

Can civilian attitudes predict insurgent violence? Ideology and insurgent tactical choice in civil war

Journal of Peace Research
2017, Vol. 54(1) 47–63
© The Author(s) 2016
Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/0022343316675909
journals.sagepub.com/home/jpr



Kentaro Hirose

Waseda Institute for Advanced Study, Waseda University

Kosuke Imai

Department of Politics and Center for Statistics and Machine Learning, Princeton University

Jason Lyall

Department of Political Science, Yale University

Abstract

Are civilian attitudes a useful predictor of patterns of violence in civil wars? A prominent debate has emerged among scholars and practitioners about the importance of winning civilian ‘hearts and minds’ for influencing their wartime behavior. We argue that such efforts may have a dark side: insurgents can use pro-counterinsurgent attitudes as cues to select their targets and tactics. We conduct an original survey experiment in 204 Afghan villages and establish a positive association between pro-International Security Assistance Force attitudes and future Taliban attacks. We extend our analysis to 14,606 non-surveyed villages and demonstrate that our measure of civilian attitudes improves out-of-sample predictive performance by 20–30% over a standard forecasting model. The results are especially strong for Taliban attacks with improvised explosive devices. These improvements in predictive power remain even after adjusting for possible confounders, including past violence, military bases, and development aid.

Keywords

civil war, public opinion, out-of-sample prediction, survey experiment

Introduction

Can civilian attitudes help predict insurgent attacks? The past decade has witnessed the renewal of a debate over the importance of winning ‘hearts and minds’ in counterinsurgency wars such as Afghanistan and Iraq. Billions of dollars have been spent by militaries and development agencies on reducing insurgent violence by finding the mix of aid, services, and protection that persuades fence-sitting civilians to side with the government (World Bank, 2012; Sambanis, Schulhofer-Wohl & Shayo, 2012; Beath, Christia & Enikolopov, 2011; Berman, Shapiro & Felter, 2011; Department of the Army, 2007).

This discussion rests on a simple premise: civilian attitudes are a reliable guide to subsequent behavior.

Most of the theoretical literature on civil war violence, however, remains deeply skeptical of a link between attitudes and behavior (see, for example, Kalyvas, 2006, 2012; Stoll, 1993; Leites & Wolf, 1970). Civilians are often cast as strategic actors, shifting allegiances frequently as circumstance requires. In this view, attitudes are likely endogenous to a host of wartime dynamics, including the relative distribution of control and economic assistance, and thus hold little independent power when explaining behavior. In addition, given the difficulties of accurately measuring wartime attitudes, it is

Corresponding author:

jason.lyall@yale.edu

unsurprising that the role of attitudes in shaping patterns of violence has been relatively unexplored.

Given this skepticism, we take up the challenge of demonstrating that civilian attitudes can be used to predict future patterns of insurgent violence. We argue that the distribution of attitudes toward the counterinsurgent act as 'cues' for insurgents that facilitate their decision-making about where and how to stage attacks. The spatial distribution of attitudes helps insurgents prioritize their attacks against counterinsurgent forces given resource constraints while dictating how discriminate their violence must be given the local population's prevailing views. Efforts to win 'hearts and minds' may therefore have an unintended consequence: these efforts can attract increased insurgent attacks in areas where counterinsurgents have made the deepest inroads.

Our interest in civilian attitudes as a predictor of insurgent violence joins a renewed call for prediction in the social sciences (e.g. Hill & Jones, 2014; Schrodt, 2014; Metternich et al., 2013; Montgomery, Hollenbach & Ward, 2012; Braithwaite & Johnson, 2012; Goldstone et al., 2010; Weidmann & Ward, 2010; Bohorquez et al., 2009; King & Zeng, 2001; Beck, King & Zeng, 2000). While existing prediction efforts are typically cross-national, we take a disaggregated approach to data collection and predict village-level violence over variable spatial and temporal windows as fine-grained as one kilometer and one day, respectively.

We also move beyond existing forecasting models by introducing contextual information rather than relying simply on prior violence to predict future attacks, as is often current practice (Johnson et al., 2011; Zammit-Mangion et al., 2012; Yonamine, 2013). Specifically, we draw on a survey experiment in 204 villages in Afghanistan, along with two datasets recording insurgent attacks against the International Security Assistance Force (ISAF) and civilians, to test the association between attitudes and subsequent violence. Wary of the dangers of over-fitting, we also extend our analysis to out-of-sample prediction (Ward, Greenhill & Bakke, 2010; King & Zeng, 2001; Beck, King & Zeng, 2000) using 14,606 non-surveyed villages to examine how incorporating attitudes improves predictive accuracy across multiple categories of insurgent targets and tactics.

Three main findings emerge. First, we find that pro-counterinsurgent attitudes significantly improve the accuracy of predicting the location of insurgent direct attacks and the use of improvised explosive devices (IEDs) for up to 10 months after our survey. Specifically, our attitudinal measure has a robust in-sample association with future violence patterns while improving the

predictive performance of our out-of-sample models by 20–30%. Second, we find little evidence that pro-counterinsurgent attitudes are associated with 'found' IEDs, suggesting that winning hearts and minds may not translate into actionable intelligence. Finally, these findings hold after adjusting for confounding variables such as prior insurgent violence, the location of ISAF and Afghan National Security Forces (ANSF) bases, and development aid.

Ideology and violence in civil war settings

We argue that ideological considerations play a prominent role in guiding insurgent decisionmaking. This contrasts with existing theories that remain deeply skeptical of ideological arguments (on this skepticism, see Drake, 1998; Gutiérrez-Sanin & Wood, 2014). Many core works sideline ideology (including Leites & Wolf, 1970; Kalyvas, 2006, 2012; Christia, 2012; Staniland, 2014) or only invoke it to explain residual variation (Weinstein, 2007; Mampilly, 2011: 77–78). Most theoretical accounts view ideology as either a post hoc rationalization for insurgent actions or as easily sacrificed on the altar of military expediency. It is unsurprising, then, that one sweeping review of existing theories of civil war violence concludes that 'how and why ideology matters is not clear' (Gutiérrez-Sanin & Wood, 2014: 217).

We contend that the location, type, and lethality of insurgent violence are all shaped by the underlying spatial distribution of civilians' relative support for combatants. While an emerging literature examines variation in ideology across and within organizations to explain violence (Gutiérrez-Sanin & Wood, 2014), we focus instead on how insurgent tactical choice is conditioned by the spatial distribution of ideological support for combatants among local populations. Insurgents target locations that exhibit relatively higher support for the counterinsurgent, maximizing the punishment and deterrent effects of their violence. Civilian attitudes toward the combatants are stable enough to drive targeting and to exhibit independent effects on the prediction of future attacks that are not captured by structural or situational factors.

'Hearts and minds' and insurgent violence

A staple of counterinsurgency theory now holds that winning over civilian 'hearts and minds' is the key to defeating insurgents. This approach counsels the adoption of a mix of economic assistance, service delivery, and protection to convince fence-sitting populations to support the counterinsurgent. Civilian behavior follows

attitudes; win over hearts and minds, and civilians will provide tips about insurgents hiding among them. Individuals in this account are calculating actors, often supporting whichever side promises the most benefit while disregarding prior ethnic or political loyalties. In brief, the greater the aid distribution, the more information gleaned from the population, and the sharper the corresponding decrease in attacks by insurgent organizations.

These claims generate several observable implications. Most centrally, if these claims are correct, we should observe a *decrease* in insurgent attacks as relative support for the counterinsurgent grows. The US Army's Field Manual (May 2014) makes this connection explicit: military offensives and aid campaigns are designed to create safe spaces for the population by reducing insurgent attacks. Each phase of the Army's 'shape-clear-hold-build-transition' framework is associated with decreased insurgent attacks, though not necessarily their complete absence. Indeed, the first metric proposed for measuring hearts and minds effectiveness is the reduction of insurgent attacks and casualties among COIN forces and civilians (Department of the Army, 2014: Section 12-29; see also Figure 9-1).

Scholars remain divided, however, over the central mechanism driving this result. Insurgent violence may decrease, for example, as the provision of tips from civilians better enables counterinsurgents to identify and destroy rebel leaders and networks (Kalyvas, 2006; Department of the Army, 2007; Berman, Shapiro & Felter, 2011). Cash-for-work programs and other forms of economic assistance may also raise the opportunity costs for participating in the insurgency, making recruitment difficult and driving would-be insurgents from the ranks (Blattman & Annan, 2016). Increased contact with counterinsurgents may also convince wary civilians to change their beliefs and to support the counterinsurgents' cause. Finally, an increase in counterinsurgent troop strength in an area may act as a deterrent, forcing rebels to seek targeting opportunities elsewhere (Department of the Army, 2007: 1–13).

Insights drawn from hearts and minds theorizing have made impressive inroads among academics and policy-makers. Yet there are reasons for skepticism. Most existing studies consist of a single-shot impact evaluation with simple pre- and post-aid distribution differences in insurgent attacks. Civilian attitudes are typically not measured; increased support is inferred from reduced insurgent violence (for an important exception, see Beath, Christia & Enikolopov, 2011). Mechanisms, too, remain untested; measures for the quantity and quality of tips from locals are usually unavailable or incomplete,

leaving scholars unable to stitch together the causal sequence of aid delivery, changed attitudes, increased information, and then reduced insurgent violence.¹

Civilian attitudes as cues for insurgent targeting

We contend that insurgent targeting in civil war is a reflection of the underlying spatial distribution of relative support for the combatants. In this view, the insurgent group's ideology shapes not only the end goal of the insurgency but also the nature of the targets (the 'who'), the location of these attacks (the 'where'), and the nature of the tactics and weapons used (the 'how'). Following Gutiérrez-Sanin & Wood (2014: 214), we define *ideology* as 'a set of more or less systematic ideas that identify a constituency, the challenges the group confronts, the objectives to pursue on behalf of that group, and a (perhaps vague) program of action'. In particular, ideology informs tactical choice by establishing the set of appropriate means and targets for achieving desired political goals. *Tactics* are defined here as the specific techniques that insurgent organizations use to destroy, degrade, or displace their opponents to secure operational objectives (see Drake, 1998).

We anticipate that populations suspected of disloyalty to the insurgent cause are likely to absorb the brunt of insurgent attacks. More specifically, we focus on the political allegiance of local populations, where support is defined as the relative leanings of a given populace toward the insurgency and the counterinsurgent. Contrary to the expectations of 'hearts and minds' theorizing, we expect to observe a *positive* association between relatively higher support for the counterinsurgent and the level of subsequent insurgent attacks. Insurgents may of course be opportunistic in their execution of attacks, but they do so within a bounded set of choices and means delineated by their ideological precepts.

Insurgents' knowledge of the distribution of civilian attitudes will influence their tactical choice, especially the mix between discriminate and indiscriminate weaponry. Tactics that require networks within villages – notably, the use of improvised explosive devices (IEDs) and suicide bombings – become wasting assets if locals begin to side with the counterinsurgent. Insurgents will feel compelled by a 'use it or lose it' dynamic to privilege these tactics in an effort to roll back pro-counterinsurgent support. These tactics are typically less discriminate than other forms of violence, raising the possibility of civilian

¹ For a skeptical take on wartime informing, see Lyall, Shiraito & Imai (2015).

casualties. Roadside IEDs often miss their intended military targets and instead inflict casualties on locals who had the misfortune of trailing a military convoy, for example. Yet insurgents will relax their prohibitions (if any) on killing civilians if the violence is largely restricted to the pro-counterinsurgent village. In this case, indiscriminate violence underscores the risks associated with abandoning the insurgency. Similarly, civilians who actively side with the counterinsurgent ('traitors') will be viewed as legitimate targets by insurgents and their own supporters, mitigating the damage to the insurgent organization's reputation commonly associated with killing innocents.

These tactical choices are driven by the insurgents' overall end-goal and are designed to have several effects on counterinsurgent forces and local populations. First, insurgents are seeking to punish, not persuade, pro-counterinsurgent villages. In these situations, they are less concerned about negative effects of killing 'traitors'. Civilians in swing villages may actually support these actions if they share the view that pro-counterinsurgent individuals or villages have betrayed a particular political cause. Insurgents may even gain support via such actions: if the counterinsurgent is provoked into retaliatory actions that harm civilians, the net effect is likely a shift in relative support *toward* the insurgent organization. Counterinsurgents are disproportionately punished by populations for civilian victimization if they are drawn from groups outside the insurgent group's own members (Lyll, Blair & Imai, 2013).

Second, insurgents attack pro-counterinsurgent villages to influence fence-sitting villages. Punishing pro-counterinsurgent villages can forestall the loss of other villages by demonstrating the costs of siding openly with the counterinsurgent. Insurgent targeting against these villages has a twin punish-deter logic; punishing these pro-counterinsurgent villages will have a deterrent effect on other, perhaps wavering, 'swing' villages. While insurgents are certainly interested in swing villages, they are less likely to use violence here, hoping to avoid the absolute loss of support that can accompany rebel victimization of civilian populations (Lyll, Blair & Imai, 2013). Insurgents are more likely to use redistributive mechanisms ('carrots') in these swing villages; pro-counterinsurgent villages, on the other hand, are more likely to be punished violently ('sticks').²

Third, targeting pro-counterinsurgent villages inflicts costs on the counterinsurgent publicly, helping to stem the erosion of insurgent support that can mean the loss of information, resources, and recruits needed to generate the insurgency's combat power. Loss of these villages may also drive a wedge into the insurgency by creating a 'loyalist' faction that undercuts insurgent recruitment along ideological or ethnic lines (see Kalyvas, 2008; Lyll, 2010). As a result, striking pro-counterinsurgent locations preserves the insurgency's momentum while increasing the counterinsurgent's cost of defending these centers. If the insurgents are able to create the perception that the counterinsurgent cannot credibly defend the populace, then erstwhile counterinsurgent supporters may curtail their collaboration or defect (back) to the rebels.

This argument rests on two assumptions, namely, that civilian preferences toward the combatants are stable enough to be meaningful and that insurgents devote resources to monitoring these attitudes. We believe these are reasonable assumptions. In particular, the assumption that insurgents possess informational advantages relative to counterinsurgents is shared by nearly all theories of civil war violence (Lyll & Wilson, 2009). It is well recognized that counterinsurgents struggle to solve the 'identification problem' – namely, correctly identifying insurgents hiding among civilians (Kalyvas, 2006: 89). We flip the identification problem, however, by exploring how insurgent tactical choice is guided by the spatial distribution of combatant support.

Our theory produces several empirical expectations. Our core claim is that insurgents will disproportionately target locations that express relatively higher pro-counterinsurgent sympathies. Both civilians and military forces should be targeted in these locations at rates higher than in similar, but less pro-counterinsurgent, areas. We expect that these attacks are quite localized (i.e. at the village level). That is, insurgent violence should cluster around a pro-counterinsurgent village; the predictive value of civilian attitudes should diminish as we move away from these villages. Insurgents should also emphasize indiscriminate tactics when attacking these villages given their ability to maximize damage among counterinsurgent forces, their shock value among targeted civilians, and their deterrent value among neighboring villages. Pro-counterinsurgent attitudes should therefore be an important predictor of the use of indiscriminate weapons such as improvised explosive devices. Finally, attitudes should remain important for explaining future attacks even when controlling for alternative explanations such as base location, the distribution of control and aid, and prior patterns of violence.

² This claim is also broadly consistent with Kalyvas's (2006) argument that populations under contested insurgent-counterinsurgent control are least likely to witness violence.

Empirical strategy

We use a survey experiment conducted in Afghanistan in 2011 to measure attitudes for our predictive exercise. We first briefly outline our survey design, including the multi-stage sampling strategy. We then detail the four indirect endorsement experiments that, when combined with a statistical model, provide a village-level measure of relative support for ISAF. Next, we introduce the declassified data used to construct measures for three types of insurgent violence. We also introduce variables to operationalize alternative predictors for insurgent violence, including: (1) prior insurgent attacks; (2) the location of ANSF and ISAF bases; and (3) the location of counterinsurgents' aid projects. Taken together, these three covariates constitute our 'base' model for estimating the location and timing of insurgent attacks over various spatial and temporal intervals. We provide qualitative evidence of the link between civilian attitudes and Taliban tactical choice and take up the question of external generalizability in the Online appendix.

The survey

We measure support for ISAF and the Taliban using a survey of 2,754 adult male respondents from 204 villages in 21 districts of five Pashtun-dominated provinces of Afghanistan (Logar, Kunar, Uruzgan, Helmand, and Khost).³ After two large-scale pilot surveys in these provinces, the survey was conducted between 18 January and 3 February 2011. These pilots were conducted from 25 September to 5 October 2010 and 22 November to 5 December 2010. They were used to test sensitivity to question order, potential design effects (e.g. ceiling and floor effects for the experiment modules), and the nature of logistical challenges.

The sample was constructed using a multistage sampling method. Figure 6 in the Online appendix illustrates the location of the 204 surveyed villages (blue dots) and the 14,606 non-surveyed villages (red dots). First, the five provinces, whose borders are highlighted in the map, were randomly sampled from the 13 Pashtun-majority provinces (colored gray on the map).⁴ Within each of these five provinces, districts, and then villages, households, and finally individuals, were randomly sampled. Households were chosen using the random walk method. Owing to security considerations and

cultural constraints, only male respondents aged 16 years and older were randomly selected from the sampled households using a Kish grid.

We obtained an 89% participation rate (2,754 respondents out of 3,094 approached individuals). The average respondent was a 32-year old Pashtun male who was likely married (77%), possibly employed full-time (58%), and possessed little or no formal (government) schooling and only an average of 18 months of madrassa education. Nearly all respondents were Pashtun by ethnicity (93%). Our sampling frame encompassed some of the poorest areas in Afghanistan; respondents reported possessing only 90 minutes of electricity per day, and daily income hovered between US\$1 and US\$6.

Of the original 204 villages, only four proved inaccessible due to a combination of Taliban hostility, the presence of criminal elements and, in two cases, the inability of enumerators to find the village. In all cases, village elders were first approached by our survey firm, Opinion Research Center of Afghanistan (ORCA), to describe the survey and to obtain permission for enumerators to enter the village. All enumerators were locals; nearly all surveys were conducted in Pashto. Informed consent was obtained in all cases orally; special permission was granted by Yale's IRB (Institutional Review Board) to waive written consent requirements to avoid linking respondents to particular surveys. This helped minimize the risk of reprisal if the written forms were intercepted by the Taliban. Similarly, the names of respondents were not recorded.⁵

Measuring support

We employ a battery of four endorsement experiments to mitigate the sensitivity of measuring wartime support for ISAF and the Taliban. This indirect approach reduces social desirability bias and item non-response when asking questions about sensitive issues (Bullock, Imai & Shapiro, 2011). Direct questions, by contrast, can endanger enumerators and respondents alike and often result in high non-response rates and biased answers.

To take one example, a recent wave of ISAF's own Afghan National Quarterly Assessment Report (ANQAR) in November–December 2011 recorded nearly 50% non-response rate as potential respondents refused to participate. Our overall refusal rate was about 5% for our endorsement experiments, a difference we ascribe to the indirect questioning method we employed.

³ These 21 districts represent 39% of the total districts in these provinces (21/54).

⁴ The remaining eight provinces are Ghazni, Kandahar, Laghman, Nangahar, Paktia, Paktika, Wardak, and Zabul.

⁵ Further details about survey implementation can be found in the Online appendix and Lyall, Blair & Imai (2013).

Table III in the Online appendix shows the pattern of non-response rates by province. The highest non-response rate is obtained in Helmand province, which has the highest level of insurgent violence and support. We address this non-response problem in our statistical analysis by assuming that the pattern of non-response is random conditional on all observed data.

The mechanics of an endorsement experiment are straightforward. We first randomly divided a sample of respondents into groups. In the 'control' group, respondents were asked to rate the level of their support for a particular policy. For those in the 'treatment' group, the identical question was asked except that the policy was said to be endorsed by an actor of interest. We then take advantage of subtle cues induced by endorsements and interpret the difference in responses between the treatment and control groups as evidence of support (or lack thereof) for this actor of interest. This is based on the psychological literature, which finds that people tend to more positively (negatively) evaluate an item when paired with another item they like (dislike). In our application, we have two actors of interest, ISAF and the Taliban, and thus the sample was randomly divided into three groups of equal size – Taliban treatment, ISAF treatment, and control – across individual respondents within each sampled village.

Typically, multiple policies in the same domain are selected so that the measurement does not rely on a single instrument. Statistical power is also increased by analyzing multiple items together. In our survey experiment, we ask questions concerning four domestic policies: prison reform, direct election of district councils, a reform of the Independent Election Committee, and the strengthening of anti-corruption policies. The exact question wording appears in the Online appendix.

A Bayesian hierarchical factor analytic model is used to (partially) pool the responses to these four questions together, creating an estimate of individual-level support for ISAF and the Taliban (Bullock, Imai & Shapiro, 2011). We then aggregate individual-level support values by computing the sample mean to create an estimate of village-level support for each combatant. When modeling individual-level support, we use as regressors village-level random effects and individual-level covariates such as the respondent's age, income level, years of education, years of madrassa schooling, direct exposure to violence by the Taliban or ISAF, experience of encountering the Taliban or ISAF, the frequency in which the respondent encounters ISAF, and whether his tribe was pro-Taliban. The model we use is described in the Online appendix and is fitted through the open-source R package *endorse* (Shiraito & Imai, 2012).

The resulting measures of village-level support are numerical estimates for each combatant. The support level for the Taliban ranges from -1.14 to 1.42 while that for ISAF ranges from -1.75 to 0.43 with positive (negative) values, indicating support for (opposition against) the combatant. Figure 7 in the Online appendix presents the distribution of each support measure as well as the distribution of the difference between the two. The figure clearly shows that a majority of the Afghan population supports the Taliban over ISAF. These support measures are standardized on a latent variable scale (so-called 'ideal points' and 'ability' parameters in the political methodology and psychometrics literatures, respectively), and so only their sign and relative magnitude can be interpreted. In addition, we validate these support measures against another measure based on the item count technique and find these two indirect questioning methods provide essentially identical findings (Blair, Imai & Lyall, 2014).

In our main analysis, we operationalize relative support for ISAF as the difference between ISAF and Taliban support levels and use this measure as the key predictor of insurgent violence. Relative support for ISAF ranges from -3.01 to 0.91 , suggesting that Afghans are far more supportive of the Taliban than ISAF. Figure 1 presents the spatial distribution of relative ISAF support measure for the 204 surveyed villages in the five sampled provinces. In general, Taliban support is strongest in Uruzgan and especially Helmand, two provinces long associated with the post-2001 re-emergence of the Taliban. The eastern provinces of Khost, Kunar, and Logar present a more mottled picture, with substantial pockets of support for the Taliban and ISAF.

Of particular interest is the fact that neighboring villages sometimes have different relative levels of support for these combatants. Even within the same district, some villages are supportive of ISAF while others are opposed. This local variation underscores the need to embrace disaggregated data; ISAF, by contrast, continues to track support at much more aggregate district or even provincial levels.

Measuring violence

We first measure insurgent violence using declassified event data from ISAF's Combined Information Data Network Exchange (CIDNE). These data record the date, location (using the Military Grid Coordinate System), and nature of insurgent attacks against ISAF forces and installations throughout Afghanistan. Distinct from WikiLeaks' Afghan War Diary, these data represent the

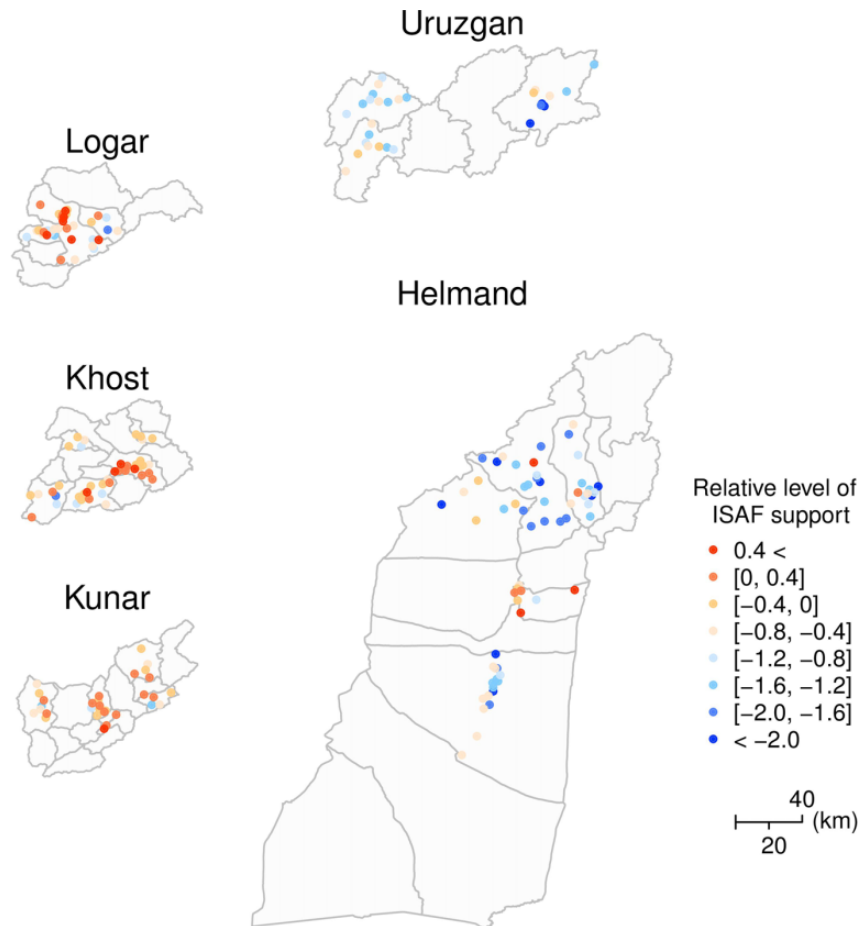


Figure 1. Spatial distribution of relative support for International Security Assistance Force (ISAF) in the 204 surveyed villages from five randomly sampled Pashtun-dominated provinces

The warmer color (red) represents villages which are more supportive of ISAF than the Taliban; the colder color (blue) indicates less supportive villages.

most comprehensive account of insurgent attacks to date, though they are not without their limitations. These ‘Significant acts’ (SIGACTs) rarely cover violence against Afghan National Security Forces and exclude violence against civilians. As such, our main analysis focuses on insurgent attacks against ISAF alone. We use data from 10 months before and after our January–February 2011 survey for our prediction models; a total of 69,841 insurgent attacks were recorded over this period.

CIDNE tracks 11 discrete types of insurgent attacks that are relevant for our purposes here. We aggregate these types into three broad categories. First, we constructed an ‘Improvised explosive device (IED) attacks’ category that records 11,577 events of IED explosions, mine strikes, and premature IED detonations. IEDs represent the most lethal form of insurgent attack against ISAF forces, accounting for 54% of all soldier fatalities since 2007 (iCasualties.org, 2013), and ISAF has

devoted billions of dollars in a cat-and-mouse effort to mitigate this threat.⁶

Second, we created an ‘IED found’ category that includes 19,093 events where ISAF forces found and cleared IEDs or mines.⁷ We make a clear distinction between ‘found’ IEDs and ‘detonated’ IEDs to determine whether tips from locals in pro-ISAF villages translate into higher than average timely discovery of IEDs, as anticipated by hearts and minds theory.⁸ This is an imperfect measure: found IEDs may reflect improved

⁶ Specific CIDNE categories are (1) IED Explosion, (2) Premature IED Detonation, and (3) Mine Strike.

⁷ Specific CIDNE categories are (1) IED Found and Cleared, (2) IED Threat, (3) IED Cache/Find, and (4) Mine Found and Cleared.

⁸ Without disaggregating the IED category, we would also be left unable to determine whether a relative increase in IED counts in pro-ISAF villages was due to greater insurgent targeting (‘detonated

detection capabilities via technical means by ISAF that occurred independent of civilian support. Absent classified data on IED detection, we cautiously use this variable as a proxy for tips flowing from civilians.

Lastly, we created a 'Non-IED attacks' category that includes 39,171 insurgent attacks using small arms fire, indirect fire (e.g. mortars), and rocket fire against ISAF forces and installations.⁹ We distinguish between IED and non-IED attacks because, as explained earlier, the notion of attitudes-as-targeting-cues suggests an especially strong association between civilian attitudes and certain classes of tactics with maximum destructive value such as IEDs.

In our analysis, these three categories are operationalized as count variables recording the number of relevant events in specified temporal windows before and after the survey in each village. We also test across different spatial radii around villages. We therefore aim to predict the aggregate number of attacks of each category within defined spatial and temporal windows around sampled and then non-sampled villages. Table I in the Online appendix illustrates the pre- and post-survey distribution of insurgent violence across five provinces. The locations of ANSF and ISAF bases and aid projects are also noted. Among the randomly sampled five provinces, most insurgent violence (57% ~ 80%) is centered on Helmand, while most ISAF bases (31%) and aid projects (48%) are clustered in Kunar.

Other predictors

The core test of our argument lies in observing improvements in predictive accuracy while adjusting for alternative explanations for insurgent violence.

First, the conflict modeling literature has demonstrated that a leading predictor of future violence is simply the prior distribution of violence (e.g. Yonamine, 2013; Zammit-Mangion et al., 2012; Montgomery, Hollenbach & Ward, 2012; Bohorquez et al., 2009). To account for this explanation, we draw on the same declassified CIDNE data as used for our dependent variable and simply mirror the spatial and temporal windows on either side of the survey date for a given village. For example, if we examine the effect of village-level support on the number of IED attacks occurring within a x km radius around each village during y months post-survey,

we also include the number of past IED attacks that have occurred within the same spatial and temporal windows.

Second, insurgent targeting may reflect the distribution of counterinsurgent bases and forces within a given area; that is, after all, where the majority of counterinsurgent targets are found. We therefore use geo-referenced base location data to adjust for the number of ANSF and ISAF military installations present in the pre-survey period (June 2007–December 2010) within a three-kilometer radius of each village. The choice of three kilometers is somewhat arbitrary. As a result, in our analysis we use varying radii to examine the robustness of our findings.

We lack a direct measure of ISAF and ANSF force size or patrol rate; these data are either classified or do not exist. That said, we believe our data on base location do a credible job in controlling for ISAF presence. Many of these bases are small combat outposts (COPs) that patrolling forces sortie to and from; a five-kilometer radius is likely to capture all but the most distant (and infrequent) patrol. The Taliban have also expressed a preference for ambushing tired counterinsurgent forces as they arrive back at their bases (Meyerle & Malkasian, 2009: 6). Far from viewing these military bases as oases of calm, they have historically been the site of concerted Taliban efforts to harass or breach their defenses, often via suicide bombing.

Figure 8 in the Online appendix presents the bivariate relationship between ISAF support and the number of ANSF/ISAF bases. We observe that regardless of the choice of radii there is a small degree of positive relationship between them; villages which have a larger number of bases tend to be more supportive of ISAF.¹⁰

Third, pro-counterinsurgent attitudes may be associated with the presence of aid programs designed to influence fence-sitting populations. Insurgents may in turn orchestrate attacks to derail these aid projects. This dynamic was observed in the Philippines, for example, in response to the government's KALAHI-CIDSS Community Driven Development (CDD) anti-poverty program. Insurgents timed their attacks to the announcement and initial phase of aid programming in CDD-villages to prevent the erosion of their popular support (Croft, Felter & Johnston, 2014). To account for this possibility, we adjust for the number of USAID Community Development Programs (CDP) using

IEDs'), increased discovery thanks to locals, or to both mechanisms working simultaneously.

⁹ Specific CIDNE categories are (1) Direct Fire, (2) Attack, (3) Raid/Ambush, and (4) Indirect Fire.

¹⁰ Our findings are robust to the inclusion of an alternative measure of control provided by ISAF's own district-level assessments. See Figure 18 in the Online appendix.

geo-referenced data for all projects initiated within one kilometer of each village. The large-scale, \$250 million CDP program was designed to foster stability using cash-for-work initiatives in areas where insurgents had been forcibly evicted through military operations during 2010–12. In Figure 9 in the Online appendix, we present the bivariate association between the aid projects (with varying radii from villages) and village level support for ISAF. While there appears to be quite weak correlation between the two variables, we nonetheless include this variable in our base model.

Empirical results

We begin our empirical analysis by establishing a robust association between civilian attitudes and future insurgent violence against ISAF based on our original sample of 204 villages. We then assess the improvement gained by introducing civilian attitudes to our predictive models of insurgent violence by extending the analysis to nearly 15,000 out-of-sample villages. The latter analysis cross-validates the predictive power of our civilian attitudes measure.

Using this same two-step process, we then examine the association between civilian attitudes and insurgent violence against civilians. For space reasons, we report these findings in the Online appendix.

In-sample prediction performance

Throughout our in-sample analysis, we use a linear regression model to quantify the association between relative support for ISAF and subsequent insurgent violence *after* the survey. Note that each village has a slightly different start and end date given their survey order. This results in different temporal windows (of the same size) across villages. Specifically, we regress the number of future insurgent attacks within certain temporal and spatial windows on its relative support level for ISAF as well as the other possible confounders discussed in the Online appendix. We vary the length of time and distance windows in order to examine the robustness of our findings. We also explored the possibility of nonlinearity but we find a simple linear model captured most of the systematic variation.

Findings. As an illustrative example, we present our model's results using a temporal window of five months pre- and post-survey. We use a 15 km radius around each sampled village to calculate the number of insurgent attacks. Figure 2 demonstrates that a strong positive association exists between relative support for ISAF and

future insurgent attacks even after accounting for prior violence, counterinsurgent bases, and CDP aid projects. This is particularly true for IED attacks (leftmost plot), where the coefficient is most precisely estimated. A village that has modest relative support for ISAF (equivalent to a 0.5 value) is predicted to have an additional 13 IED attacks on average over the next five months (with a 95% confidence interval of [7, 20] using a heteroskedasticity-consistent standard error) when compared to a village strongly opposed to ISAF (equivalent to a -2.5 value).¹¹

We observe a similar pattern in the association between relative support and non-IED attacks (the rightmost plot). A village with modest relative support is expected to have 34 more non-IED attacks on average than a village with strong opposition, although the 95% confidence interval [3, 64] is quite large. The number of 'found' IEDs in the middle plot does not follow this pattern, however. Shifting a village's support from modestly pro-ISAF to strongly opposed actually yields an average *increase* of eight found IEDs, though the estimate is not statistically significant different from zero. More generally, the association between relative support and found IEDs is relatively weak, raising doubt about the linkage between pro-counterinsurgent attitudes and the provision of tips. To be sure, not all IEDs are discovered through tips, and we cannot rule out the possibility that local ANSF and ISAF vary in their ability to detect such devices without local assistance. Yet the weakness of this finding suggests that pro-counterinsurgent attitudes may not necessarily translate into positive (from the counterinsurgent's view) local action.

The positive association between relative support and future IED attacks is robust to the choice of temporal and spatial windows. We repeat the analysis by varying the temporal window from one to ten months and changing the radius around the surveyed village from 1 km to 60 km. Figure 3 presents contour plots of the t -statistics for the estimated coefficient of the relative ISAF support measure where the dark blue areas represent large values of t -statistics. Overall, we continue to observe a positive and statistically significant association between ISAF support and insurgent IED attacks while

¹¹ Figure 10 in the Online appendix presents the correlations without adjustment. While a positive association exists between these measures, the correlation is weak, especially for IED attacks and IEDs found. We believe, however, that adjusting for the past level of violence is essential when assessing the importance of support measures in predicting future violence.




Figure 2. Positive association between the number of future insurgent attacks and relative ISAF support

The plots illustrate the statistically significant association between the number of insurgent attacks that have occurred within 15 km of each village during the five months *after* the survey in each village (vertical axis) and its relative level of ISAF support (horizontal axis) while adjusting for the number of insurgent attacks that have occurred (again within 15 km around each village) five months *prior* to our survey, the number of counterinsurgent bases within a 3 km radius of each village, and the number of aid programs within a 1 km radius of each village. The results are based on the linear regression model estimated separately for each of the three violence categories where the number of future insurgent attacks is regressed on the relative level of ISAF support and the other covariates. The dashed lines represent 95% confidence intervals based on robust standard errors.

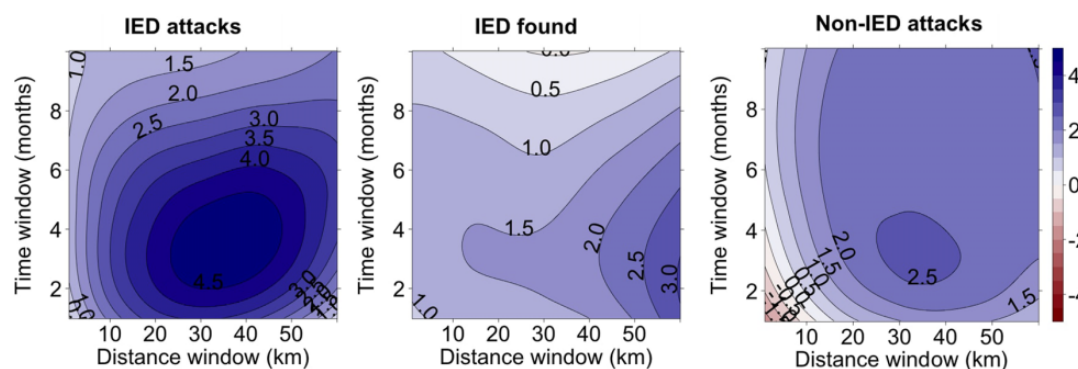


Figure 3. t -statistic of the estimated coefficient for the relative level of ISAF support across a wide range of temporal and spatial window sizes

The dark blue areas represent large values of t -statistics. The estimated coefficient corresponds to its marginal effect on the number of future insurgent attacks while adjusting for prior insurgent violence and the number of ANSF/ISAF bases and USAID aid projects. The linear regression models which produced the results displayed in Figure 2 are repeatedly estimated using broad time and distance windows (from 1 to 10 months before/after the survey and from 1 to 60 km of each village's perimeter). The results illustrate the robustness of the positive association between the number of future insurgent attacks (especially with IEDs) and relative ISAF support.

controlling for prior insurgent attacks, counterinsurgent bases, and CDP projects.

As Figure 3 illustrates, we observe important variation in the model's ability to predict insurgent violence. The positive relationship between pro-ISAF sentiment and insurgent IED attacks is strongest at the four-month interval and about 40 km around the surveyed village as indicated by the darkest blue area. Similarly, the relatively large t -statistics for non-IED attacks can be seen over similar temporal and spatial windows. Comparing t -statistics across attack types, it is apparent that our model is especially well-suited to predicting IED

emplacement, with peak t -statistics approaching 4.5 compared to a still sizable 2.5 for non-IED attacks.¹² In addition, our results also hold if we analyze the absolute (rather than relative) level of support for each group separately. We find that insurgent violence is positively associated with the absolute measure of ISAF support,

¹² For IED and non-IED attacks, the variation in t -statistics mainly stems from the variation in the effect sizes rather than the standard errors. By contrast, the variation in t -statistics for IED found comes from both the effect sizes and the standard errors.

while it is negatively correlated with the absolute levels of Taliban support (see Figures 12 and 13 in the Online appendix).

The positive association between relative ISAF support and insurgent IED attacks is robust to various modeling assumptions. In Figure 14 of the Online appendix, we present additional results based on matching to enable flexible covariate adjustment rather than simple linear adjustment (Ho et al., 2007). Using Mahalanobis metric matching, villages are first paired according to prior insurgent attacks, counterinsurgent bases, and the number of CDP projects. We then regress the pairwise difference in future violence on the pairwise differences in relative ISAF support, prior violence, counterinsurgent bases, and CDP projects. The association between ISAF support and IED attacks remains strong even using this non-parametric analysis. By contrast, the results for the other two types of attacks do not appear to hold for this matching analysis, indicating the lack of robustness for IED found and non-IED attacks.¹³

We also conducted robustness checks with various distance-to-base windows. In Figure 15 of the Online appendix, we conduct the same analysis as above but using a set of different distance windows (1 km, 5 km, and 10 km) when counting the number of bases around each village. Similarly, we examine the robustness of our finding to the choice of distance window for the presence of aid programs. In Figure 16 of the Online appendix, we use alternative distance windows (3 km, 5 km, and 10 km) and repeat the same analysis. All of these robustness checks support our conclusion that a strong positive association exists between pro-ISAF sentiment and future insurgent IED attacks.

Finally, we show that our results are robust to the inclusion of political ideology of respondents aggregated at the village level. The measure is obtained directly from the estimated IRT endorsement experiment model and is based on four policy reform questions. A greater value of estimated ideal point implies more support for policy reform in general. In Figure 19 of the Online appendix, we observe that there is a large variation in this measure of ideology across villages. In Figure 20 of the Online appendix, we show that our main results are essentially unchanged even after adjusting for this measure of ideology.

¹³ We also explored the non-parametric estimation of the effects of relative ISAF support on insurgent attacks using natural cubic splines (the results not shown). However, the cross-validation indicated that the dataset is too small to reliably estimate such a non-parametric model, favoring a simple linear model.

Out-of-sample prediction performance

To further examine the predictive power of civilian attitudes, we evaluate their out-of-sample predictive performance using 14,606 non-surveyed villages. Doing so reduces the possibility of false discovery by ensuring that the results of our in-sample results are not due to overfitting to a particular sample of surveyed villages.¹⁴ We first predict relative ISAF support for these out-of-sample villages using village-level covariates. We then forecast the number of insurgent attacks with these predicted support levels and then compare our forecasts with actual insurgent attacks. As before, our aim is to examine whether these predicted support levels improve forecasting performance of future insurgent violence even while controlling for prior insurgent attacks and the presence of counterinsurgent bases and aid projects.

The out-of-sample prediction proceeds in two steps. First, using the 204 surveyed villages, we estimate the *ISAF support model* by regressing the relative ISAF support level and a set of available village- and district-level covariates. These covariates include village population size and elevation as well as several district-level factors, including ISAF's own measure of its relative control in that district, the existence of Taliban-run sharia courts (as a measure of Taliban control), and whether the district bordered Pakistan (to control for differences arising from cross-border spillover). Using the same set of covariates, we can now estimate ISAF support level for each of the out-of-sample villages based on the fitted model. We rescale the relative support estimate for out-of-sample villages so that their standard deviation is identical to that of the original village sample.

Second, we estimate the *Insurgent violence model*, which is the same linear regression model as used in the Online appendix. Together with the estimated ISAF support from the previous step, we can now forecast future insurgent attacks for out-of-sample villages.

¹⁴ This is an important consideration especially given the small number of villages we surveyed. Indeed, accounting for spatial correlation even at the district level is difficult because each district contains a small number of villages. In particular, Figure 21 in the Online appendix shows that the relative level of ISAF support is positively, but not strongly, associated with the number of non-IED attacks when spatial correlations among villages within each district are taken into account (using a random effect model). Figure 22 shows moderate correlations between districts and covariates, explaining why including district-level random effects into regression models weakens the association between relative ISAF support and future insurgent attacks. Given this data limitation, the robustness of our findings is best examined using out-of-sample data.

Finally, we compare the predicted insurgent attacks with actual attacks for each out-of-sample village to evaluate our model's predictive performance.

Figure 4 uses a diagram to summarize our procedure for out-of-sample prediction. Below, we also provide the details about the entire prediction procedure.

Step 1: Building prediction models using in-sample villages

Step 1a: Regress ISAF support S on village-level covariates Z to estimate the *ISAF support model* (blue arrow in the left panel), $g(S|Z; \theta)$

Step 1b: Regress future violence Y on ISAF support S and other control variables X to estimate the *Insurgent violence model* (red arrows in the left panel), $f(Y|S, X; \beta, \lambda)$

Step 2: Predicting future violence using out-of-sample villages and the estimated models from Step 1

Step 2a: Using the village-level covariates Z for out-of-sample villages and the estimated *ISAF support model* from Step 1a, $g(S|Z; \hat{\theta})$, obtain the predicted ISAF support for out-of-sample villages \hat{S} (blue squiggly arrow in the right panel)

Step 2b: Using the predicted ISAF support from Step 2a, \hat{S} , and control variables X , obtain the predicted insurgent violence for each out-of-sample village, \hat{Y} , based on the estimated *Insurgent violence model* from Step 1a, $f(Y|S, X; \hat{\beta}, \hat{\lambda})$ (red squiggly arrows in the right panel)

Step 2c: Compare the predicted insurgent violence \hat{Y} against the actual insurgent violence Y for each out-of-sample village

We assess the accuracy of our out-of-sample prediction by comparing our forecast with the actual level of insurgent attacks. To measure the degree to which political attitudes improve forecasting performance, we compute the mean absolute forecasting error (MAFE) for two models: (1) one with prior attacks, ANSF/ISAF bases, and CDP projects as the predictors of future violence and (2) one with these three covariates and the estimated support level for ISAF as an additional predictor. This quantity is formally defined as $\sum_{i=1}^N |Y_i - \hat{Y}_i|/N$, where Y_i represents the number of observed future insurgent attacks for an out-of-sample village i and \hat{Y}_i is its prediction from a forecasting model.

We then compute the percentage improvement obtained by adding the estimated relative ISAF support level to the model with only prior insurgent attacks, ANSF/ISAF bases, and aid programs. Specifically, we compute $(MAFE_2 - MAFE_1)/MAFE_1 \times 100\%$, where

$MAFE_1$ and $MAFE_2$ are obtained from the model with and without the estimated relative ISAF support level variable, respectively.

Findings. We demonstrate the additional predictive gains from incorporating our measure of relative ISAF support level in the upper panel of Figure 5. Similar to Figure 3, we examine our forecasting performance across a wide range of temporal and spatial windows using contour plots for each category of insurgent violence. The inclusion of our estimate of relative support for ISAF improves predictions of the location and timing of IED attacks by up to 30%, a substantial improvement. This improvement is especially apparent in areas within 30 km of the village's center in the three months following the survey. Our model also improves, though to a lesser extent, predictions of the location and timing of non-IED attacks by up to 14%, again with the greatest improvement occurring near the 30 km distance mark. These patterns are consistent with our in-sample analysis (see Figure 3).

Do these improvements stem from the introduction of village- and district-level covariates Z , which we use, along with the survey data, to predict the estimated support level? To address this concern, we plot the percentage improvement attributable to these additional covariates alone in the lower panel of Figure 5. These covariates in fact lead to over-fitting and add little out-of-sample predictive power. For IED attacks, there is a modest improvement in predictive performance with a small temporal window, though the magnitude of the improvement is much less than the model with our measure of ISAF support. For non-IED attacks and found IEDs, the inclusion of these covariates actually *worsens* the predictive power of the models. Our analysis shows that while the estimated ISAF support for out-of-sample villages is a linear function of these covariates, these covariates are relatively poor predictors of future violence.

We emphasize that this comparison is possible only with respect to *out-of-sample* prediction performance. The Gauss-Markov theorem implies that the inclusion of these (or any other) covariates always improves the *in-sample* predictive performance (Hastie, Tibshirani & Friedman, 2009). Nevertheless, this does not mean that out-of-sample predictive performance always improves by adding more covariates. Indeed, what is striking about this finding is that these noisy covariates, which by themselves worsen out-of-sample predictive performance, can improve it when combined with our survey data.

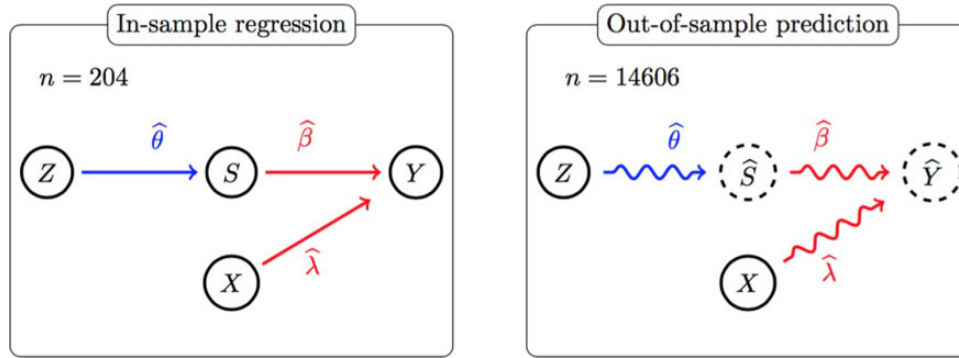


Figure 4. Out-of-sample prediction procedure

In Step 1 (left panel), using in-sample villages, we estimate the *ISAF support model* (blue arrow) and the *Insurgent violence model* (red arrow). For the former model, we regress ISAF relative support S on village- and district-level covariates Z such as log population, log elevation, ISAF control, Pakistan border, and Taliban Sharia. For the latter model, we regress future violence Y on S as well as village-level control variables (X) such as past violence, ANSF/ISAF bases, and aid projects. In Step 2 (right panel), we predict the ISAF support level for out-of-sample villages and then future violence using the models fitted in Step 1. For each out-of-sample village, we first estimate the ISAF relative support level \hat{S} (blue squiggly arrow) using the covariates Z and then predict Y with this estimated support level \hat{S} and the other covariates X .

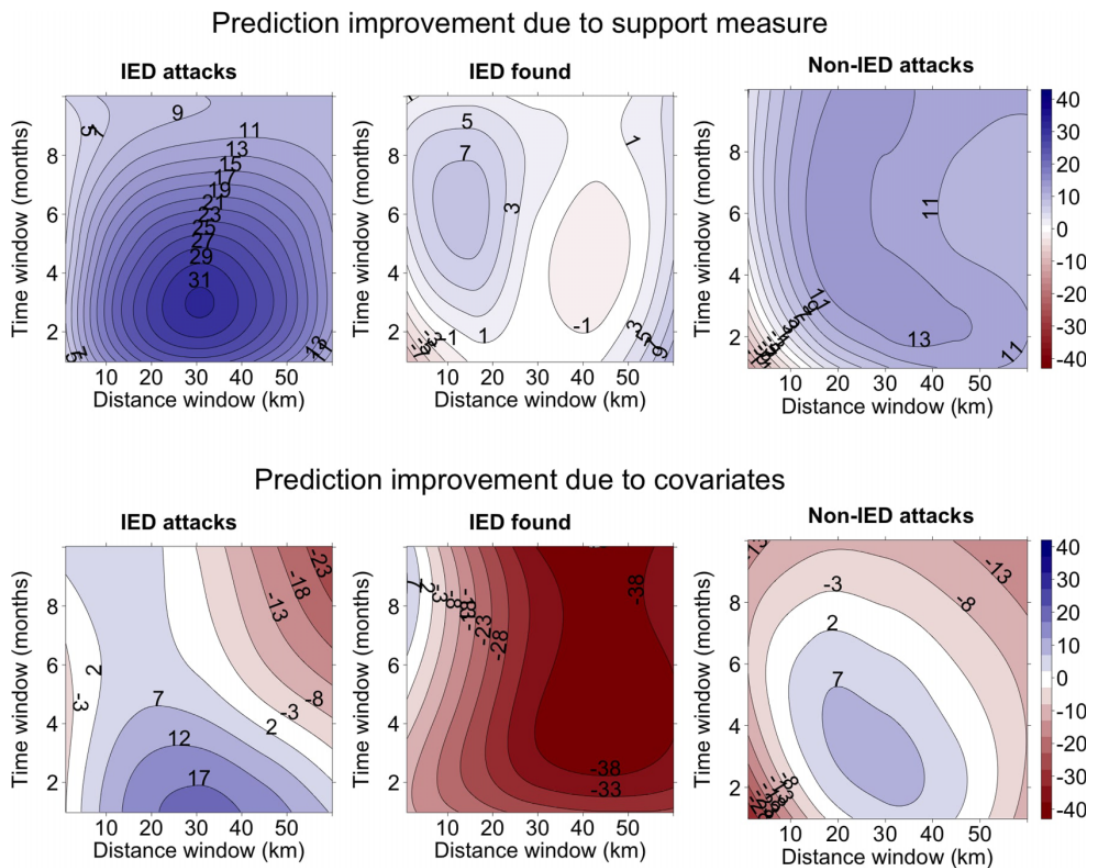


Figure 5. Out-of-sample forecasting performance

The upper panels depict forecasting improvement rates from adding estimated ISAF relative support level to the baseline model with past violence, counterinsurgent bases, and aid projects. Prediction improvement is measured by mean absolute forecasting errors derived from the baseline model ($MAFE_2$) and the model with the estimated support level ($MAFE_1$) – i.e. $(MAFE_2 - MAFE_1)/MAFE_1 \times 100\%$. The lower panels depict forecasting improvement rates from adding village- and district-level covariates to the baseline model.

These out-of-sample predictive improvements are once again robust to different modeling assumptions. In the Online appendix, we present corresponding results for models that interact relative ISAF support with prior violence, ISAF bases, and aid (see Figures 24, 25, and 26, respectively, in the Online appendix). For IED attacks, the inclusion of interaction terms does not improve predictive accuracy of future violence. By contrast, the use of the interaction term between support and past non-IED attacks and the interaction term between support and counterinsurgent bases increases our ability to predict future non-IED attacks.

We also consider the robustness of our finding using an alternative method of assessing predictive improvement. In the Online appendix (see Figure 23), we present results based on the root mean squared forecasting error (RMSFE). This measure is formally defined as $\sqrt{\sum_{i=1}^N (Y_i - \hat{Y}_i)^2 / N}$. Though RMSFE is more sensitive to outliers, the results largely agree with those presented above for IED attacks. The prediction improvement, however, does not exist for IED found and non-IED attacks. As before, adding village and district covariates instead of the predicted support measure to the base model generally does not increase predictive performance. In fact, these covariates *worsen* the base model's predictive accuracy for all three types of attacks.

Despite its clear methodological advantages, the use of out-of-sample testing remains relatively rare in the study of civil war violence. Yet this approach can yield powerful insights that are missed by relying solely on statistical significance to assess the importance of different covariates (Ward, Greenhill & Bakke, 2010). In particular, our out-of-sample modeling reveals that standard covariates in studies of civil war violence – including population size and terrain – add little to our predictive success. Instead, much of the predictive improvement stems from fine-grained covariates, including civilian attitudes, that track the spatial distribution of combatants and their actions.

Conclusion

Civilian attitudes are an important predictor of multiple types of insurgent violence in Afghanistan. A single study cannot, of course, definitively prove a link between attitudes and behavior; we take up the question of external generalizability in Section A.7 and Figure 22 of the Online appendix. But our evidence does suggest that the location, type, and lethality of insurgent violence are all shaped by the underlying spatial distribution of civilians' relative support for combatants. This association is

especially clear in the case of improvised explosive devices against ISAF forces, where villages with pro-counterinsurgent leanings attract a disproportionate share of their use.

These findings suggest several theoretical, methodological, and policy avenues for further research. Our contention that civilian attitudes inform insurgent targeting joins a nascent 'strong program' (Gutiérrez-Sanin & Wood, 2014: 213) of broader inquiry about how ideology shapes civil war dynamics. Scholars have only recently begun the difficult work of theorizing how (and when) combatant violence or development assistance affects civilian attitudes and subsequent behavior. In particular, the twin issues of how civilians form and update their views of combatants remain largely neglected in our theories. As a result, we call for not only a broader investigation into the dynamics of ideology during wartime but also a shift in emphasis to consider how states and rebels also condition their own tactical choice on the distribution of ideological support among the populace. In that vein, our framework could be extended by examining whether the location, type, or lethality of government- or ISAF-initiated actions also hinge on the spatial distribution of combatant support (Lyll, 2015). Ideology, after all, not only informs the preferences and strategies of individuals and organizations but also shapes how potential allies and rivals interpret the ideational terrain of the conflict and craft their own responses.

Methodologically, these findings underscore an additional advantage of micro-level data: disaggregating our units of observation facilitates out-of-sample testing. These tests not only avoid over-fitting our models to (sparse) data but also reveal that some well-known covariates, despite their statistical significance in current civil war research, actually worsen our predictive accuracy. Our combination of in- and out-of-sample testing could be extended to panel data on civilian attitudes, providing greater causal insight into the relationship between civilian attitudes and insurgent violence. Our use of survey experiments and predictive analysis could also be extended to a host of other sensitive issues – including interethnic relations, perceptions of government legitimacy, and corruption – and associated wartime behavior. Attitudinal indicators could form the basis of an early warning/early response (EW/ER) system that could help predict conflict onset or 'hot spots' within an ongoing war (Blair, Blattman & Hartman, 2014). In-sample findings could be cross-validated with out-of-sample predictions, leveraging the insights from a limited number of randomly chosen locations into thousands, or even

tens of thousands, of non-surveyed villages at reduced cost and potential harm to enumerators and respondents.

Those seeking policy recommendations may be initially disheartened by these findings. We are not dismissing ‘hearts and minds’ efforts wholesale. Instead, ours is a cautionary tale. Hearts and minds programs can generate unexpected consequences by creating targeting cues for insurgents among exactly the same populations where such efforts are finding the most traction. Unfortunately, both civilian and military aid planners have gravitated toward quick impact projects that provide little or no protection for civilians – including, for example, the US Army’s massive Commanders’ Emergency Response Program (CERP) in Iraq and Afghanistan and numerous USAID programs. Planners should anticipate the possibility that violence will increase, not decrease, at least in the early stages of programming. Instead of decamping for more favorable sites, aid programs should either work to maximize civilian protections (e.g. by close coordination with military forces) or avoid programming in areas where protection cannot be credibly extended.

Replication data

The dataset, codebook, and *R* files for the empirical analysis in this article, including all figures and supplemental on-line analyses, can be found at <http://www.prio.org/jpr/datasets> and on Harvard’s Dataverse at <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/DBSEQC>. All analyses were conducted using *R* and the freely available *endorse* package.

Acknowledgements

We thank the Editor and two anonymous reviewers as well as seminar participants at the University of Sydney and Princeton University for helpful comments on earlier versions.

Funding

We gratefully acknowledge financial support for the survey from the Institution for Social and Policy Studies’ Field Experiment Initiative and the Macmillan Center for International and Area Studies, both at Yale University. Additional support was also provided by the Air Force Office of Scientific Research (Lyll; Grant #FA9550-09-1-0314) and the National Science Foundation (Imai; Grant SES-0849715). This research was approved by Yale’s Human Subjects Committee under IRB protocol #1006006952.

References

- Beath, Andrew; Fotini Christia & Ruben Enikolopov (2011) Winning hearts and minds? Evidence from a field experiment in Afghanistan. MIT Political Science working paper no. 2011-14 (http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1809677).
- Beck, Nathaniel; Gary King & Langche Zeng (2000) Improving quantitative studies of international conflict: A conjecture. *American Political Science Review* 94(1): 21–35.
- Berman, Eli; Jacob Shapiro & Joseph Felter (2011) Can hearts and minds be bought? The economics of counterinsurgency in Iraq. *Journal of Political Economy* 119(4): 766–819.
- Blair, Robert; Christopher Blattman & Alexandra Hartman (2014) Predicting local violence. Working paper (http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2497153).
- Blair, Graeme; Kosuke Imai & Jason Lyall (2014) Comparing and combining list and endorsement experiments: Evidence from Afghanistan. *American Journal of Political Science* 58(4): 1043–1063.
- Blattman, Christopher & Jeannie Annan (2016) Can employment reduce lawlessness and rebellion? A field experiment with high-risk men in a fragile state. *American Political Science Review* 110(1): 1–17.
- Bohorquez, Juan; Sean Gourley, Alexander Dixon & Michael Spagat (2009) Common ecology quantifies human insurgency. *Nature* 462(17 December): 911–914.
- Braithwaite, Alex & Shane Johnson (2012) Space-time modelling of insurgency and counterinsurgency in Iraq. *Journal of Quantitative Criminology* 28(1): 31–48.
- Brandt, Ben (2011) The Taliban’s conduct of intelligence and counterintelligence. *CTC-Sentinel* 4(6): 19–23.
- Bullock, Will; Kosuke Imai & Jacob Shapiro (2011) Statistical analysis of endorsement experiments: Measuring support for militant groups in Pakistan. *Political Analysis* 19(4): 363–384.
- Chivers, Christopher J (2011) In eastern Afghanistan, at war with the Taliban’s shadowy rule. *New York Times* 6 February, A1.
- Christia, Fotini (2012) *Alliance Formation in Civil Wars*. Cambridge: Cambridge University Press.
- Crost, Benjamin; Joseph Felter & Patrick Johnston (2014) Aid under fire: Development projects and civil conflict. *American Economic Review* 104(6): 1833–1856.
- Department of the Army (2007) *U.S. Army Field Manual No. 3-24*. Chicago, IL: University of Chicago Press.
- Department of the Army (2014) *Field Manual 3-24: Insurgencies and Countering Insurgencies*. Washington, DC: Headquarters, Department of the Army.
- Drake, Charles M (1998) The role of ideology in terrorists’ target selection. *Terrorism and Political Violence* 10(2): 53–85.
- Fearon, James & David Laitin (2003) Ethnicity, insurgency, and civil war. *American Political Science Review* 97(1): 75–90.
- Flynn, Michael; Matt Pottinger & Paul Batchelor (2010) *Fixing Intel: A Blueprint for Making Intelligence Relevant in Afghanistan*. Washington, DC: Center for a New American Security.

- Giustozzi, Antonio (2013) Military adaptation by the Taliban, 2001–2011. In: Theo Farrell, Frans Osinga & James A Russell (eds) *Military Adaptation in Afghanistan*. Stanford, CA: Stanford University Press, 242–262.
- Goldstone, Jack; Robert Bates, David Epstein, Ted Robert Gurr, Michael Lustik, Monty Marshall, Jay Ulfelder & Mark Woodward (2010) A global model for forecasting political instability. *American Journal of Political Science* 54(1): 190–208.
- Gopal, Anand (2014) *No Good Men Among the Living: America, the Taliban, and the War Through Afghan Eyes*. New York: Macmillan.
- Gutiérrez-Sanin, Francisco & Elisabeth Wood (2014) Ideology in civil war: Instrumental adoption and beyond. *Journal of Peace Research* 51(2): 213–226.
- Hastie, Trevor; Robert Tibshirani & Jerome Friedman (2009) *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. 2nd edition. New York: Springer.
- Hill, Daniel, Jr & Zachary Jones (2014) An empirical evaluation of explanations for state repression. *American Political Science Review* 108(3): 661–687.
- Ho, Daniel; Kosuke Imai, Gary King & Elizabeth Stuart (2007) Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. *Political Analysis* 15(3): 199–236.
- iCasualties.org (2013) IED Fatalities (<http://icasualties.org/oef/>).
- Johnson, Thomas (2013) Taliban adaptations and innovations. *Small Wars and Insurgencies* 24(1): 3–27.
- Johnson, Thomas & Matthew DuPee (2012) Analyzing the new Taliban code of conduct (Layeha): An assessment of changing perspectives and strategies of the Afghan Taliban. *Central Asian Survey* 31(1): 77–91.
- Johnson, Neil; Spencer Carran, Joel Botner, Kyle Fontaine, Nathan Laxague, Philip Nuetzel, Jessica Turnley & Brian Tivnan (2011) Pattern in escalations in insurgent and terrorist activity. *Science* 333(6038): 81–84.
- Kalyvas, Stathis (2006) *The Logic of Violence in Civil War*. Cambridge: Cambridge University Press.
- Kalyvas, Stathis (2008) Ethnic defection in civil wars. *Comparative Political Studies* 41(8): 1043–1068.
- Kalyvas, Stathis (2012) Micro-level studies of violence in civil war: Refining and extending the control–collaboration model. *Terrorism and Political Violence* 24(4): 658–668.
- King, Gary & Langche Zeng (2001) Improving forecasts of state failure. *World Politics* 53(4): 623–658.
- Leites, Nathan & Charles Wolf (1970) *Rebellion and Authority: An Analytic Essay on Insurgent Conflicts*. Chicago, IL: Markham.
- Lyall, Jason (2010) Are co-ethnics more effective counterinsurgents? Evidence from the second Chechen war. *American Political Science Review* 104(1): 1–20.
- Lyall, Jason (2015) Bombing to lose? Airpower and the dynamics of coercion in counterinsurgency wars. Working paper (http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2422170).
- Lyall, Jason & Isaiah Wilson (2009) Rage against the machines: Explaining outcomes in counterinsurgency war. *International Organization* 63(1): 67–106.
- Lyall, Jason; Graeme Blair & Kosuke Imai (2013) Explaining support for combatants in wartime: A survey experiment in Afghanistan. *American Political Science Review* 107(4): 679–705.
- Lyall, Jason; Yuki Shiraito & Kosuke Imai (2015) Coethnic bias and wartime informing. *Journal of Politics* 77(3): 833–848.
- Mampilly, Zachariah (2011) *Rebel Rulers: Insurgent Governance and Civilian Life During War*. Ithaca, NY: Cornell University Press.
- Metternich, Nils; Cassy Dorff, Max Gallop, Simon Weschle & Michael Ward (2013) Antigovernment networks in civil conflicts: How network structures affect conflictual behavior. *American Journal of Political Science* 57(4): 892–911.
- Montgomery, Jacob; Florian Hollenbach & Michael Ward (2012) Improving predictions using ensemble Bayesian model averaging. *Political Analysis* 20(3): 271–291.
- Ruttig, Thomas (2010) *How Tribal Are the Taliban?* Kabul: Afghanistan Analysts Network (<http://www.afghanistan-analysts.org/wp-content/uploads/downloads/2012/10/20100624TR-HowTribalAretheTaliban-FINAL.pdf>).
- Sambanis, Nicholas; Jonah Schulhofer-Wohl & Moses Shayo (2012) Parochialism as a central challenge in counterinsurgency. *Science* 336(6083): 805–808.
- Schrodt, Philip (2014) Seven deadly sins of contemporary quantitative political analysis. *Journal of Peace Research* 51(2): 287–300.
- Shiraito, Yuki & Kosuke Imai (2012) endorse: R package for analyzing endorsement experiments (<http://imai.princeton.edu/software/endorse.html>).
- Staniland, Paul (2014) *Networks of Rebellion: Explaining Insurgent Cohesion and Collapse*. Ithaca, NY: Cornell University Press.
- Stoll, David (1993) *Between Two Armies in the Ixil Towns of Guatemala*. New York: Columbia University Press.
- Task Force 3-10 (2012) *State of the Taliban (January 6, 2012): Detainee Perspectives*. Bagram: NATO (<http://s3.documentcloud.org/documents/296489/taliban-report.pdf>).
- United Nations Assistance Mission Afghanistan (2011) *Afghanistan: Annual Report on Protection of Civilians in Armed Conflict, 2010*. Kabul: UNAMA Human Rights Unit (https://unama.unmissions.org/sites/default/files/engi_version_of_poc_annual_report_2011.pdf).
- Ward, Michael; Brian Greenhill & Kristin Bakke (2010) The perils of policy by p-value: Predicting civil conflicts. *Journal of Peace Research* 47(2): 363–375.
- Weidmann, Nils & Michael Ward (2010) Predicting conflict in space and time. *Journal of Conflict Resolution* 54(6): 883–901.
- Weinstein, Jeremy (2007) *Inside Rebellion: The Politics of Insurgent Violence*. Cambridge: Cambridge University Press.

World Bank (2012) *World Development Report 2011*. Washington, DC: World Bank.

Yonamine, James (2013) Predicting future levels of violence in Afghanistan districts using GDELT. Working paper (<http://data.gdeltproject.org/documentation/Predicting-Future-Levels-of-Violence-in-Afghanistan-Districts-using-GDELT.pdf>).

Zammit-Mangion, Andrew; Michael Dewar, Visakan Kadiramanathan & Guido Sanguinetti (2012) Point process modeling of the Afghan War Diary. *Proceedings of the National Academy of Sciences* 109(31): 12414–12419.

KENTARO HIROSE, b. 1981, PhD in Political Science (University of Chicago, 2013); Assistant Professor, Waseda University (2015–); Postdoctoral Fellow in Political Science,

Princeton University (2013–15); research area: military conflict and political methodology.

KOSUKE IMAI, b. 1975, PhD in Political Science (Harvard University, 2003); Professor, Department of Politics and Center for Statistics and Machine Learning, Princeton University (2003–); research area: political methodology, applied statistics; forthcoming book: *Quantitative Social Science: An Introduction* (Princeton University Press, 2017).

JASON LYALL, b. 1975, PhD in Government (Cornell University, 2005); Associate Professor, Yale University and Director, Political Violence FieldLab, Yale University (2013–); research area: violence, development, experiments.