shocking violation of trust, values, and international law. A surprise attack might thus easily be interpreted as an act of treachery, thereby provoking outrage and leading to demands for retaliation. By contrast, once a conflict has begun, casualties are not as shocking because they occur within the institutional context of warfare. The state responsible for the casualties is already established as a military adversary, and

the practice of attempting to harm or kill the other side's troops is accepted as normal. Thus, unless evidence emerges of other forms of treachery – for instance, a violation of international humanitarian law – casualties should not necessarily provoke demands for escalation.

Reputation and credibility

The second mechanism involves the public's concern about the state's reputation for credibility in foreign policy. Attacks on troops deployed abroad might represent challenges to the commitments embodied by those deployments. If troops are deployed in ways that seem intended to deter an attack by an adversary, then observers might believe that an attack requires an escalatory response in order to reassert the credibility of that commitment, and perhaps the state's other commitments more broadly. Since the public might reasonably (if incorrectly) believe that the state's reputation for credibility should be defended, this logic could explain why attacks on certain kinds of foreign deployments might generate support for escalatory action. We stress that this is not a judgment about the credibility of a state's reputation by other states but rather about the expectations of a domestic audience about how other states will view their government's reputation.

Similar logic led Fearon to classify “tripwire” forces deployed near threatened areas as potential sources of audience costs. Fearon (1994, 580) argues that troop deployments “engage the national honor,” which raises the salience of crises in which they are implicated, thus “exposing leaders to risk of criticism or loss of authority if they are judged to have performed poorly by the relevant audiences.” He then suggests that “backing down after making a show of force is often most immediately costly for a leader because it gives *domestic* political opponents an opportunity to deplore the *international* loss of credibility, face, or honor” (581). While this does not fully specify a link between *attacks* on forward deployed troops and support for escalation, it reflects the same logic laid out above. Other authors make similar allusions to a reputational mechanism. Fearon is not alone in specifically arguing about the importance of domestic audiences' judgments about the anticipated reputational costs as separate from any actual shift in international reputation. Slack (2018), for example, invokes the significance of the public's “demand that the state uphold its honor” in the wake of an attack on troops deployed abroad.

Like the revenge mechanism, the reputational mechanism can explain why pre-conflict casualties might prompt support for escalation while casualties incurred during a conflict might lead to support for withdrawal. When an attack clearly challenges an international commitment to deter an adversary, backing down or failing to escalate may seem to harm the state's reputation for

credibility. After a conflict begins, the calculation changes. Concerns about the credibility of threats and promises are likely to recede: to the extent that they were implicated in the crisis leading up to the conflict, they should have been satisfied by the state's initial decision to escalate. In the context of an ongoing conflict, reputational concerns are thus likely to be superseded over time by concerns about cost and feasibility.

Hypotheses, Research Design, and Results

Having clarified the mechanisms that could produce tripwire effects through the public opinion channel, we proceed to explore how to test them. While research on forward deployment is increasingly common, prior studies have not directly investigated its function as a hands-tying signal. O'Mahony et al. (2018), for instance, use a large-N observational design to explore the association between forward deployed U.S. troops and militarized conflict. This study leaves unclear *how* forward deployment contributes to the empirical patterns the authors discover. In another observational study, Jakobsen and Jakobsen (2019) find that the presence of U.S. troops appears to make citizens of the host country less willing to fight to defend themselves. This study highlights one of the potential *costs* of forward deployment, but says nothing about the ability of those deployments to function in the way policymakers and analysts imagine they do: as pre-commitment devices.

Observational research designs, whether qualitative or quantitative, generally face challenges in assessing the effects of forward deployment because forward deployment is not random. States deploy troops to areas of the world that they value highly, which may make establishing credibility less challenging. More generally, patterns of deployment are influenced by factors that drive variation in the probability of conflict – like the perceived likelihood of challenges by potential aggressors. And forward deployments are rarely used on their own – rather, they typically appear as part of a package of tools (including alliance commitments and public threats) used together to deter adversaries. This makes it difficult to observationally identify the effect of deployments on outcomes of interest.

Moreover, the infrequency of attacks on forward deployed troops means that observational studies aimed at analyzing whether or not these actually produce public demands for escalation will inevitably confront a relative paucity of data. While it is possible that this fact *reflects* beliefs about the operation of the mechanisms developed above, it cannot be taken as *evidence* that these mechanisms are sound. The relative infrequency of attacks on forward deployed troops could easily be explained

by other factors which simultaneously co-vary with deployment patterns and independently influence the likelihood of armed challenges.

To overcome these kinds of obstacles, researchers investigating related signaling processes have long used survey experiments. We follow this example, using survey experiments that randomly vary information about forward deployed troops and attacks on them in order to investigate whether and how these factors influence attitudes toward conflict. Because the (at least implicit) subject of most debate regarding forward-deployed troops has been the United States, we focus on the reactions of U.S.-based respondents. We test these hypotheses using two types of online experiments: three parallel online vignette experiments that explore concrete, realistic, scenarios and conducted soon after the 2020 attacks on U.S. troops in Iraq and a separate a conjoint study intended to examine the effect of variation in the *presence* of troops in a way that realistic vignette designs preclude.⁵

Hypotheses

While both mechanisms developed above suggest that attacks on American troops deployed abroad should increase public support for escalation, they imply different patterns of evidence related to the features of attacks and the conditions under which they occur. We exploit these different implications to design a number of hypothesis tests that allow us to discriminate between their effects.

Casualties

First, the two mechanisms differ in their assessments of how casualties matter. The revenge mechanism implies that attacks that result in higher numbers of casualties should be more provocative and thus prompt higher levels of support for escalation. The reputational mechanism, by contrast, implies that what matters most is the challenge to the commitment embodied by the presence of troops. It does not clearly imply that support for escalation should be strongly linked to the number of casualties. We test the following hypothesis related to casualties:

⁵ For simplicity's sake, we focus on the classic case of forward-deployed ground troops, rather than deployments of naval or air forces. Though such deployments might serve similar purposes (Blankenship and Lin-Greenberg 2021), we leave this question to future research.

H1: The higher the number of casualties resulting from an attack, the greater the support for escalation.

Purpose of Deployment

Second, the two mechanisms differ in their evaluation of whether the stated purpose of a deployment should affect public opinion. The revenge mechanism does not imply that the purpose of a deployment should influence the reaction to an attack on deployed forces. As its name suggests, the reputation mechanism implies that attacks on deployments described as publicly committing a state to deter an adversary should produce stronger support for escalation. This leads to hypothesis 2:

H2: Attacks on deployments that are explicitly intended to deter an adversary should prompt stronger support for escalation than attacks on deployments serving other purposes.

Civilian/Military Status of Target Attacked

Third, the mechanisms differ in their assessments of the importance of the military or civilian status of the targets of an attack. The revenge mechanism does not imply that this distinction should matter: attacks against U.S. civilians (such as embassy staff) and attacks against U.S. military personnel might very well produce similar reactions. French strategist Pierre Gallois even suggested that “American school children would play this role [that of a ‘tripwire’] just as well as armored divisions” (Freedman and Michaels 2019, 359 note 6). Separately, Schelling (1960, 136, note 13) conjectured that the “wives and children” of U.S. soldiers may “have been a more persuasive commitment or ‘trip wire’ than the troops themselves”. By contrast, civilians living abroad are not typically understood as signals of resolve or commitment. Accordingly, the reputation mechanism implies that attacks on military personnel should increase support for escalation more than attacks on civilian personnel. We examine these contrasting expectations by testing hypothesis 3:

H3: Attacks on U.S. military personnel deployed abroad should prompt higher levels of support for escalation than do attacks on U.S. civilians abroad.

Nationality of Target Attacked

Another relevant distinction concerns the *nationality* of troops under attack. U.S. troops are often stationed overseas alongside local troops, military personnel from other allied militaries, or both.

The revenge mechanism implies that attacks on U.S. troops should prompt higher levels of support for escalation than attacks on foreign troops. The reputation mechanism is not as clear on this point. The significant element for the reputation mechanism is that a public commitment to a partner has been challenged. It is possible to imagine an assault on a country hosting U.S. troops in which U.S. troops do not directly come under attack (or suffer casualties) that nonetheless appears to challenge the deterrent commitment symbolized by the deployment. This scenario might still trigger concerns about the state's reputation for credibility, and producing support for escalation. We assess the significance of the nationality of the targets of an attack by testing Hypothesis 4:

H4: Attacks on U.S. troops deployed overseas should produce greater support for escalation than attacks on foreign troops deployed alongside them.

Location of Deployment

Fifth, the mechanisms produce distinct expectations about the significance of the *location* of troops under attack. Conventional understandings of “tripwire” forces – as in Fearon’s discussion – refer to troops deployed in areas directly threatened by a potential adversary. These deployments are significant because they embody a specific commitment to deterring challenges against a localized status quo. However, U.S. troops are often deployed in ways that do not fit the definition of a ‘tripwire’, as with overseas bases that function as staging areas for operations in the region or serve a more general presence function. Examining the effect of hypothetical attacks on these non-tripwire deployments allows us to distinguish between the two mechanisms developed above. The reputation mechanism implies that attacks on troops deployed as tripwires should increase support for escalation more than attacks on troops deployed elsewhere because, in theory, tripwire deployments constitute more direct signals of resolve and commitment than do other kinds of deployments. By contrast, from the perspective of the revenge mechanism, this difference is immaterial: casualties are casualties, wherever they occur. We test Hypothesis 5, related to the location of forward deployed troops:

H5: Attacks on troops deployed near a threatened or contested area should lead to greater levels of support for escalation than attacks on troops deployed outside of a threatened or contested area.

Presence

Finally, we examine the effect of the *presence* of U.S. troops in a partner country under attack. The reputation mechanism implies that information about the presence of U.S. troops within the borders of a country under attack should increase support for intervention, even in the absence of casualties. After all, if forward deployment embodies a commitment to deter, then attacks on countries hosting U.S. troops should prompt greater concerns about reputation and thus greater support for intervention than attacks on countries not hosting U.S. troops. The revenge mechanism, by contrast, does not imply that the presence of a deployment itself should influence support for escalation. This leads to hypothesis 6:

H6: The presence of U.S. troops should increase support for escalation after an attack.

Testing Hypothesis 6 requires a different experimental design than do hypotheses 1-5 as we explain below.

Table 1 summarizes the relation between these six hypotheses and the two mechanisms. As both mechanisms could operate simultaneously (or additively), it is not the case that evidence that supports (or undermines) one mechanism necessarily undermines (or supports) the other (Zaks 2017). Rather, some hypotheses provide clear tests of one mechanism, while providing little leverage for assessing the other. Taken together, these hypotheses allow us to distinguish whether the pattern of evidence presented below is broadly consistent with one, both, or neither of the mechanisms developed above.

Table 1: Patterns of Evidence Expected by Revenge and Reputation Mechanisms

Hypothesis	If evidence consistent with H is found...	If evidence inconsistent with H is found...
<i>H1: Attacks that result in higher numbers of casualties increase support for escalation (compared to attacks that result in lower numbers of casualties)</i>	Supports revenge mechanism Irrelevant for reputation mechanism	Undermines revenge mechanism Irrelevant for reputation mechanism

Hypothesis	If evidence consistent with H is found...	If evidence inconsistent with H is found...
<i>H2: Attacks on deployments intended to deter increase support for escalation (compared to attacks on deployments serving other purposes)</i>	Irrelevant for revenge mechanism Supports reputation mechanism	Irrelevant for revenge mechanism Undermines reputation mechanism
<i>H3: Attacks on military personnel increase support for escalation (compared to attacks on civilian personnel)</i>	Irrelevant for revenge mechanism Supports reputation mechanism	Irrelevant for revenge mechanism Undermines reputation mechanism
<i>H4: Attacks on American troops increase support for escalation (compared to attacks on foreign troops)</i>	Supports revenge mechanism Irrelevant for reputation mechanism	Undermines revenge mechanism Irrelevant for reputation mechanism
<i>H5: Attacks on deployments near threatened areas increase support for escalation (compared to attacks on remotely located deployments)</i>	Irrelevant for revenge mechanism Supports reputation mechanism	Irrelevant for revenge mechanism Undermines reputation mechanism
<i>H6: Attacks on countries hosting U.S. troops increase support for escalation (relative to attacks on countries not hosting U.S. troops)</i>	Irrelevant for revenge mechanism Supports reputation mechanism	Irrelevant for revenge mechanism Undermines reputation mechanism

Study 1: Three Realistic Vignettes

To test hypotheses 1-5, we fielded three parallel survey experiments using Amazon's Mechanical Turk platform during February 2020. This study assessed the effect of different features of attacks against U.S. security partners on the attitudes of Americans toward escalatory policies. To maximize experimental realism, we presented respondents with vignettes rooted in three real-world scenarios involving locations (South Korea, Iraq, and the Baltics) in which U.S. troops were then forward-deployed and escalations were plausible. We conducted the study in the wake of attacks (widely attributed to militias supported by the Islamic Republic of Iran) on U.S. troops based in Iraq in late December 2019. Our hope in doing so was to improve the "external validity" of our results – salience and realism may be important in this context given the revenge mechanism's theoretical reliance on emotional reactions to casualties. Thus, Aguinis and Bradley's (2014, 361) observation that "improving realism by increasing the similarity between the experimental and natural settings increases the observed effects" is especially relevant for our purposes here.

To measure the dependent variable (attitudes toward escalation), after the treatment we presented respondents with a list of potential policy responses that we designed to span a spectrum ranging from de-escalatory to highly escalatory. We first asked respondents to assess, on a 0-100 scale, how much each action would escalate tensions between the United States and the named adversary. We then asked respondents how much, on a 0-100 scale, they would support each of these options and a "do nothing" option in response to the attack described in the treatment, and which option they would most prefer. Table 2 lists both scores, in order of perceived escalatory potential. Finally, we also asked respondents to identify their most-preferred response option.

[Table 2 about here]

Table 2. Mean scores of respondent perceptions of escalation and support.

Retaliatory Option	Iran (N=2,103)		Korea (N=2,033)		Russia (N=2,963)	
	Escalation	Support	Escalation	Support	Escalation	Support
<i>Withdraw Troops</i>	27	50.3	28.5	38.4	24.4	45.5
<i>Condemn Attack</i>	35.7	74	36.6	74.2	37.1	75
<i>Levy Sanctions</i>	49.7	70.2	47.8	73.1	50	70.8
<i>Send Reinforcements</i>	52.9	47.9	52.1	56.6	50.9	54.3
<i>Use Cyber Attacks</i>	56.6	57.8	57.4	58.3	59.5	52.6
<i>Launch Airstrikes (Factory Targets)</i>	72.1	51.3	75.1	46.9	75.8	40.4
<i>Launch Airstrikes (Military Targets)</i>	74.7	48.4	77.3	47	77.8	41
<i>Invade</i>	82.4	34.6	83.5	34.8	83.4	28.9

Escalation scales run from 0 (marked “Would not escalate tensions at all”) to 100 (“Would escalate tensions as high as possible”), with 50 marked “Would escalate tensions somewhat”. Support scales run from 0 to 100, with 0 marked “Strongly oppose”, 25 marked “Oppose”, 50 marked “Neither support nor oppose”, 75 marked “Support”, and 100 marked “Strongly support”. N: Iran = 2,103, Korea = 2,033, Russia = 2,963.

Our aim in presenting respondents with a range of response options (rather than simply asking about support for intervention in general) was to assess the ability of pre-conflict attacks to influence support for interventions at different levels of escalation. Though not all analysts explicitly claim that attacks on forward deployed troops are capable of triggering demands for highly

escalatory responses, it is clear from Snyder and Schelling's writing that some prominent theorists do. The question of whether this is true has important implications for theory and policy. Asking for both evaluations of support for each type of retaliatory option and, separately, for respondents' most-preferred response option allowed us to measure these responses in greater detail.

Our treatments are rooted in vignettes that describe realistic scenarios involving three areas in which U.S. troops were deployed: the Baltics, the Korean Peninsula, and Iraq.⁶ Respondents completed a brief demographic survey and then read a short background summary mentioning a recent (accurate) news report mentioning the presence of U.S. troops in each scenario. This background section also included a treatment testing Hypothesis 2, which randomly varied the stated purpose of the troop deployment: whether or not "deterrence" of the named adversary was given as an explicit part of the mission (as opposed to the control condition, "protecting" American allies and interests). Respondents were administered a manipulation check between the background passage and the main treatment to assess whether they could identify the country hosting U.S. troops.

Next, respondents were asked to imagine that the relevant adversary (North Korea, Iran, or Russia) had launched an attack. To test Hypothesis 1, we randomized the number of casualties resulting from the attack (0, 1, 10, or 200 killed). We also varied information about the nature and location of the targets of the attack in different vignettes to test hypotheses 3, 4, and 5. In the Iran vignette, we described attacks as having been launched against American troops based in Iraq, American embassy officials in Baghdad, or Iraqi troops. Comparing the effect of attacks against American troops vs. American embassy officials facilitates a test of Hypothesis 3. Comparing the effect of attacks against American troops vs. Iraqi troops contributes to a test of Hypothesis 4. In the Russia vignette, we described attacks as having been launched against American troops, Lithuanian troops, or foreign NATO troops. Comparing the effect of attacks against American troops vs. Lithuanian and NATO troops contributes to a test of Hypothesis 4. In the North Korea vignette, we described attacks as having been launched against American troops in the DMZ, American troops stationed at a base in Okinawa, or South Korean troops. Comparing the effect of attacks against American troops vs. South Korean troops contributes to a test of Hypothesis 4.

⁶ Respondents received \$0.30 for completing the survey. The Iran scenario received 2,442 responses, the Korea scenario received 2,323, and the Russian scenario received 3,339, for a total of 8,104 respondents. Details are available in the appendix.

Comparing the effect of attacks against American troops in the DMZ vs. American troops on Okinawa enables a test of Hypothesis 5.

To explore the relationship between treatment conditions and support for different responses, we present results from OLS regression models that include all treatments and control for scenario (where appropriate), demographics, party identification, and attentiveness (as measured by the attention and manipulation check). Additional information is available in the appendix.

[Figure 1 about here]




Figure 1.

Figure 1 displays results for different numbers of casualties. For each response option and casualty level (with zero casualties as the base category), we show three different coefficient estimates. The first compares casualties across attacks against all types of targets; the second excludes local targets (Lithuanian, South Korean, and Iraqi troops). The third, and most restrictive, includes only targets that fit conventional definitions of a “tripwire”: U.S. troops stationed near a threatened area.⁷

⁷ We focus on results for these target types for purposes of presentation. While other ways of aggregating target types are possible, they do not substantially alter the results discussed here.

Results are generally consistent with Hypothesis 1. Regardless of which target category we focus on, increasing casualties lead to higher support for interventions that are on the more escalatory side of the response continuum. This effect also appears to strengthen as we restrict target type – it is weakest when we include local troops in the analysis, and strongest when we focus only on U.S. troops.⁸ It is important to note, though, that the effect of casualties on support for escalation is limited. The maximum measured effect is quite modest: an attack that causes 200 U.S. military casualties produces an increase of only about 13 points on a 100-point scale in terms of support for airstrikes against the adversary’s military bases, or just over half the distance between each of the five marked categories. Other results are much smaller, such as the approximate 5-point increase in support for airstrikes against bases caused by 10 casualties.

These results hold when we examine respondents’ most-preferred responses. Figure 2 displays the results of multinomial logit regressions to predict the most favored response across different casualty levels. As Figure 2 shows, across all scenarios, suffering 200 U.S. military casualties only brings the airstrike option into a rough three-way tie with withdrawal and condemnation for the second-most preferred response (behind sanctions). This is surprising, given that 200 casualties would be an exceptionally bloody day for the American military in the context of recent history. Moreover, even large numbers of casualties do not appear capable of substantially increasing support for the most escalatory response: invasion. Here, the increase in the level of support is only around 7 points, and invasion remains a low-ranked option. Our evidence is thus broadly inconsistent with Snyder and Schelling’s claims that attacks on “tripwire” deployments would result in widespread demands for highly escalatory policy responses.

[Figure 2 about here]

⁸ This general pattern is broadly consistent across all three scenarios analyzed separately.




Figure 2. Values represent the predicted share of respondents choosing each response, by intervention and casualty level. Represents pooled results for all three scenarios; full results available in the appendix.

Figure 3 displays estimates of the effect of varying the stated purpose of the deployment. We present coefficient estimates for this treatment for each response option, aggregated across scenarios, for the three target types described above. Results are inconsistent with Hypothesis 2. Support for escalation is not responsive to information about whether or not deployments are intended to deter. This undermines the reputation mechanism.

[Figure 3 about here]




Figure 3.

Figure 4 presents results of the influence of the civilian or military status of the target of an attack. Our analysis here is limited to results from the Iran scenario, in which the attack targeted Iraqi troops, U.S. troops, or Baghdad embassy officials. To test Hypothesis 3, we compare the two treatment conditions involving Americans. Results are inconsistent with Hypothesis 3: attacks on U.S. military personnel do not increase support for escalation compared to attacks on American civilian diplomats. Like the results presented in Figure 3, this undermines the reputation mechanism.

[Figure 4 about here]




Figure 4.

Figure 5 displays estimates of the effect of the nationality of the target of an attack. This analysis compares treatment conditions from all three scenarios that specify attacks against American targets vs. foreign (local or NATO) targets. We also show results from the Russia scenario separately, comparing the effect of attacks on American troops vs. attacks on NATO troops – an especially useful comparison because both of these targets play the same “tripwire” role in the context of the vignette. Results show that support for the most escalatory responses is several points higher when the victims are fellow Americans compared to when they are foreign. This evidence is consistent with Hypothesis 4, and supports the revenge mechanism.

[Figure 5 about here]

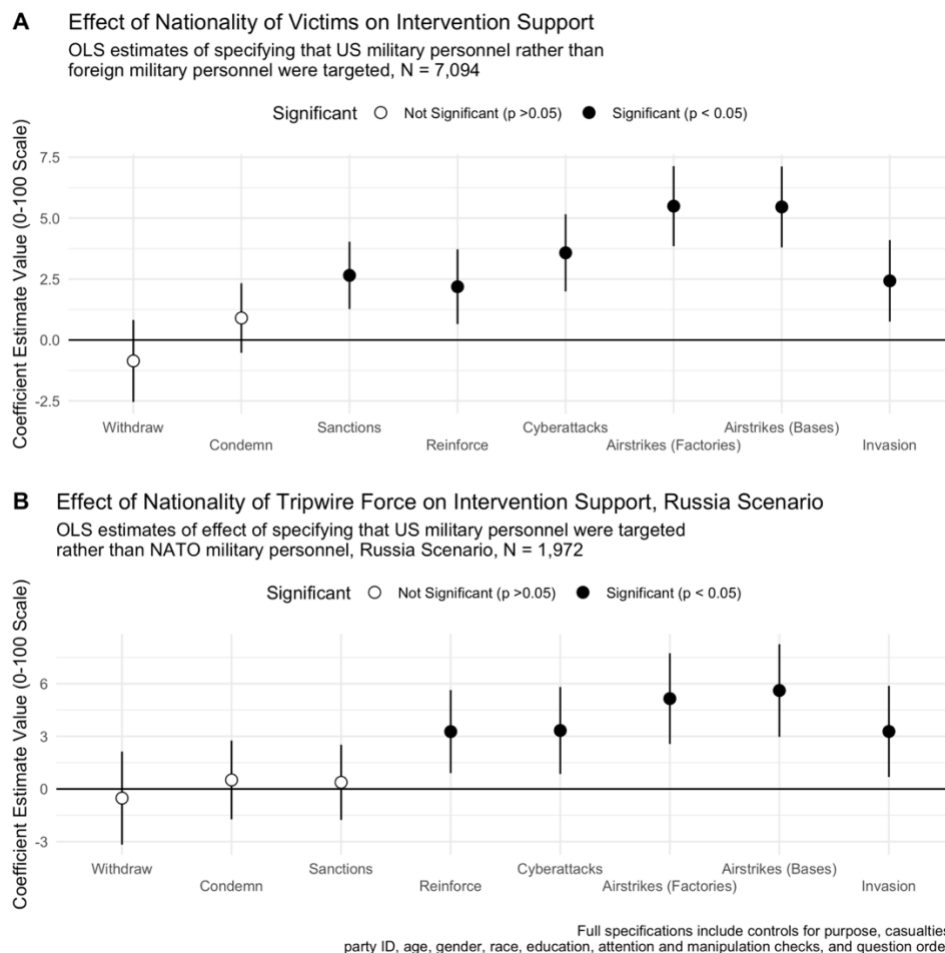


Figure 5. Panel A displays estimates of the effect of the nationality of victims on support for different response options across all three scenarios. Panel B repeats this analysis, but compares only US military vs. NATO military victims from the Russia scenario.

Figure 6 displays coefficient estimates for the location of the attack. This analysis relies on the North Korea scenario, comparing treatment conditions describing an attack on U.S. troops stationed in the DMZ, and an attack on U.S. troops stationed on Okinawa. Results are inconsistent with Hypothesis 5. Attacks on the “tripwire” deployment do not increase support for escalation compared to attacks on the non-tripwire deployment. This evidence, like our tests of Hypotheses 2 and 3, undermines the reputation mechanism.

[Figure 6 about here]




Figure 6.

Study 2: The Effect of U.S. Military Presence

The experimental design that we presented in the previous section cannot be used to examine the effect of the *presence* of U.S. troops on public support for intervention after an attack on a partner country. We were unable to vary information about the presence or number of troops deployed in our three scenarios, since these values were constrained by real-world conditions. It might be possible to ask respondents to imagine a *future* situation in which U.S. troops were not present in

these scenarios, but this would risk introducing extraneous information. Respondents might reason that other – unmentioned – elements of the scenario must also have changed (perhaps as a cause or consequence of the removal of American troops). Any measured differences in attitudes toward conflict would thus be impossible to attribute to variation in deployment, rather than to other inferences respondents might draw from information about deployment.

To overcome this problem, we designed a conjoint survey experiment. Conjoint designs allow researchers to simultaneously vary the separate influence of many different factors on respondents' choices (Hainmueller, Hopkins, and Yamamoto 2014). This is especially attractive for our purposes because we can directly control for and measure the effect of factors about which respondents might otherwise draw inferences based on information about deployment. And, though our approach here necessarily relies on more abstract scenarios than our vignette study, recent research confirms that abstract designs still yield useful insights about the real world (Brutger et al. 2020). Conjoint studies in particular also perform well even if they vary information about many different dimensions (Bansak et al. 2019).

Our conjoint survey presented respondents with pairs of scenario profiles that varied along a number of dimensions (which we describe below). Respondents were told that these were scenarios in which one country had attacked another, and that the United States could choose to intervene militarily (or not) on behalf of the victim. They were asked to choose the scenario in which they would be more willing to support intervention, and to rate how likely they would be to support intervention in each scenario.⁹ Each respondent completed this task five times, in each case comparing and choosing between two randomly generated “attack profiles.”¹⁰ This design allows us to estimate the effect of each variable included in the profile description on support for intervention, independent of all of the other included variables.

[Table 3 about here]

⁹ The additional complexity inherent in the conjoint setting precluded us from disaggregating intervention as we did in the vignette study.

¹⁰ Bansak et al. (2018) show that conjoint surveys perform well even if respondents are asked to complete many different tasks.

Table 3. Complete list of conjoint attributes and levels.

Attribute	Levels
US Military Presence in Victim Country	<ul style="list-style-type: none"> • No US military personnel present • A few hundred military personnel • A few thousand military personnel • Tens of thousands of military personnel
Casualties Suffered Already (<i>Experiment 2 only</i>)	<ul style="list-style-type: none"> • No • Several • Dozens <p>(If no personnel are present in previous condition, no casualties suffered)</p>
Probability of US Intervention Succeeding	<ul style="list-style-type: none"> • Almost certainly will not succeed • Probably will not succeed • About even chance of succeeding • Probably will succeed • Almost certainly will succeed
Likely US Military Casualties in an Intervention	<ul style="list-style-type: none"> • Severe (More than 10,000 US casualties) • Moderate (Several thousand US casualties) • Light (No more than several hundred US casualties)
Region	<ul style="list-style-type: none"> • The Middle East • Eastern Europe • Asia
Aggressor Regime Type	<ul style="list-style-type: none"> • Autocratic • Democratic
Aggressor Trade Relationship with US	Major/minor/not a significant US trade partner
Aggressor Nuclear Status	<ul style="list-style-type: none"> • Does not have nuclear weapons • Has a limited nuclear arsenal • Has a substantial nuclear arsenal
Victim Regime Type	<ul style="list-style-type: none"> • Autocratic • Democratic
Victim Nuclear Status	<ul style="list-style-type: none"> • Does not have nuclear weapons • Has a limited nuclear arsenal • Has a substantial nuclear arsenal
Victim Country Trade Relationship with US	Major/minor/not a significant US trade partner
US Civilians at Risk in Victim Country	<ul style="list-style-type: none"> • Fewer than 100

Attribute	Levels
	<ul style="list-style-type: none"> • Several thousand • More than 100,000
Risk of Aggressor Retaliation Against US Homeland	<ul style="list-style-type: none"> • Almost no chance • Probably no chance • An even chance • A very good chance
International Community Stance	Likely/Unlikely the UN Security Council will support US intervention on behalf of the victim
Congressional Authorization	Likely/Unlikely that Congress will formally approve US intervention

Table 3 summarizes the variables included in profile descriptions, along with the range of values each could take on. As our primary objective is to estimate the effect of the presence of U.S. troops in an attacked country on support for intervention, we varied information about whether or not – and if so, how many – U.S. troops were deployed inside the victim country. This ranged from zero troops deployed, to tens of thousands, with two possible intermediate values: a few hundred troops, and a few thousand troops. These levels capture the range of real-world variation in the size of American overseas military contingents.

We included a number of other variables in our profile descriptions, six of which are especially significant. These included information about whether or not the United States had an alliance with the victim country; how many U.S. casualties any future intervention was expected to result in; the probability that the intervention would succeed; the risk that the aggressor would be able to retaliate against the U.S. homeland; whether or not the UN Security Council was likely to approve the intervention; and whether or not the U.S. Congress was likely to approve the intervention. These variables are important because they convey information about the United States' relationship with the victim country, as well as about the costs, risks, legitimacy, and difficulty of the intervention that might otherwise be implicitly communicated by information about the presence or absence of U.S. troops. Their inclusion also allows us to assess the *magnitude* of the effect of forward deployment relative to other salient factors.

Results are based on the responses of 912 respondents recruited using MTurk in 2018 (full details are available in the appendix).¹¹ Respondents completed a short survey about their personal background and views on international relations generally, and then five conjoint tasks.

[Figure 7 about here]




Figure 2. Conjoint Experiment 1 core results. $N = 4,570$ (914 respondents for five tasks each).

¹¹ Respondents received \$0.75 for completing the survey. Recent discussions have questioned the reliability of Mechanical Turk respondents because of concerns about recruitment during a period that overlapped with this study. Including or excluding users with suspect IP addresses does not change our results.

Figure 7 presents our results. There is some evidence that troop deployments increase support for intervention, though two caveats are in order. First, only relatively large deployments make any difference: deployments of at least a few thousand troops increase support for intervention, but smaller deployments make no difference relative to the no-troop base category. Second, the effect is small: even the presence of tens of thousands of troops increases support for intervention by only about 3 percentage points, smaller than the effect of other factors. Treaties increase support for intervention by almost twice as much, as does anticipated Congressional and UN Security Council authorization. Other variables are even more important: a very high risk of retaliation reduces support by almost ten percentage points relative to very low risk. Meanwhile – in line with prior work on support for war – the expected costs and outcome of the intervention are by far the most important factors. The expectation of “severe” casualties reduces support by almost 20 percentage points relative to the “light casualty” base category; and reading that the intervention will almost certainly fail reduces support by over 20 percentage points relative to almost certain success. Taken together, these results lend at best weak support for Hypothesis 6. While there is evidence that relatively large contingents of forward deployed troops increase support for intervention, the effect is small and ranks as the least important factor that we included in the profile descriptions.

Discussion

We summarize our findings in Table 3. Our evidence paints a consistent picture. Forward deployment increases public support for intervention through the revenge mechanism rather than a reputational one. Only in the conjoint experiment, testing H6, do we find evidence that is consistent with the reputation mechanism, and then in a test irrelevant for the revenge mechanism; elsewhere, the revenge mechanism passes each of its tests and the reputation mechanism fails each of its tests. Taken together, the balance of the evidence provides greater support for Schelling and Snyder’s conjectured public-opinion effect driving a tripwire response: demands for escalation following an attack on forward deployed troops are likelier to be driven by an outraged response to the spectacle of unanticipated casualties among co-nationals than by a rational calculation that the state’s reputation has been challenged and must be restored.

Table 3: Pattern of Evidence

Hypothesis	Findings
<i>H1: Attacks that result in higher numbers of casualties increase support for escalation (compared to attacks that result in lower numbers of casualties)</i>	Supported weakly (Figures 1 and 2); supports revenge mechanism, irrelevant for reputation mechanism
<i>H2: Attacks on deployments intended to deter increase support for escalation (compared to attacks on deployments serving other purposes)</i>	Rejected (Figure 3); Irrelevant for revenge mechanism, undermines reputation mechanism
<i>H3: Attacks on military personnel increase support for escalation (compared to attacks on civilian personnel)</i>	Rejected (Figure 4) Irrelevant for revenge mechanism; undermines reputation mechanism
<i>H4: Attacks on American troops increase support for escalation (compared to attacks on foreign troops)</i>	Supported (Figure 5): Supports revenge mechanism; irrelevant for reputation mechanism
<i>H5: Attacks on deployments near threatened areas increase support for escalation (compared to attacks on remotely located deployments)</i>	Rejected (Figure 6): Irrelevant for revenge mechanism; undermines reputation mechanism
<i>H6: Attacks on countries hosting U.S. troops increase support for escalation (relative to attacks on countries not hosting U.S. troops)</i>	Supported weakly (Figure 7) Irrelevant for revenge mechanism; supports reputation mechanism

Our evidence allows us not only to establish the mechanism most frequently linking public opinion and attacks on forward deployed forces but also to measure the effect size. Part of the reason that tripwire deployments seem attractive is because they should be, in theory, capable of

achieving deterrence without triggering security dilemmas. If even small numbers of casualties among American troops are capable of producing strong demands for escalation, then only small deployments that cannot significantly shift the balance of capabilities (and thus threaten a potential adversary) might be necessary in order to generate an effective pre-commitment mechanism. Our analysis suggests that this story may be too good to be true. We find that the link between attacks on forward deployed troops (or countries hosting them) and public support for escalation is weaker than analysts have often assumed. In our conjoint study, even the very small increase in support for intervention linked to the presence of U.S. troops in an attacked state only appears when deployments are relatively large. In our vignette studies, large numbers of casualties were necessary before we observe substantial increases in support for escalatory intervention. Although the positive effect of increasing pre-conflict casualties on support for escalation is interesting and important – especially considering prior work on casualties and support for war – it is unclear whether it is large enough, on its own, to force a reluctant leader into escalation, or to guarantee public support for such a policy. Moreover, respondents only became substantially more supportive of *moderately* escalatory options like airstrikes rather than highly escalatory options, like an invasion. Thus, contra Snyder’s claim about public support for nuclear retaliation, it seems unlikely that a tripped “tripwire” could produce strong demands for disproportionate – and costly – escalation.

Conclusion

Theorists, analysts, and policymakers have long thought that forward deployed troops functioned as strong pre-commitment devices, sending effective deterrent signals by guaranteeing public support for escalation in case of an attack on the host. In this paper, we have conducted the first empirical test of this claim’s key mechanism. Our findings suggest two broad conclusions, and also point toward important questions for future research.

First, there is evidence that attacks on troops deployed abroad *might* produce support for escalation, probably by sparking demands for revenge. This important finding constitutes the first empirical support for the widespread notion that attacks on forward deployed troops can increase American public support for war, and adds nuance to the well-established finding that casualties suffered *during* a conflict reduce support for intervention. Pre-conflict casualties are evidently different – they induce support for retaliation, rather than caution or pessimism about further involvement.

Second, our analysis suggests that forward deployment is not as strong a pre-commitment device as theorists and policymakers sometimes claim. This is due partly to the modest size of even the largest measured increases in support for escalation; to the fact that very large numbers of casualties are required before we observe substantial increases in support for escalation; and to the fact that increases in support seem to be limited to only moderately escalatory responses. The nature of the revenge mechanism calls for further caution about claims that tripwire effects will be automatically produced. Hall (2017, 8) and Stein (2015, 559) contend that public outrage and demands for revenge are best understood as political resources that policymakers can foster and harness to mobilize support for belligerent policies that they already prefer, rather than as strong constraints that force policymakers to act in ways that they might not otherwise. If this is true, it calls into question even more seriously the notion that forward deployment can reliably tie leaders' hands.

Of course, forward deployment could serve as a strong pre-commitment device through some mechanism that theorists have not yet specified and that analysts do not customarily invoke. For instance, audiences of foreign policy elites or military advisors might respond more strongly than the broader public does to attacks on troops deployed abroad. Thus, leaders may still be strongly constrained by narrower (but more influential) groups whose attitudes toward conflict and intervention are systematically different than are those of our survey respondents. Alternatively, leaders might act out of the incorrect *anticipation* of strong domestic constraints (rooted, perhaps, in the outsized influence of writers like Schelling). These possibilities lie outside the scope of this study, but are worth exploring further.

Future work should also investigate the beliefs, attitudes, and responses of *foreign* audiences and actors in relation to forward deployment. How might different audiences within NATO member states react to attacks on their own forward deployed troops? How do policymakers in host states *expect* attacks on troops deployed within their borders to influence escalation dynamics? And how have policymakers in potential adversaries perceived the significance and function of deployments hosted in nearby states? Although our analysis here takes an important step toward better understanding the function of forward deployment, there is much more to learn about how this important and much-used, but understudied foreign policy tool actually works.

Works Cited

- Aguinis, Herman, and Kyle J. Bradley. 2014. "Best Practice Recommendations for Designing and Implementing Experimental Vignette Methodology Studies." *Organizational research methods* 17(4): 351–71.
- Bandow, Doug. 2018. "Why Trump Is Right to Withdraw Troops." *Cato Institute*.
<https://www.cato.org/commentary/why-trump-right-withdraw-troops> (February 3, 2021).
- Bansak, Kirk, Jens Hainmueller, Daniel J. Hopkins, and Teppei Yamamoto. 2019. "Beyond the Breaking Point? Survey Satisficing in Conjoint Experiments." *Political Science Research and Methods*: 1–19.
- Baum, Matthew A., and Philip BK Potter. 2019. "Media, Public Opinion, and Foreign Policy in the Age of Social Media." *The Journal of Politics* 81(2): 747–56.
- Blankenship, Brian, and Erik Lin-Greenberg. 2021. "Rethinking Reassurance: The Importance of Military Capabilities in Credibly Assuring Allies." *Available at SSRN*.
- Brooks, Stephen G., and William Curti Wohlforth. 2016. *America Abroad: The United States' Global Role in the 21st Century*. Oxford University Press.
- Brutger, Ryan et al. 2020. *Abstraction and Detail in Experimental Design*. Working Paper 1-47. Accessed June 4, 2020. <https://api.semanticscholar.org>
- Burns, Sarah, and Andrew Stravers. 2020. "Obama, Congress, and Audience Costs: Shifting the Blame on the Red Line." *Political Science Quarterly* 135(1): 67–101.
- Casler, Don, and Richard Clark. Forthcoming. "Trade Rage: Audience Costs and International Trade." *Journal of Conflict Resolution*.
- Chaudoin, Stephen. 2014. "Promises or Policies? An Experimental Analysis of International Agreements and Audience Reactions." *International Organization*: 235–56.
- Cooley, Alexander, and Daniel H. Nexon. 2013. "The Empire Will Compensate You": The Structural Dynamics of the U.S. Overseas Basing Network." *Perspectives on Politics* 11: 1034–50.
- Davies, Graeme AM, and Robert Johns. 2013. "Audience Costs among the British Public: The Impact of Escalation, Crisis Type, and Prime Ministerial Rhetoric." *International Studies Quarterly* 57(4): 725–37.
- Fazal, Tanisha M. 2021. "Life and Limb: New Estimates of Casualty Aversion in the United States." *International Studies Quarterly* 65(1): 160–72.
- Fearon, James D. 1994. "Domestic Political Audiences and the Escalation of International Disputes." *American Political Science Review* 88(3): 577–92.
- . 1997. "Signaling Foreign Policy Interests: Tying Hands versus Sinking Costs." *Journal of Conflict Resolution* 41(1): 68–90.

- Freedman, Lawrence, and Jeffrey Michaels. 2019. *The Evolution of Nuclear Strategy: New, Updated and Completely Revised*. Springer.
- Fuhrmann, Matthew, and Todd S. Sechser. 2014. "Signaling Alliance Commitments: Hand-Tying and Sunk Costs in Extended Nuclear Deterrence." *American Journal of Political Science* 58(4): 919–35.
- Gartner, Scott Sigmund. 2008. "The Multiple Effects of Casualties on Public Support for War: An Experimental Approach." *American political science review*. 95–106.
- Gartner, Scott Sigmund, and Gary M. Segura. 1998. "War, Casualties, and Public Opinion." *Journal of conflict resolution* 42(3): 278–300.
- Gelpi, Christopher, and Joseph M. Grieco. 2015. "Competency Costs in Foreign Affairs: Presidential Performance in International Conflicts and Domestic Legislative Success, 1953–2001." *American Journal of Political Science* 59(2): 440–56.
- Geys, Benny. 2010. "Wars, Presidents, and Popularity: The Political Cost (s) of War Re-Examined." *Public Opinion Quarterly* 74(2): 357–74.
- Gordon, Philip H. 2020. "Trump's Sudden and Dangerous Troop Withdrawal From Germany." *Council on Foreign Relations*. <https://www.cfr.org/in-brief/trumps-sudden-and-dangerous-troop-withdrawal-germany> (February 3, 2021).
- Hainmueller, Jens, Daniel J. Hopkins, and Teppei Yamamoto. 2014. "Causal Inference in Conjoint Analysis: Understanding Multidimensional Choices via Stated Preference Experiments." *Political Analysis* 22: 1–30.
- Hall, Todd H. 2017. "On Provocation: Outrage, International Relations, and the Franco–Prussian War." *Security Studies* 26(1): 1–29.
- Jakobsen, Jo, and Tor G. Jakobsen. 2019. "Tripwires and Free-Riders: Do Forward-Deployed US Troops Reduce the Willingness of Host-Country Citizens to Fight for Their Country?" *Contemporary security policy* 40(2): 135–64.
- Kertzer, Joshua D., and Ryan Brutger. 2016. "Decomposing Audience Costs: Bringing the Audience Back into Audience Cost Theory." *American Journal of Political Science* 60(1): 234–49.
- Kriner, Douglas L., and Francis X. Shen. 2012. "How Citizens Respond to Combat Casualties: The Differential Impact of Local Casualties on Support for the War in Afghanistan." *Public Opinion Quarterly* 76(4): 761–70.
- Levendusky, Matthew S., and Michael C. Horowitz. 2012. "When Backing down Is the Right Decision: Partisanship, New Information, and Audience Costs." *The Journal of Politics* 74(2): 323–38.
- Levy, Jack S., Michael K. McKoy, Paul Poast, and Geoffrey PR Wallace. 2015. "Backing out or Backing in? Commitment and Consistency in Audience Costs Theory." *American Journal of Political Science* 59(4): 988–1001.

- Lin-Greenberg, Erik. 2021. "Soldiers, Pollsters, and International Crises: Public Opinion and the Military's Advice on the Use of Force." *Foreign Policy Analysis* 17(3).
- MacKinnon, Mark. 2017. "Canadians in Divided Latvia Will Defend against Any Russian Aggression." *The Globe and Mail*. <https://www.theglobeandmail.com/news/world/canada-in-latvia-russia-border/article34695831/> (February 4, 2021).
- Mueller, John E. 1973. *War, Presidents, and Public Opinion*. New York: Wiley.
- Nomikos, William G., and Nicholas Sambanis. 2019. "What Is the Mechanism Underlying Audience Costs? Incompetence, Belligerence, and Inconsistency." *Journal of Peace Research* 56(4): 575–88.
- O'Mahony, Angela et al. 2018. *US Presence and the Incidence of Conflict*. RAND Corporation Santa Monica United States.
- Quek, Kai. 2021. "Four Costly Signaling Mechanisms." *American Political Science Review*: 1–13.
- Reiter, Dan, and Paul Poast. 2021. "The Truth About Tripwires: Why Small Force Deployments Do Not Deter Aggression." 4(3). <https://tnsr.org/2021/06/the-truth-about-tripwires-why-small-force-deployments-do-not-deter-aggression/>.
- Rovner, Joshua, and Caitlin Talmadge. 2014. "Hegemony, Force Posture, and the Provision of Public Goods: The Once and Future Role of Outside Powers in Securing Persian Gulf Oil." *Security Studies* 23(3): 548–81.
- Sander, Gordon F. 2017. "When the U.S. Almost Went to War With North Korea." *POLITICO Magazine*. <https://www.politico.com/magazine/story/2017/09/14/north-korea-1976-axe-murder-incident-215605> (February 10, 2021).
- Schelling, Thomas C. 1960. *The Strategy of Conflict*. Cambridge, Mass: Harvard University Press.
- . 1966. *Arms and Influence*. New Haven, CT: Yale University Press.
- Schmidt, Sebastian. 2020. *Armed Guests: Territorial Sovereignty and Foreign Military Basing*. Oxford University Press, USA.
- Shandler, Ryan, Michael L. Gross, Sophia Backhaus, and Daphna Canetti. 2021. "Cyber Terrorism and Public Support for Retaliation—a Multi-Country Survey Experiment." *British Journal of Political Science*: 1–19.
- Slack, Matthew C. 2018. "Conventional Tripwire Deterrence." Dissertation. University of Chicago.
- Snyder, Glenn Herald. 2015. *Deterrence and Defense*. Originally published 1961. Princeton University Press.
- Snyder, Jack, and Erica D. Borghard. 2011. "The Cost of Empty Threats: A Penny, Not a Pound." *American Political Science Review* 105(3): 437–56.

- Stein, Rachel M. 2015. "War and Revenge: Explaining Conflict Initiation by Democracies." *American Political Science Review* 556–73.
- Tomz, Michael. 2007. "Domestic Audience Costs in International Relations: An Experimental Approach." *International Organization* 61(4): 821–40.
- Tomz, Michael, and Jessica Weeks. Forthcoming. "Military Alliances and Public Support for War." *International Studies Quarterly*.
- Tomz, Michael, Jessica LP Weeks, and Keren Yarhi-Milo. 2020. "Public Opinion and Decisions about Military Force in Democracies." *International Organization* 74(1): 119–43.
- Trachtenberg, Marc. 1989. "Strategic Thought in America, 1952-1966." *Political Science Quarterly* 104(2): 301–34.
- Trager, Robert F., and Lynn Vavreck. 2011. "The Political Costs of Crisis Bargaining: Presidential Rhetoric and the Role of Party." *American Journal of Political Science* 55(3): 526–45.
- Vanaga, Nora, and Toms Rostoks. 2018. *Deterring Russia in Europe: Defence Strategies for Neighbouring States*. Routledge.
- Zaks, Sherry. 2017. "Relationships Among Rivals (RAR): A Framework for Analyzing Contending Hypotheses in Process Tracing." *Political Analysis* 25(3): 344–62.