

The Tripwire Effect: Experimental Evidence Regarding U.S. Public Opinion

PAUL MUSGRAVE 

Department of Political Science and Legal Studies, University of Massachusetts Amherst, Amherst, MA 01003, USA

AND

STEVEN WARD 

Department of Politics and International Studies, University of Cambridge, Alison Richard Building, Cambridge CB3 9DT, UK

Claims about the deterrent role of forward military deployments often depend on the argument that attacks on troops stationed abroad will activate a “tripwire effect.” It is assumed that this effect would generate strong domestic pressure for further military intervention by the country whose troops have come under attack. In theory, the anticipation of a strong tripwire effect prospectively ties the hands of leaders, thereby bolstering the reliability of extended deterrent threats and promises. In this paper, we define the tripwire effect and use both conjoint and vignette survey experiments to evaluate its operation and magnitude among Americans. Results suggest that the tripwire effect is, at best, far weaker than many analysts and policymakers commonly assume. This finding raises serious questions about a core logic underpinning the United States’ forward military posture and highlights the need for more research on the deterrent functions of forward deployment.

Las afirmaciones con respecto al papel disuasorio que tienen los despliegues militares avanzados dependen, con frecuencia, del argumento de que los ataques contra las tropas estacionadas en el extranjero activarán un «efecto trampa». Se supone que este efecto generaría una fuerte presión interna en relación con una mayor intervención militar por parte del país cuyas tropas han sido atacadas. En teoría, el hecho de anticiparse a este fuerte efecto trampa coarta, prospectivamente, la libertad de los gobernantes, lo cual refuerza, de esta forma, la fiabilidad que tienen tanto las amenazas como las promesas disuasorias extendidas. En este artículo, definimos este efecto trampa y utilizamos experimentos de encuesta conjunta y estudios de viñeta con el fin de evaluar su funcionamiento y su magnitud entre los estadounidenses. Los resultados sugieren que el efecto trampa es, en el mejor de los casos, mucho más débil de lo que muchos analistas y muchos formuladores de políticas asumen habitualmente. Esta conclusión plantea serias dudas sobre la lógica fundamental que sustenta la postura militar avanzada de los Estados Unidos y destaca la necesidad de llevar a cabo una mayor investigación sobre las funciones disuasorias del despliegue avanzado.

Le rôle de dissuasion des déploiements militaires avancés se justifie souvent par les effets en cascade qu’une attaque sur des troupes stationnées à l’étranger déclencherait. Ces effets généreraient une importante pression

Paul Musgrave is an Associate Professor of political science at the University of Massachusetts Amherst. He studies U.S. foreign policy, international relations theory, and their intersection.

Steven Ward is an Assistant Professor of international relations in the Department of Politics and International Studies at the University of Cambridge.

Musgrave, Paul, and Steven Ward. (2023) The Tripwire Effect: Experimental Evidence Regarding U.S. Public Opinion. *Foreign Policy Analysis*, <https://doi.org/10.1093/fpa/orad017>

© The Author(s) (2023). Published by Oxford University Press on behalf of the International Studies Association. All rights reserved. For permissions, please e-mail: journals.permissions@oup.com

chez le pays touché en vue d'intensifier les interventions militaires. En théorie, l'anticipation de tels effets immobiliserait les dirigeants ; les menaces et promesses à des fins de dissuasion seraient donc fiables et pourraient être largement utilisées. Dans cet article, nous définissons ces effets en cascade avant d'utiliser des sondages conjoints et capsules pour évaluer leur fonctionnement et leur ampleur chez les Américains. D'après les résultats, ces effets en cascade seraient dans le meilleur des cas bien moins puissants que nombre d'analystes et de décideurs politiques le supposent généralement. Cette conclusion soulève de sérieuses questions quant à la logique fondamentale du positionnement militaire avancé des États-Unis. Elle souligne aussi la nécessité de davantage de travaux de recherche sur les fonctions dissuasives des déploiements avancés.

Successful extended deterrence requires that threats be credible. Yet credibility may be difficult to achieve if the interests to be protected from attack are geographically distant, strategically unimportant in themselves, and difficult to defend. These conditions describe several current U.S. security partners—such as the Baltic states—that rely on U.S. security guarantees.

Policymakers address this dilemma by contriving ways to enhance the credibility of security guarantees. One common tool involves the deployment of troops within the security partner's territory. This practice has been at the center of U.S. foreign policy since the end of the Second World War (Schmidt 2020), but has gained new relevance recently as questions have arisen about the reliability of U.S. security guarantees in Eastern Europe. After Russia's annexation of Crimea in 2014, the United States and other NATO allies implemented the Enhanced Forward Presence (EFP) program, which involved the stationing of multinational battalion-sized battle groups in Latvia, Lithuania, Estonia, and Poland. The EFP highlights the idea that forward deployment is sometimes not intended to deter by materially changing the balance of military forces between aggressor and defending power. As military analysts have noted, EFP deployments (before 2022, at least) would not have been adequate to defend the Baltics from a Russian attack (Shlapak 2018). Instead, they were meant to deter by guaranteeing further escalation in case the Baltic states were attacked (Zapfe 2017).

This function of forward deployment, often labeled “tripwire” deterrence, has long played a central role in U.S. grand strategy. The logic has also arguably attained canonical status in the discipline of international relations, not least due to Thomas Schelling's colorful and influential analysis of the Berlin Brigade's role in preventing Soviet aggression (Schelling 1966, 47). Yet surprisingly little has been written about how tripwire deterrence works. Although scholars have thoroughly investigated other kinds of credibility-enhancing measures (such as formal alliances and public threats), there has been little sustained theoretical or empirical attention to the mechanics of forward deployment as a deterrent.

One important oversight relates to the intuition—central to the logic of tripwire deterrence—that attacks on forward-deployed forces would produce a significant increase in domestic political support for further intervention. Although claims about tripwire deterrence often assume that casualties suffered before the outbreak of open conflict will lead to demands for escalation, this proposition has never been systematically empirically examined. This is especially troubling because related research casts some doubt on this expectation. Indeed, the relationship could plausibly run in the opposite direction, as increasing casualties *during* wartime have been found to reduce U.S. public support for conflict.

We aim to address this core theoretical and empirical oversight. First, we argue that tripwire deterrence is often at least implicitly (and sometimes explicitly) theo-

rized to work via a domestic political mechanism. Analysts often assume that attacks on forward-deployed troops will generate a strong demand among domestic audiences for an escalatory response, which we call the “tripwire effect.” In theory, the anticipation of a strong tripwire effect ties the hands of leaders in the state whose troops are deployed, thereby deterring potential adversaries by producing the expectation of nearly automatic escalation. Second, we report the results of a series of survey experiments designed to measure the tripwire effect. Although, in principle, tripwire dynamics may apply beyond the U.S. case, we concentrate on assessing the reactions of Americans to attacks on forward-deployed U.S. troops because the United States makes more use of such forward deployments than any other country. Our findings suggest that the tripwire effect is, at best, modest and potentially too small to function as a reliable pre-commitment device. Indeed, other pre-commitment devices derided by early generations of tripwire theorists, such as treaties of mutual defense, may be more effective.

The paper proceeds as follows. We first briefly review various logics through which forward deployment may bolster extended deterrent commitments, with special attention to developing tripwire deterrence as a distinct mechanism. We then explicate the theoretical foundations of tripwire deterrence. We identify the tripwire effect—the production of demands for escalation among domestic audiences—as a key part of the logic of tripwire deterrence. The next section introduces and describes a series of survey experiments designed to measure the tripwire effect. We report results of both conjoint and vignette experiments aimed at identifying the causal effect on American attitudes toward military intervention of attacks on U.S. troops deployed abroad. In the conclusion, we highlight the implications of our findings for theoretical and policy debates and suggest some directions for future research on tripwire deterrence.

Forward Deployment and Deterrence

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 (Vanaga and Rostoks 2018; Schmidt 2020). Although the practice may serve a range of purposes, forward deployments are often meant—at least in part—to deter potential adversaries (Reiter and Poast 2021). They could produce a deterrent effect through three distinct mechanisms. First, forward-deployed troops could function as a form of deterrence by defense or denial. Deployments could shift the local balance of military capabilities in a way that convinces a potential adversary that challenging the status quo is unlikely to succeed (Reiter and Poast 2021, 47; Wilner and Wenger 2021). In practice, however, forward deployments are usually too small to materially influence the local balance of military capabilities. Thus, arguments about how they achieve deterrence often rely on other logics.

The second mechanism through which forward deployment might produce deterrence is as an investment that signals the extent of the deploying state’s interest in the maintenance of the status quo. This is what Fearon referred to as the logic of “sunk costs.” Potential adversaries might distinguish credible from incredible deterrent threats by observing differences in the behavior of the deterrer (Fearon 1997). If only committed deterrers would be willing to incur the costs associated with basing, say, a thousand soldiers on another country’s territory, then credible threats might be distinguished from incredible threats by observing the presence or absence of contingents of forward-deployed troops. Yet this logic, too, is questionable as a full explanation for the deterrent capacity of forward deployment. Theoretically, “sunk cost” signals may easily be interpreted as bluffs, which means that the costs must be quite large in order to achieve reliable deterrence in the presence of even small levels of uncertainty about intentions (Fearon 1997, 75–78). Empirically, as noted above, forward deployments are typically composed of

relatively small numbers of personnel. They often do not involve the construction or maintenance of permanent base facilities, which would constitute a more costly and visible investment in a partner's security (Lostumbo et al. 2013). Indeed, because some allies underwrite basing costs, forward deployment may actually not be all that expensive when compared to the cost of maintaining a similar level of operational capacity based from the United States (Brooks, Ikenberry, and Wohlforth 2012, 16–17).

The limitations of these mechanisms help to explain the popularity among academics and practitioners of a third logic: that forward-deployed military personnel may function as pre-commitment devices, increasing the cost and difficulty of refusing to escalate if they come under attack (Fearon 1997). This function of forward deployment is often referred to as “tripwire” deterrence. Tripwire logic undergirds policy-oriented arguments both in favor of and against the maintenance of a forward U.S. military posture, and has often been invoked by analysts to assert that the credibility of security commitments can be substantially enhanced by stationing only a very small contingent of troops within a security partner's borders. Schelling memorably illustrated this claim with an argument about how the few troops of the U.S. Army's Berlin Brigade could seemingly hold back the entire Soviet military. The role of the Berlin Brigade was not to fight the Red Army, he claimed, but to “die. . . in a manner that guarantees that the action cannot stop there” (Schelling 1966, 47).

Although these claims are frequently adduced, the logic by which the tripwire mechanism enhances credibility is rarely fully specified. We argue below that the most compelling formulation of this theoretical link relies on domestic public opinion, particularly on the assumption that attacks on troops deployed overseas are likely to produce strong public support for escalation. We refer to this phenomenon as the tripwire effect. The tripwire effect (though not limited in scope to attacks on small deployments) is what, in theory, gives the kind of small deployments that are often labeled tripwires their deterrent force. However, it has never been systematically empirically assessed.

Hands-Tying Signals, Domestic Audience Costs, and The Tripwire Effect

Tripwire deterrence depends on the logic of hands-tying. According to this logic, credibility results from the visible construction of a situation in which backing down from a threat (or reneging on a promise) would be prohibitively costly. The existence of this sort of pre-commitment device should, in theory, change a leader's cost-benefit calculus in ways that incentivize upholding the commitment if challenged, even if other conditions have changed in ways that reduce the desirability of following through. This distinguishes hands-tying signals from sunk-cost signals, which Fearon suggests are less effective (Fearon 1997, 71).

The most well-studied pre-commitment device in international relations (IR) is the public threat, which is at the center of research about audience costs. Public threats deter by creating an expectation of follow-through on the part of domestic audiences. Leaders who fail to follow through will, according to audience cost theory, be punished even if conditions since the threat was made have changed in ways that reduce the desirability of escalation. A similar deterrent mechanism, which Fearon also categorized as a hands-tying signal, is the public alliance commitment. Broken alliance commitments might create public concerns about military reputation, non-military reputation, and whether or not the state is fulfilling its moral obligations (Tomz and Weeks 2021, 3). Audiences may even exact costs from leaders who exit alliances in peacetime (Levin and Kobayashi 2022) and from those who fail to support alliance partners in situations not explicitly made part of an alliance (Guenther and Musgrave 2022).

Both hands-tying signals work by creating an expectation that a leader will pay a cost *domestically* for failing to follow through on the commitment. Fearon explicitly suggests that this is typical of hands-tying signals more generally: “a tying-hands signal typically works by creating audience costs that the leadership would suffer due to the reaction of domestic political audiences to a perceived failure in the management of foreign policy” (Fearon 1997, 70).

Fearon suggested that a third example of hands-tying signal was the stationing of “small ‘trip-wire’ forces” in a “threatened area” (Fearon 1997, 70). On this view, tripwire deterrence should function—like public threats and alliance commitments—by creating an expectation that a leader will be punished domestically for failing to escalate in response to an attack on forward-deployed troops. The crucial role of the domestic audience in this logic has perhaps been obscured because not all analysts fully specify how tripwire deterrence works. For example, Schelling’s famous description of the Berlin Brigade does not contain a clear, complete description of the causal linkages between attack and escalation. More recent applied work also occasionally omits the intervening causal steps. Zapfe, for instance, argues that “the main function of NATO’s EFP is to help deter a conventional Russian attack by providing a tripwire, the engagement of which would all but guarantee that the Alliance as a whole would respond in some way,” but does not explain how the tripwire guarantees a response (Zapfe 2017, 150).

Yet some analysts do more fully describe a causal logic linking attacks on forward-deployed troops to escalation, and these typically rely on domestic political pressure. Glenn Snyder argued during the Cold War 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” (Snyder 2015, 130–31). Rovner and Talmadge note that “light presence” deployments “often deliberately create a ‘tripwire’ ensuring that regional aggression will necessarily entail early engagement with the hegemon.” They argue that Cold War-era U.S. deployments in Germany worked by “guaranteeing public support for European allies” in case of a Soviet attack (Rovner and Talmadge 2014, 554). Fuhrmann and Sechser argue that while forward-deployed nuclear forces are “militarily superfluous,” they may nevertheless deter because their destruction during an initial attack on the partner might produce “domestic political pressure for the patron to enter the war wholeheartedly” (Fuhrmann and Sechser 2014, 923). Slack 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” (Slack 2018, 25, 31). A RAND study asserts that forward-deployed troops “deter potential adversaries” because attacks 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). Glaser notes that “[i]t would be very difficult politically for the United States to renege on a security guarantee if U.S. troops were already caught up in the fighting” (Glaser 2017, 5).

A large “tripwire effect” on public support for escalation thus constitutes the most commonly articulated theoretical foundation for the strong, nearly automatic link between attacks on forward-deployed troops and escalation that many analyses and invocations of tripwire deterrence posit. Though alternative logics—involving, for instance, concerns about international reputation, or the influence of elites—are conceivable, they are rarely explicitly invoked. This may be partially because (as Fearon’s argument suggests with reference to “international” audience costs) these mechanisms are less compelling theoretical bases for strong claims about effective tripwire deterrence that are valid regardless of variation in the identities or dispositions of particular leaders or advisors (Fearon 1994, 581).

Empirical (In)attention to the Tripwire Effect

Given the dependence of tripwire deterrence on the audience cost mechanism, it is striking that there has been little sustained empirical attention to the question of whether public opinion would demand escalation in the case of an attack on forward-deployed troops.

That deficit contrasts starkly with a relative abundance of research on other kinds of hands-tying signals. A substantial amount of empirical work has investigated the effect of public threats on public support for escalation (Trager and Vavreck 2011; Levendusky and Horowitz 2012; Davies and Johns 2013; Chaudoin 2014). Alliances have also received significant empirical attention, including two recent studies that directly explored the influence of alliance commitments on public support for military intervention (Tomz and Weeks 2021; Berejikian and Justwan 2022).

Forward deployment has received less attention. A few studies have explored the determinants of forward deployment patterns, and some have investigated the association between forward deployment and phenomena such as the incidence of interstate conflict, free-riding behavior, public opinion, and other dynamics relating to the *host* state (Machain and Morgan 2013; Allen, VanDusky-Allen, and Flynn 2016; Jakobsen and Jakobsen 2019; Allen et al. 2020, 2022, 2023; Altman 2020, 511–14). One recent experimental study investigated whether different kinds of forward-deployed U.S. military assets reassure foreign elites (Blankenship and Lin-Greenberg 2022).

What is missing from this literature is a direct attempt to explore the question at the heart of tripwire deterrence: Do attacks on forward-deployed troops increase support for escalatory military intervention among audiences in the state whose troops are deployed? If so, is this tripwire effect strong enough to plausibly support the kinds of claims about a tight link between attacks on forward-deployed troops and escalation that are often implicated in articulations of the logic of tripwire deterrence?

There are good reasons to question the size—and even the existence—of the tripwire effect. Evidence from research on military casualties and public support for war sits awkwardly with 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 an *ongoing* conflict by becoming less supportive of continuing the intervention—not by demanding further escalation (Mueller 1973; Gartner and Segura 1998; Gartner 2008; Kriner and Shen 2012; Fazal 2021). These findings are broadly consistent with a portrait of a U.S. public that weighs questions related to the use of military force in a way that is reasonably sensitive to costs and stakes—Americans are more likely to support interventions when the ratio between stakes and costs is high (Eichenberg 2005; Gelpi, Feaver, and Reifler 2005). Tripwire deterrence, by contrast, is premised on the notion that attacks on forward-deployed troops should trigger a reaction that overrides these kinds of calculations about the intrinsic importance or difficulty of protecting a partner. It is true that the context within which the tripwire effect is supposed to function (prior to the outbreak of open conflict) likely differs in important ways from the context of an ongoing conflict. These contextual differences may change how the public would evaluate the costs and benefits of using force overseas. However, these differences remain untheorized and empirically unexamined.

Moreover, while observational analysis of the tripwire effect is complicated by a relative paucity of data, states hosting U.S. troops or the troops themselves have occasionally been attacked. These cases include North Korea's seizing of the *U.S.S. Pueblo* in 1968 (1 sailor killed and 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); and recent Iranian attacks on U.S. troops stationed in both Iraq and Syria.

None of these attacks prompted irresistible U.S. public demands for escalation. For instance, the 1976 attack in the DMZ occurred during a U.S. presidential election, when the pressures of public opinion should have been especially intense. Despite that context, President Gerald Ford responded not with armed retaliation but with a show of force and a demand for a North Korean apology (Sander 2017). Similarly, after the October 23, 1983, bombing of the Marine barracks in Beirut, an October 25 ABC News/*Washington Post* poll found that 51 percent of Americans disagreed with the statement “The United States should find and militarily punish the groups of nations responsible for the terrorist attack on the Marines even if it means risking a larger war,” while only 39 percent agreed (ABC News/*Washington Post* 1983). In an October 28 poll by the same firm, only 17 percent favored sending more troops to Lebanon, while 37 percent favored removing the troops from Lebanon (41 percent favored leaving the number of troops in Lebanon about the same)—even though 73 percent agreed that if the U.S. Marines were withdrawn, “Syria, backed by the Soviet Union, will take over Lebanon” (ABC News 1983).

Some research has addressed related issues, but no prior studies have been primarily or explicitly designed to explore the tripwire effect. Tomz investigated the influence of different levels of escalation on the disapproval stemming from an empty threat (Tomz 2007). To do so, he included experimental scenarios in which a U.S. president sent troops to a threatened region and launched an attack that resulted in casualties among U.S. troops. Results suggested that increasing levels of escalation lead to increased disapproval when the president backs down from a threat, and that disapproval was highest when casualties were incurred. While potentially consistent with a relatively large tripwire effect, this study does not directly evaluate the mechanism underlying tripwire deterrence. It does not describe a foreign attack on pre-positioned troops, it does not measure change in preferences over policy, and its design does not allow for the clean identification of the influence on presidential disapproval of the failure to respond to casualties incurred among troops deployed abroad (Nomikos and Sambanis 2019).

More recently, Allison, Herzog, and Ko found in a survey experiment that hypothetical U.S. military casualties during a North Korean attack do not change U.S. public support for nuclear war (Allison, Herzog, and Ko 2019). Although relevant to debates over tripwire deterrence, the study has limitations as a test of the tripwire effect. First, the authors do not report results related to the effect of casualties on public support for *conventional* escalation. The tripwire effect might be insignificant in the context of decisions about launching nuclear weapons, but still be important for decisions about less drastic forms of escalation. Second, the tripwire manipulation is appended to the description of a “mass casualty” event, such as might result from a nuclear attack (Allison, Herzog, and Ko 2019, 17). This allows only for an inference about the size of the tripwire effect in that particular context. It could be that information about the occurrence of “mass casualties” overwhelms information about the fact that some of those casualties occurred among military personnel. This is especially likely if, as some analysts have suggested, civilian casualties also produce public support for escalation (Freedman and Michaels 2019, 359 note 6; Schelling 1960, 136 note 13). Third, the treatment is dichotomous and does not provide subjects with information about the *number* of military casualties, which could plausibly be important. Finally, the design does not allow for an assessment of the influence of the *presence* of U.S. troops within an attacked state, only of the effect of any casualties incurred. Some formulations of the logic of tripwire deterrence imply that what matters most is *involvement*, or that the public commitment embodied by the presence of troops has been challenged (Fearon 1997; Zapfe 2017).

Hypotheses

The foregoing discussion has demonstrated that tripwire deterrence is a widely accepted theoretical dynamic implicated in a key feature of U.S. foreign policy—its forward military posture. It may also be based on questionable assumptions about the tripwire effect, which, in turn, is almost entirely unexamined empirically and which some historical examples suggest may be negligible or oppositely signed. Establishing the magnitude and direction of the tripwire effect is thus important.

To do so, we test the following hypotheses.

Casualties Produce Higher Support for Involvement

Schelling's classic early statement of tripwire logic suggested that the heroic deaths of members of the Berlin Brigade could produce pressure for escalation. As noted above, other analysts have made similar claims about the function of forward-deployed troops in other contexts. We operationalize this as a general hypothesis linking military casualties and support for war:

H1: *The higher the number of casualties resulting from an attack on U.S. troops, the greater the support among Americans for escalation.*

It is worth noting that this hypothesis stands in stark contrast to the consistent finding from work on casualties and public support for *ongoing* military interventions described above. Results related to this hypothesis may thus also have implications for that body of research. However, because there are potentially significant theoretical differences between the context within which the tripwire effect is thought to function and the context of an open military conflict, we are hesitant to explicitly frame our study as a test of the relationship between casualties and support for war in general.

Troop Presence Produces Higher Support For Involvement

While Schelling's theoretical treatment suggests that the valiant deaths of heroic troops trigger public support for intervention, Fearon proposes a reputational tripwire logic. For Fearon, tripwire deployments (like public threats and formal alliances) serve as public markers of commitment that “engage the national honor” (Fearon 1994, 581). Audience costs arise, in this model, when adversaries challenge commitments and leaders do not respond—this generates concerns among domestic audiences about a “loss of credibility, face, or honor” (Fearon 1994, 581). The logic of this mechanism does not rely on the attack producing casualties—indeed, Fearon (1994, 1997) never mentions military casualties in his treatments of the link between forward deployment and audience costs. Rather, the significant element is that a public commitment to protect a partner (signaled by the presence of troops on the ground within the partner's territory) has been called into question. Thus, Fearon's tripwire logic—unlike Schelling's—suggests a relationship between troop presence in an attacked country and support for escalation, regardless of the occurrence of casualties:

H2: *The presence of U.S. troops in a country should increase support among Americans for escalation after an attack on that country.*

Experimental Testing

To test these hypotheses, we employ survey experiments.¹ Survey experiments have long been used to evaluate the mechanisms underlying other hands-tying signals.

¹The experiments contained herein were reviewed by IRB and declared exempt from IRB review as protocol #1803007840 by the Cornell Institutional Review Board. The CES experiments were reviewed as part of the general CES review procedures.

One of the earliest direct applications of experimental methods to IR aimed to measure the strength of audience costs generated by public threats (Tomz 2007). More recent experimental research has continued to explore public threats and the influence of alliance commitments on public support for war (Levendusky and Horowitz 2012; Davies and Johns 2013; Levy et al. 2015; Kertzer and Brutger 2016; Thomson 2016; Lin-Greenberg 2019; Nomikos and Sambanis 2019; Croco, Hanmer, and McDonald 2021; Tomz and Weeks 2021; Berejikian and Justwan 2022).

Survey experimental designs are valuable in this field for two reasons. First, the mechanisms underlying each of these hands-tying signals directly implicate public opinion, pointing toward empirical approaches capable of assessing the drivers of change in public support for military intervention. Second, because of problems arising from strategic behavior, inconsistent data availability, and the co-occurrence of multiple forms of tied-hands signaling, observational designs are limited by endogeneity and other concerns that threaten the identification of causal effects. Survey experiments solve this problem by randomly assigning treatments, thereby facilitating causal inference.

We employ two different experimental designs: a pair of conjoint experiments and a pair of vignette designs. Using different sorts of designs enables us to address theoretically relevant factors from multiple angles. Conjoint experiments, for example, enable researchers to vary many types of causal factors, but are costly and difficult to field on nationally representative samples. Vignette designs, by contrast, may be easier to understand and are comparatively easy to field on nationally representative samples, but are limited in the number of factors they can test. We describe the designs and report the results in turn.

Conjoint Experiments

Conjoint designs simultaneously vary many different factors that may influence respondents' choices (Hainmueller, Hopkins, and Yamamoto 2014; Hainmueller and Hopkins 2015). In doing so, they allow researchers to specify the values of a variety of variables that respondents might otherwise believe co-vary (which could threaten inference by undermining the assumption of information equivalence—see Dafoe, Zhang, and Caughey 2018). Conjoint studies “allow [researchers] to simultaneously manipulate the litany of factors . . . about which IR scholars have offered a variety of—often contradictory—theoretical expectations” (Kertzer, Renshon, and Yarhi-Milo 2021). Conjoint studies in particular perform well even if they vary information about a large number of dimensions (Bansak et al. 2021) and can recover estimates that closely match those returned by more conventional survey methods as well as real-world behavioral measures (Hainmueller, Hangartner, and Yamamoto 2015). Despite occasional concerns about the use of conjoint estimates to yield information about preferences (Ganter 2021), these are not universally shared (Bansak et al. 2022). Applications of conjoint methodologies are increasingly popular in research on international relations and foreign policy (Clary and Siddiqui 2021; Escribà-Folch, Muradova, and Rodon 2021; Kertzer, Renshon, and Yarhi-Milo 2021; Leal and Musgrave 2022, 2023).

Our conjoint surveys presented respondents with pairs of scenario profiles that varied along several theoretically relevant dimensions (specified 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. Each respondent completed this task five times. This description accurately casts these tasks as abstract, but recent research confirms that abstract designs yield useful insights about the real world (Brutger

et al. 2022).² Importantly, this design offers a significant advantage over the more traditional experiments described later: The conjoint method allows us to explicitly manipulate information that respondents in our vignette surveys may infer from the different contexts of those scenarios, bolstering our ability to cleanly identify the effect of troop presence and casualties on support for escalation and allaying concerns about information equivalence.

Conjoint Experiment 1

We fielded our first conjoint experiment on Amazon’s Mechanical Turk website between April 13 and April 26, 2018. Mechanical Turk is widely used for social science surveys (Krupnikov, Nam, and Style 2021) and has been shown to replicate results from representative surveys (Coppock 2019). We administered the survey via Qualtrics using Strezhnev et al.’s Conjoint Survey Tool (Strezhnev et al. 2014). A total of 1,226 respondents completed the survey. Because our survey was fielded during a time of heightened concern regarding “bots” or low-quality respondents on MTurk (Dreyfuss 2018), we took steps to identify (on a very conservative basis) potentially suspicious respondents. These left us with 912 valid responses.³ Each respondent completed five tasks, giving us a total N of 4,560. Further information, including a discussion of respondent quality, is available in the online appendix; including the excluded respondents does not substantively change the results of the analyses described below.

Table 1 summarizes the variables included in profile descriptions. Our primary objective in this survey was to estimate the effect of the presence of U.S. troops in an attacked country on support for intervention (H2). We thus varied information about whether and how many U.S. troops were deployed inside the victim country. This ranged from zero troops deployed to tens of thousands. These levels capture the range of real-world variation in the size of meaningful American overseas military contingents (Allen et al. 2020). Our smallest non-zero level (“a few hundred military personnel”) is about equivalent to the current U.S. military presence in Latvia; our second level (“a few thousand military personnel”) is akin to troop levels in Qatar or Kuwait; and our highest level (tens of thousands) corresponds to deployments in Japan, Germany, and South Korea.

We included several other variables in our profile descriptions. These included information about U.S. trade relationships with the victim and the aggressor, a way of measuring interests; 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 (Mueller 1973; Gelpi, Feaver, and Reifler 2005; Howell and Pevehouse 2005; Voeten 2005; Grieco et al. 2011). These variables 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.

Using a conjoint design allows us to simultaneously manipulate all of these factors. This facilitates the separation—and identification—of any potential tripwire effect from the influence of variation in the costs, stakes, and likely outcomes of inter-

²Subsequent advances in conjoint methodology may mean that future work might wish to consider using non-uniform profile distributions for similar experiments (de la Cuesta, Egami, and Imai 2022).

³A scholarly discussion of these troubles can be found in Ahlers et al. (2021). Their suggested solutions, including the use of IP addresses to block known low-quality workers, approximate the way that we approached these solutions in real time. Their other solutions, including using more vetted survey platforms, similarly parallel our choice of SSI for the second conjoint experiment. Notably, we made these adjustments in real time during the early stages of the bot panic.

Table 1. Complete list of conjoint attributes and levels

Attribute	Levels
U.S. military presence in victim country	<ul style="list-style-type: none"> • No U.S. military personnel present • A few hundred military personnel • A few thousand military personnel • Tens of thousands of military personnel
Casualties suffered already (<i>Conjoint Experiment 2 only</i>)	<ul style="list-style-type: none"> • No • Several • Dozens
Probability of U.S. intervention succeeding	<p data-bbox="659 662 1129 710">(If no personnel are present in previous condition, no casualties suffered)</p> <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 U.S. military casualties resulting from an intervention	<ul style="list-style-type: none"> • Severe (More than 10,000 U.S. casualties) • Moderate (Several thousand U.S. casualties) • Light (No more than several hundred U.S. 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	<ul style="list-style-type: none"> • Major • Minor • Not a significant U.S. 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

Table 1. Continued

Attribute	Levels
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	<ul style="list-style-type: none"> • Major • Minor • Not a significant U.S. trade partner
U.S. civilians at risk in victim country	<ul style="list-style-type: none"> • Fewer than 100 • Several thousand • More than 100,000
Risk of aggressor retaliation against U.S. 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 U.S. intervention on behalf of the victim
Congressional authorization	Likely/unlikely that Congress will formally approve U.S. intervention

ventions (Eichenberg 2005; Gelpi, Feaver, and Reifler 2005; Nomikos and Sambanis 2019), which might otherwise be inferred from information about forward deployment. If the conventional wisdom about tripwire deterrence is correct, commitment devices should work independent of information about prospective losses or gains resulting from the conflict. Indeed, the entire point of the tripwire mechanism is that, at the moment of decision, the option of failing to respond has been made substantially less attractive regardless of variable factors like the intrinsic importance of the victim and the probable outcome of the conflict. These considerations primarily motivated our decisions about what factors to include in the conjoint survey. However, including these variables also allows us to assess the magnitude of the effect of forward deployment relative to other salient factors.

Figure 1 presents our results. Our core quantity of interest is the coefficient on “U.S. Troop Presence in Victim Country.” We find that there is no difference between having no U.S. troops and a few hundred military personnel in a victim country, but that there is a small increase (about three percentage points, and statistically distinguishable from zero at conventional levels) associated with moving from zero to a few thousand military personnel ($p = 0.03$) or to tens of thousands of personnel ($p = 0.04$).

This is consistent with H2, but some caveats are in order. First, treaties increase support for intervention by almost twice as much, a difference of 5.4 percentage points ($p = 0.00$). So does anticipated Congressional approval (4.3 percentage points, $p = 0.00$) and UN Security Council authorization (5.8 percentage points,




Figure 1. Results for the first conjoint experiment, administered on Mechanical Turk in April 2018. *N* of 4,560 observations (912 respondents). Note modest effects for tripwire deployments and relatively larger effects for treaty of alliance.

$p = 0.01$). Other variables are even more important. A very high risk of retaliation reduces support by 7.8 percentage points relative to very low risk. The expected costs and outcome of the intervention are by far the most important factors. The expectation of “severe” casualties reduces support by seventeen percentage points relative to the “light casualty” base category. Being informed the intervention will almost certainly fail reduces support by 20.7 percentage points relative to almost

certain success.⁴ Taken together, these results show that although relatively large contingents of forward-deployed troops may increase support for intervention, the effect is small.

In the online appendix, we present evidence that this pattern does not substantially vary across different subgroups, including gender, party identification, and education level (although respondents who score high on an index of internationalism and college-educated respondents are more sensitive to both treaties and troop placements).

Conjoint Experiment 2

The first conjoint experiment allows us to test H2, regarding the effect of having troops stationed in a country (along the lines of Fearon's formulation of the tripwire argument), but it does not allow us to test the effect of casualties on public opinion. Our second conjoint experiment accordingly extends Experiment 1 by adding an attribute that manipulates U.S. casualties suffered in an attack on the victim by the aggressor, ranging from none to dozens of U.S. soldiers killed, with two intermediate levels (several American soldiers wounded but none killed; and several American soldiers killed). (We restricted randomization such that countries with no troops deployed could not have any casualties.)

We recruited a sample via Survey Sampling International to be nationally representative with respect to age, gender, and political party affiliation. After screening (see online appendix), we were left with 818 respondents who successfully completed seven tasks for a total of 5,726 observations. Respondents were recruited between August 15 and August 25, 2018.

Figure 2 shows the results. Although the estimates for troop deployments remain positive, they are no longer statistically significant. The estimates for coefficients of casualties already sustained are negative, but not statistically significant. The estimated effect of a treaty commitment is slightly smaller, at 2.4 percentage points, but remains statistically significant at conventional levels ($p = 0.01$). The effects of likely future casualties on support for an intervention are attenuated ("severe" casualties now reduce support by only about 6.4 points) but remain negative and statistically significant. The chances of success continue to matter, with an invasion that "probably will not" or "almost certainly will not" succeed eleven percentage points less likely to be supported for an intervention compared to "almost certainly will succeed." So does the approval of the UN Security Council and of Congress (both of which move support for intervention by about 2.5 points, significant at conventional levels) and the victim's status as a major trading partner (3.6 points, significant at conventional levels). These patterns do not vary substantially across gender, party identification, internationalist attitudes, and education levels, although Republicans are more sensitive to expected casualties, and the college-educated and the highly internationalist respondents are more sensitive to treaty status. The online appendix features more information regarding these subgroup analyses.

The second conjoint experiment, in other words, points to a rejection of both H1 and H2. Casualties do not increase support when we specify all other factors, and troop deployments do not either. Our confidence in these results is increased by the fact that, as with the first experiment, other coefficients seem plausible. Costs and benefits, interests (in the form of trade relationships), and congressional and United Nations approval influence attitudes toward intervention in the expected

⁴Interestingly, nuclear status has no notable influence, with the exception that respondents were less likely to support intervention against aggressors with "limited" nuclear arsenals. This finding is consistent with the implications of arguments about the instability of crises involving states whose nuclear arsenals may not be large enough to guarantee a secure second strike (Narang 2014); it may also be facilitating an inference about the identity of the adversary, and thus suggest greater support for escalating against Russia (a state with a substantial nuclear arsenal) than against China or North Korea (states with more limited nuclear arsenals). However, given that the result does not replicate in the second conjoint survey (described below), we are hesitant to interpret this as more than statistical noise.




Figure 2. Results for the second conjoint experiment, administered on SSI to an *N* of 818 respondents for 5,726 observations in August 2018. Note again the low (and insignificant at conventional levels) effects of tripwire deployments and the relatively larger effects of treaties of alliance, as well as the test of casualties showing these effects to be negative but insignificant.

directions. In line with a broad swath of the literature on war and public opinion, our evidence supports explanations for variation in support for escalation based on a rough cost-benefit analysis, one that incorporates concerns about legitimacy and reputation. Tripwire effects, however, are difficult to find.

Vignette Studies

Our conjoint design cannot specify actual countries and may thus be less realistic and concrete than a vignette design. To address this concern, we also report results from a vignette experiment based on descriptions of scenarios in which U.S. troops based in South Korea and Estonia come under attack by forces from North Korea and Russia, respectively. The vignette study cannot assess the influence of the presence of U.S. troops on U.S. public support for escalation, since it is premised on troops being present, but it can assess the influence of different numbers of casualties suffered in a hypothetical attack.

We fielded this experiment within a module of the 2018 Cooperative Congressional Election Study (CCES, now the Cooperative Election Study or CES; $N = 1,000$). Full details are presented in the online appendix. This vignette was re-run in 2020, allowing us an unusual opportunity to replicate our analysis with another sample of the same size.⁵ We used a 2×3 factorial design varying the location of the scenario and the level of casualties sustained by U.S. forces in an attack:

As you know, many American military personnel are deployed to countries where the United States is not currently fighting a war, such as [*Treatment 1a*]. Imagine that fighting has broken out between [*Treatment 1a*] and [*Treatment 1b*]. [*Treatment 2*] U.S. military personnel have been killed by [*Treatment 1b*] forces.

U.S. policymakers are considering whether to commit additional forces to defend [*Treatment 1a*] or to pull back U.S. troops. Committing additional troops could lead to an escalation of the conflict, while pulling back could lead to the defeat of [*Treatment 1a*].

The dependent variable was “Would you support committing additional U.S. troops to this conflict?” Respondents could answer “Yes, I support committing additional troops” or “No, I do not support committing additional troops.” The dependent variable, support for committing additional troops, is relatively conservative and asks only for evidence that the public would favor an escalatory response, though not necessarily an exceptionally risky one.

Treatment 1 varied information about the U.S. partner (1a: South Korea or Estonia) and the aggressor (1b: North Korea or Russia). This enables us to explore whether tripwire mechanisms are influenced by contextual factors. We chose the Korean Peninsula and the Baltics because they represent regions where the United States has actually deployed troops. Consequently, these choices correspond to real-world scenarios.

Treatment 2 varies U.S. fatalities. It specifies whether none, a few, or dozens of U.S. troops have been killed. These casualty levels were chosen for their realism. A few U.S. military personnel killed would correspond to a grave but not unusual incident, such as the attacks on the USS *Cole* or the Khobar Towers apartments in Saudi Arabia, while “dozens of U.S. military personnel” would correspond to a severe shock, such as the numbers of American servicemembers killed in the Second Battle of Fallujah over weeks of fighting or the bombing of the Marine barracks in Beirut. Varying these enables direct tests of H1. We code the casualties treatment in two ways: as a three-level variable (none, a few, or dozens) or as a simple binary (none or any). This enables us to test whether *any* deaths, or only *substantial* numbers of deaths, affect public opinion.

Figure 3 presents the raw distribution of results, divided by scenario and by fatalities. The results suggest greater support overall for interventions in the Korean scenario than in the Russian scenario. The occurrence and number of fatalities resulting from the attack appear to matter very little, for the most part. In three of the

⁵For an effect size of 0.2 with 5 percent significance and a two-tailed test, 788 total respondents would be required to reach 80 percent power in a two-level treatment.

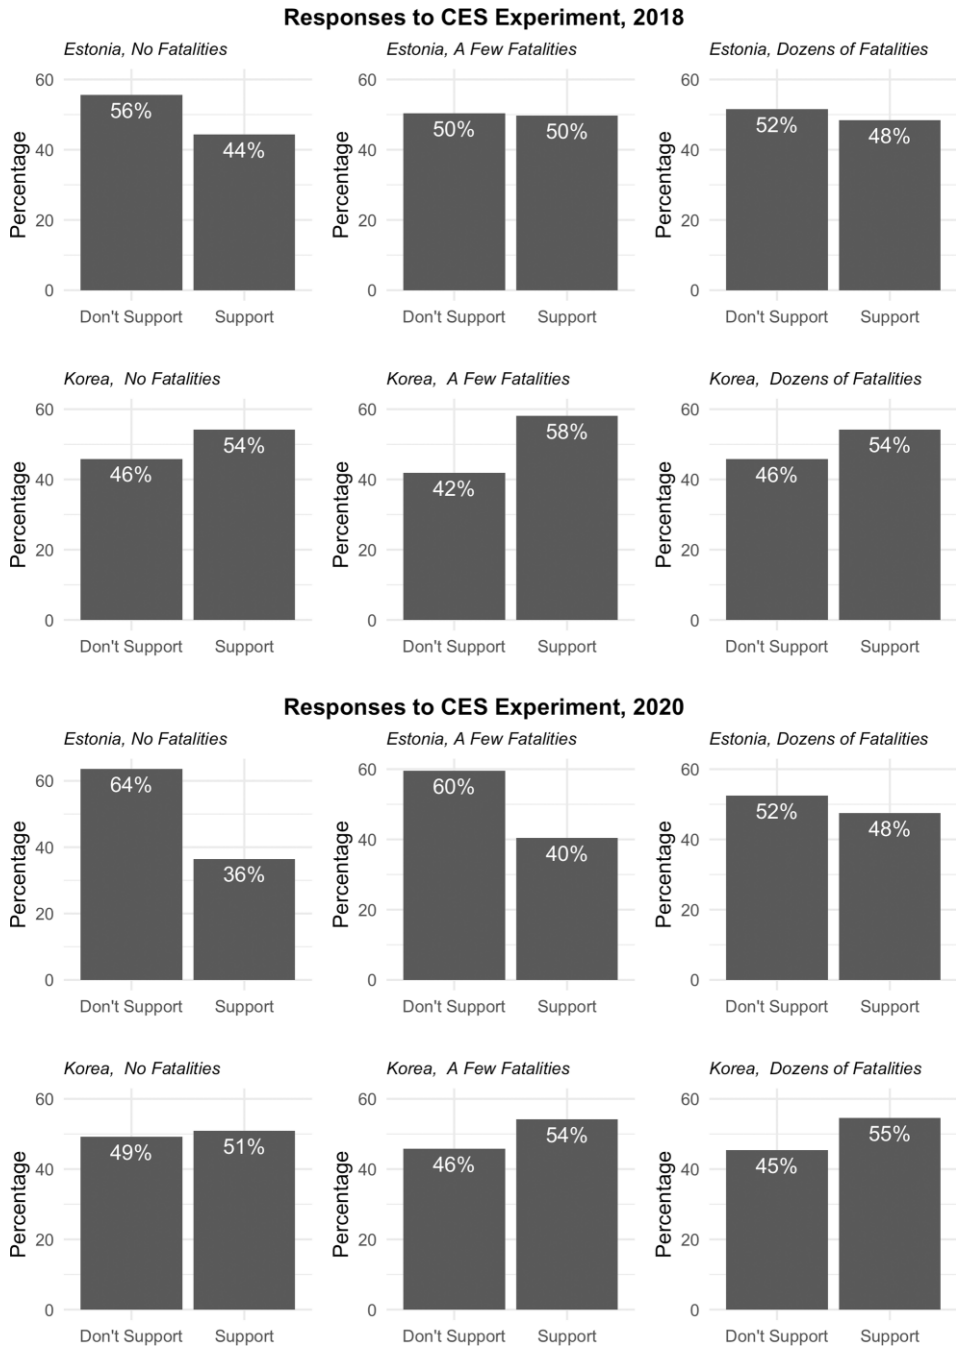


Figure 3. Raw results for the CES (formerly CCES) experiments from 2018 and 2020. $N(2018) = 979$; $N(2020) = 986$. Note the weak (at best) and inconsistent relationship between support for escalation and casualties incurred across each scenario.

four scenarios (Estonia 2018, Korea 2018, and Korea 2020), the difference between support for escalation in the “no fatalities” condition and the “dozens of fatalities condition” never exceeds four percentage points, and is not statistically distinguishable from zero. Only in the Estonia 2020 scenario do fatalities appear to boost public support for escalation, with a shift of twelve percentage points between the “no fatalities” condition and the “dozens of fatalities” condition. It is possible that some differences between the context of 2018 and the context of 2020 enhanced the effect of casualties on support for escalation. However, given the absence of a compelling theoretical explanation for this difference, and the fact that this is only one of four similar studies, the other three of which returned null results, we think it unwise to overinterpret this positive finding.


We use OLS regression to analyze these data (both pooled and separately) and present the results in [figure 4](#) (full results and alternative specifications are available in the online appendix). The OLS coefficients for the fatalities terms are positive but do not reach statistical significance. By contrast, the coefficients for scenario type do reach statistical and, more important, substantive significance in almost every specification ([Gross 2015](#)). Support for intervention in the Korean scenario is approximately ten percentage points higher than in the Russian scenario, suggesting that respondents’ preferences for interventions are shaped by contextual factors. The interaction term of scenario and casualties, however, does not reach significance. This confirms the visual results of [figure 3](#) in general, and suggests that the apparent effect of dozens of fatalities in Estonia 2020 is the result of factors other than the manipulations.

These results hold when we use the binary measurement of casualties (no deaths/any deaths), when we treat the trinary casualties treatment as a continuous variable, and when we use logistic regression instead of OLS (see online appendix). The results imply a rejection of H1. Increasing casualty levels do not appear to increase support for intervention. Further, although respondents were substantially more likely to favor involvement in Korea than against Russia, there was no difference between scenarios in terms of the effect of casualty sensitivity (as captured by the interaction term). That respondents’ answers about interventions varied with the scenario setting implies that they were in fact making calculations about costs and interests, and thereby lends credence to our null effect in estimates of H1.

In alternative specifications not including an interaction term, the story is slightly more complicated, as the coefficient for deaths occasionally *approaches* statistical significance at conventional levels. When measured as a binary, deaths are predicted to increase support by about four percentage points compared to the no-deaths condition, with a *p*-value under 0.10 but still >0.05 . When measured as a trinary variable, coefficients for “a few deaths” never approach conventional significance, but those for “dozens of deaths” are similar to the coefficient on “any deaths” in the binary model: roughly four to five percentage points in the pooled specifications. The estimate is conventionally and substantively significant in 2020, with an estimated increase in support of 7.7 percentage points for the dozens of deaths compared to the no-deaths condition. However, these results are dependent on model specification and appear only in the OLS but not the logistic regression specifications (see online appendix). In other words, although our experiment cannot conclude that the tripwire effect is *entirely* unsupported by evidence, any effect that our experiments detected was faint and substantively much less meaningful than would have been assumed.

The balance of the evidence, then, suggests that the magnitude of the tripwire effect lies somewhere between modest and null. Although the consistent positive coefficient signs for casualties mean that we cannot rule out a small effect, recall that the *claimed* size of the tripwire effect is meant to be large enough to bind a president to a conflict with a potential for nuclear war—to turn the sacrifice of the Berlin Brigade into a defense of Europe against Soviet aggression. Testing that infor-

Support for Intervention on CES Experiments, 2018 and 2020



OLS coefficients. CES sample. Ns: Base: 1967; Adjustments: 1965; 2018 only: 979; 2020 only: 986.

Figure 4. Ordinary least squares regression coefficients for tests of the CES (formerly CCES) experiments. Note that the tripwire variables and interaction terms do not reach conventional levels of statistical significance in any specification.

mal conjecture is difficult, but it seems reasonable that a quantitative translation of Schelling’s poetry would be an effect size on the order of ten to twenty percentage points. That is plainly inconsistent with our findings.

Table 2. Summary of tests

	H1: Casualties	H2: Presence
Conjoint 1	<i>Did not test</i>	Failed to reject: Some small support for intervention when victim country hosted thousands of U.S. troops
Conjoint 2	Rejected: No consistent effect of casualties on public support for intervention	Rejected: No effect of troop placement on public support for intervention
CES	Rejected: No consistent effect of casualties on public support for intervention	<i>Did not test</i>

Discussion of Experimental Results

Table 2 summarizes our findings. We find at most limited support for the existence of a tripwire effect. These results are at odds with the primary theoretical logic underlying claims about the strength of tripwire deterrence. Indeed, on balance, they are more consistent with results from the within-conflict casualty literature, which has consistently found that casualties reduce support for conflict. While pre-conflict casualties have little influence on support for intervention in either direction, results from our conjoint studies suggest that the expectation of further casualties produces large decreases in support for intervention. Though our experimental design does not allow us to evaluate this possibility, it could be that very large numbers of pre-conflict casualties—larger than those that we included in our manipulations—might *also* reduce support for escalation by increasing the expectation of future casualties.

It should be noted that our results, particularly in the conjoint experiment, have other implications. In particular, we find that treaty commitments do matter, even separate from other factors. Although our estimates of the size of the effect of treaty commitments are smaller than other recent studies, our estimate might more precisely capture the importance of the treaty itself rather than other factors that respondents may infer from the existence of a treaty (Tomz and Weeks 2021). Furthermore, our findings fit with the results of other recent experiments demonstrating that alliance commitments boost support for intervention after cyberattacks, with an effect size closer to what we find here and proportional to the apparent escalation risks (Guenther and Musgrave 2022). Formal treaty commitments, in other words, may have been somewhat unfairly dismissed by tripwire proponents.

Our findings also raise some avenues for future exploration. We find consistently higher support for intervention in Korea than in Estonia. We varied the location of our vignettes to test whether the tripwire effect differed, not to explain support for intervention per se. Although we have ruled out large differences in the tripwire effect, the differences we did find invite more research into how U.S. audiences gauge preferences for intervention—whether, for instance, they considered Russia a tougher foe or Estonia to be less important to U.S. interests. Our conjoint findings regarding the likely costs of an intervention and interests in the country being attacked suggest that both factors may be in play.

Conclusion

One attraction of tripwire deployments is their putative ability to deter 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. Across four separate survey experiments, conducted according to two very different designs, we have found—at best—small and inconsistent support for the claim that attacks on U.S. troops deployed overseas produce demands for escalation. Indeed, results from our conjoint surveys suggest that public support for intervention is driven primarily by factors related to stakes (as indicated by the influence of the victim's trade and formal alliance ties with the United States) and expected costs (as indicated by the influence of variation in likely *subsequent* casualties, assessments of the likelihood of intervention success, and assessments of the probability of retaliation by the target of the intervention). Relative to the influence of these factors—some of which prior research shows are also significant for explaining public support for ongoing military interventions—variation in the presence of U.S. troops in an attacked country and U.S. casualties resulting from an attack produces little, if any, increase in support for escalation.

We are, however, circumspect in interpreting these results. For one thing, while the logic of tripwire deterrence—as it has been articulated by academics and policymakers—typically relies on a domestic political mechanism, we do not rule out the possibility that 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 systematically respond much 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 are not adequately reflected in our sample (although recent work suggests that elites may be much like non-elites, save for compositional effects; Kertzer 2020). Alternatively, leaders might act out of the incorrect anticipation of strong domestic constraints (rooted, perhaps, in the outsized influence of writers like Schelling).

It could also be that our results suggest not that tripwire deterrence is a myth, but rather that it is simply not as powerful or universal as the conventional wisdom holds. For instance, our experimental designs were premised on the assumption that only small or moderate numbers of casualties should be sufficient to trigger a tripwire effect. However, it could be that "tripwires" are—in reality—not especially sensitive to small triggers, and that the magnitude of our treatment was thus too small. Perhaps hundreds of deaths would be necessary to produce significant demands for escalation—this would be especially interesting in light of the possibility that casualties suffered prior to the outbreak of a conflict could inform expectations about the severity of likely casualties resulting from intervention. Subsequent research might thus build on our analysis by varying the level of casualties across a greater range. Similarly, our dependent variables measured only respondents' willingness to become involved in or to commit additional troops to a conflict. It is possible that more fine-grained measurements, such as offering respondents choices between policy options characterized by different degrees of escalation, might return a different result. It could also be that contextual variables matter, and that the salience of real-world crisis conditions (something difficult to manipulate) would change our findings. Regardless of these caveats, however, our research still makes a valuable contribution by narrowing the conditions under which tripwire effects can be expected to operate.

Supplementary Information

Supplementary information is available at the *Foreign Policy Analysis* data archive.

Acknowledgments

We gratefully acknowledge comments and feedback from participants in workshops and presentations at Texas Tech University in October 2018, the George Washington University in February 2020, the Massachusetts Institute of Technology in September 2020, and Cambridge University in June 2022. We also acknowledge comments and feedback from participants in the 2019 annual meeting of the Midwest Political Science Association in Chicago, Illinois and the 2019 American Political Science Association annual meeting in Washington, DC. We also thank the editors and the anonymous reviewers for their comments and suggestions.

References

- ABC NEWS. 1983. *ABC News Poll: October 1983*. Cornell University, Ithaca, NY: Roper Center for Public Opinion Research.
- ABC NEWS/WASHINGTON POST. 1983. *ABC News/Washington Post Poll: October 1983*. Cornell University, Ithaca, NY: Roper Center for Public Opinion Research.
- AHLER, DOUGLAS J., CAROLYN E. ROUSH, AND GAURAV SOOD. 2021. "The Micro-Task Market for Lemons: Data Quality on Amazon's Mechanical Turk." *Political Science Research and Methods* FirstView, 1–20. <https://doi.org/10.1017/psrm.2021.57>.
- ALLEN, MICHAEL A., THOMAS CAMPBELL, NICOLAS HERNANDEZ, AND VALERYN SHEPHERD. 2023. "US Military Deployments and the Risk of Coup d'État." *Foreign Policy Analysis* 19(1): orac027. <https://doi.org/10.1093/fpa/orac027>.
- ALLEN, MICHAEL A., MICHAEL E. FLYNN, CARLA MARTINEZ MACHAIN, AND ANDREW STRAVERS. 2020. "Outside the Wire: US Military Deployments and Public Opinion in Host States." *American Political Science Review* 114(2): 326–41. <https://doi.org/10.1017/S0003055419000868>.
- ALLEN, MICHAEL A., CARLA MARTINEZ MACHAIN, MICHAEL E. FLYNN, AND ANDREW STRAVERS. 2022. *Beyond the Wire: US Military Deployments and Host Country Public Opinion*. Oxford: Oxford University Press.
- ALLEN, MICHAEL A., JULIE VANBUSKY-ALLEN, AND MICHAEL E. FLYNN. 2016. "The Localized and Spatial Effects of US Troop Deployments on Host-State Defense Spending." *Foreign Policy Analysis* 12(4): 674–94.
- ALLISON, DAVID M., STEPHEN HERZOG, AND JIYOUNG KO. 2019. "Under the Umbrella: Nuclear Crises, Extended Deterrence, and Public Opinion." *Journal of Conflict Resolution* 66(10): 00220027221100254.
- ALTMAN, DAN. 2020. "The Evolution of Territorial Conquest after 1945 and the Limits of the Territorial Integrity Norm." *International Organization* 74(3): 490–522. <https://doi.org/10.1017/S002081832000119>.
- BANSAK, KIRK, JENS HAINMUELLER, DANIEL J. HOPKINS, AND TEPPEI YAMAMOTO. 2021. "Beyond the Breaking Point? Survey Satisficing in Conjoint Experiments." *Political Science Research and Methods* 9(1): 53–71. <https://doi.org/10.1017/psrm.2019.13>.
- . 2022. "Using Conjoint Experiments to Analyze Election Outcomes: The Essential Role of the Average Marginal Component Effect (AMCE)." *SSRN Scholarly Paper*. Rochester, NY. <https://doi.org/10.2139/ssrn.3588941>.
- BEREJKIAN, JEFFREY, AND FLORIAN JUSTWAN. 2022. "Defense Treaties Increase Domestic Support for Military Action and Casualty Tolerance: Evidence from Survey Experiments in the United States." *Contemporary Security Policy* 43(2): 308–49. <https://doi.org/10.1080/13523260.2021.2023290>.
- BLANKENSHIP, BRIAN, AND ERIK LIN-GREENBERG. 2022. "Trivial Tripwires?: Military Capabilities and Alliance Reassurance." *Security Studies* 31(1): 92–117. <https://doi.org/10.1080/09636412.2022.2038662>.
- BROOKS, STEPHEN G., G. JOHN IKENBERRY, AND WILLIAM C WOHLFORTH. 2012. "Don't Come Home, America: The Case against Retrenchment." *International Security* 37(3): 7–51. https://doi.org/10.1162/ISEC_a_00107.
- BRUTGER, RYAN, JOSHUA D. KERTZER, JONATHAN RENSHON, DUSTIN TINGLEY, AND CHAGAI M WEISS. 2022. "Abstraction and Detail in Experimental Design." *American Journal of Political Science* EarlyView, 1–17. <https://doi.org/10.1111/ajps.12710>.
- CHAUDOIN, STEPHEN. 2014. "Promises or Policies? An Experimental Analysis of International Agreements and Audience Reactions." *International Organization* 68(1): 235–56. <https://doi.org/10.1017/S0020818313000386>.
- CLARY, CHRISTOPHER, AND NILOUFER SIDDIQUI. 2021. "Voters and Foreign Policy: Evidence from a Conjoint Experiment in Pakistan." *Foreign Policy Analysis* 17(2): orab001.

- COPPOCK, ALEXANDER. 2019. "Generalizing from Survey Experiments Conducted on Mechanical Turk: A Replication Approach." *Political Science Research and Methods* 7(3): 613–28. <https://doi.org/10.1017/psrm.2018.10>.
- CROCO, SARAH E., MICHAEL J. HANMER, AND JARED A. McDONALD. 2021. "At What Cost? Reexamining Audience Costs in Realistic Settings." *The Journal of Politics* 83(1): 8–22. <https://doi.org/10.1086/708912>.
- CUESTA, BRANDON DE LA, NAOKI EGAMI, AND KOSUKE IMAI. 2022. "Improving the External Validity of Conjoint Analysis: The Essential Role of Profile Distribution." *Political Analysis* 30(1): 19–45. <https://doi.org/10.1017/pan.2020.40>.
- DAFOE, ALLAN, BAobao ZHANG, AND DEVIN CAUGHEY. 2018. "Information Equivalence in Survey Experiments." *Political Analysis* 26(4): 399–416. <https://doi.org/doi:10.1017/pan.2018.9>.
- 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. <https://doi.org/10.1111/isqu.12045>.
- DREYFUSS, EMILY. 2018. "A Bot Panic Hits Amazon Mechanical Turk | WIRED." *Wired*, August 8, 2018. <https://www.wired.com/story/amazon-mechanical-turk-bot-panic/>.
- EICHENBERG, RICHARD C. 2005. "Victory Has Many Friends: US Public Opinion and the Use of Military Force, 1981–2005." *International Security* 30(1): 140–77.
- ESCRIBA-FOLCH, ABEL, LALA H. MURADOVA, AND TONI RODON. 2021. "The Effects of Autocratic Characteristics on Public Opinion toward Democracy Promotion Policies: A Conjoint Analysis." *Foreign Policy Analysis* 17(1): oraa016.
- FAZAL, TANISHA M. 2021. "Life and Limb: New Estimates of Casualty Aversion in the United States." *International Studies Quarterly* 65(1): 160–72. <https://doi.org/10.1093/isq/sqaa068>.
- 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. <https://doi.org/10.1177/0022002797041001004>.
- FREEDMAN, LAWRENCE, AND JEFFREY MICHAELS. 2019. *The Evolution of Nuclear Strategy: New, Updated and Completely Revised*. London: 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. <https://doi.org/10.1111/ajps.12082>.
- GANTER, FLAVIEN. 2021. "Identification of Preferences in Forced-Choice Conjoint Experiments: Reassessing the Quantity of Interest." *Political Analysis* 31(1): 98–112. December, 1–15. <https://doi.org/10.1017/pan.2021.41>.
- GARTNER, SCOTT SIGMUND. 2008. "The Multiple Effects of Casualties on Public Support for War: An Experimental Approach." *American Political Science Review* 102(1): 95–106. <https://doi.org/10.1017/S0003055408080027>.
- GARTNER, SCOTT SIGMUND, AND GARY M SEGURA. 1998. "War, Casualties, and Public Opinion." *Journal of Conflict Resolution* 42(3): 278–300.
- GELPI, CHRISTOPHER, PETER D. FEAVER, AND JASON REIFLER. 2005. "Success Matters: Casualty Sensitivity and the War in Iraq." *International Security* 30: 7–46. <https://doi.org/10.1162/isec.2005.30.3.7>.
- GLASER, JOHN. 2017. "Withdrawing from Overseas Bases: Why a Forward-Deployed Military Posture is Unnecessary, Outdated, and Dangerous." *Cato Institute Policy Analysis No. 816*, July 18, 2017.
- GRIECO, JOSEPH M., CHRISTOPHER GELPI, JASON REIFLER, AND PETER D FEAVER. 2011. "Let's Get a Second Opinion: International Institutions and American Public Support for War." *International Studies Quarterly* 55(2): 563–83.
- GROSS, JUSTIN H. 2015. "Testing What Matters (If You Must Test at All): A Context-Driven Approach to Substantive and Statistical Significance." *American Journal of Political Science* 59(3): 775–88. <https://doi.org/10.1111/ajps.12149>.
- GUENTHER, LINDSEY, AND PAUL MUSGRAVE. 2022. "New Questions for Old Alliances: NATO in Cyberspace and American Public Opinion." *Journal of Global Security Studies* 7(4): ogac024. <https://doi.org/10.1093/jogss/ogac024>.
- HAINMUELLER, JENS, DOMINIK HANGARTNER, AND TEPPEI YAMAMOTO. 2015. "Validating Vignette and Conjoint Survey Experiments against Real-World Behavior." *Proceedings of the National Academy of Sciences* 112(8): 2395–400. <https://doi.org/doi:10.1073/pnas.1416587112>.
- HAINMUELLER, JENS, AND DANIEL J HOPKINS. 2015. "The Hidden American Immigration Consensus: A Conjoint Analysis of Attitudes toward Immigrants." *American Journal of Political Science* 59: 529–48. <https://doi.org/10.1111/ajps.12138>.
- 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. <https://doi.org/10.1093/pan/mpt024>.
- HOWELL, WILLIAM G., AND JON C PEVEHOUSE. 2005. "Presidents, Congress, and the Use of Force." *International Organization* 59(1): 209–32.
- 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. <https://doi.org/10.1080/13523260.2018.1492066>.
- KERTZER, JOSHUA D. 2020. "Re-Assessing Elite-Public Gaps in Political Behavior." *American Journal of Political Science* 66(3): 539–53. <https://doi.org/10.1111/ajps.12583>.
- 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. <https://doi.org/10.1111/ajps.12201>.
- KERTZER, JOSHUA D., JONATHAN RENSHON, AND KEREN YARHI-MILO. 2021. "How Do Observers Assess Resolve?" *British Journal of Political Science* 51(1): 308–30. <https://doi.org/10.1017/S0007123418000595>.
- 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. <https://doi.org/10.1093/poq/nfs048>.
- KRUPNIKOV, YANNA, H. HANNAH NAM, AND HILLARY STYLE. 2021. "Convenience Samples in Political Science Experiments." In *Advances in Experimental Political Science*, edited by James N. Druckman and Donald P. Green, 165–83. Cambridge: Cambridge University Press.
- LEAL, MARCELO MESQUITA, AND PAUL MUSGRAVE. 2022. "Cheerleading in Cyberspace: How the American Public Judges Attribution Claims for Cyberattacks." *Foreign Policy Analysis* 18(2): orac003. <https://doi.org/10.1093/fpa/orac003>.
- . 2023. "Hitting Back or Holding Back in Cyberspace: Experimental Evidence Regarding Americans' Responses to Cyberattacks." *Conflict Management and Peace Science* 40(1): 42–64. <https://doi.org/10.1177/07388942221111069>.
- 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. <https://doi.org/10.1017/S002238161100154X>.
- LEVIN, DOV H., AND TETSURO KOBAYASHI. 2022. "The Art of Uncommitment: The Costs of Peacetime Withdrawals from Alliance Commitments." *European Journal of International Relations* 28(3): 589–615. <https://doi.org/10.1177/13540661221098221>.
- 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. <https://doi.org/10.1111/ajps.12197>.
- LIN-GREENBERG, ERIK. 2019. "Backing up, Not Backing down: Mitigating Audience Costs through Policy Substitution." *Journal of Peace Research* 56(4): 559–74. <https://doi.org/10.1177/0022343319832641>.
- LOSTUMBO, MICHAEL J., MICHAEL J. MCNERNEY, ERIC PELTZ, DEREK EATON, AND DAVID R FRELINGER. 2013. *Overseas Basing of US Military Forces: An Assessment of Relative Costs and Strategic Benefits*. Santa Monica, CA: Rand Corporation.
- MACHAIN, CARLA MARTINEZ, AND T CLIFTON MORGAN. 2013. "The Effect of US Troop Deployment on Host States' Foreign Policy." *Armed Forces & Society* 39(1): 102–23. <https://doi.org/10.1177/0095327x12442306>.
- MUELLER, JOHN E. 1973. *War, Presidents, and Public Opinion*. New York: Wiley.
- NARANG, VIPIN. 2014. *Nuclear Strategy in the Modern Era*. New York: Princeton University Press.
- 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. <https://doi.org/10.1177/0022343319839456>.
- O'MAHONY, ANGELA, MIRANDA PRIEBE, BRYAN FREDERICK, JENNIFER KAVANAGH, MATTHEW LANE, TREVOR JOHNSTON, THOMAS S. SZAYNA, JAKUB P. HLAVKA, STEPHEN WATTS, AND MATTHEW POVLOCK. 2018. *US Presence and the Incidence of Conflict*. Santa Monica: RAND Corporation.
- REITER, DAN, AND PAUL POAST. 2021. "The Truth about Tripwires: Why Small Force Deployments Do Not Deter Aggression." *Texas National Security Review* 4(3): 33–53. <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. <https://doi.org/10.1080/15325024.2014.935224>.
- SANDER, GORDON F. 2017. "When the U.S. Almost Went to War with North Korea." *POLITICO Magazine*, September 14, 2017. <https://www.politico.com/magazine/story/2017/09/14/north-korea-1976-axe-murder-incident-215605>.
- SCHELLING, THOMAS C. 1960. *The Strategy of Conflict*. Cambridge, MA: 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: Oxford University Press.
- SHLAPAK, DAVID A. 2018. “The Russian Challenge.” *RAND Arroyo Perspective*. Santa Monica, CA: RAND Corporation.
- SLACK, MATTHEW C. 2018. “Conventional Tripwire Deterrence.” *Dissertation*. Chicago, IL: University of Chicago.
- SNYDER, GLENN HERALD. 2015. *Deterrence and Defense*. Originally published 1961. Princeton, NJ: Princeton University Press.
- STREZHNEV, ANTON, JENS HAINMUELLER, DANIEL J. HOPKINS, AND TEPPEI YAMAMOTO. 2014. “Conjoint SDT.” May 16, 2014. <https://github.com/astrezhnev/conjointsdt>.
- THOMSON, CATARINA P. 2016. “Public Support for Economic and Military Coercion and Audience Costs.” *The British Journal of Politics and International Relations* 18(2): 407–21. <https://doi.org/10.1177/1369148115615030>.
- TOMZ, MICHAEL, AND JESSICA WEEKS. 2021. “Military Alliances and Public Support for War.” *International Studies Quarterly* 65(3): 811–24. <https://doi.org/10.1093/isq/sqab015>.
- TOMZ, MICHAEL. 2007. “Domestic Audience Costs in International Relations: An Experimental Approach.” *International Organization* 61(4): 821–40. <https://doi.org/10.1017/S0020818307070282>.
- 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. <https://doi.org/10.1111/j.1540-5907.2011.00521.x>.
- VANAGA, NORA, AND TOMS ROSTOKS. 2018. *Detering Russia in Europe: Defence Strategies for Neighbouring States*. New York: Routledge.
- VOETEN, ERIK. 2005. “The Political Origins of the UN Security Council’s Ability to Legitimize the Use of Force.” *International Organization* 59(3): 527–57. <https://doi.org/10.1017/S0020818305050198>.
- WILNER, ALEX S., AND ANDREAS WENGER. 2021. *Deterrence by Denial: Theory and Practice*. Amherst, NY: Cambria Press.
- ZAPFE, MARTIN. 2017. “Deterrence from the Ground Up: Understanding NATO’s Enhanced Forward Presence.” *Survival* 59(3): 147–60. <https://doi.org/10.1080/00396338.2017.1325604>.