

Can Civilian Attitudes Predict Civil War Violence?*

Kentaro Hirose[†]

Kosuke Imai[‡]

Jason Lyall[§]

December 1, 2013

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” as a prelude to influencing their wartime behavior. We use an original survey experiment in 204 villages to establish the robust association between civilian attitudes toward the Taliban and International Security Assistance Force (ISAF) and the timing and location of future village-level insurgent attacks in Afghanistan. We then extend our analysis to 14,606 non-surveyed villages to demonstrate how including 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 that use improvised explosive devices (IEDs). These improvements in predictive power remain even after accounting for the well-known ability of past attacks to predict future violence.

Key Words: Civil War; Public Opinion; Survey Experiment; Out-of-sample Prediction

*Financial support for the survey from Yale’s Institute for Social and Policy Studies’ Field Experiment Initiative and the Macmillan Center for International and Area Studies is gratefully acknowledged. Additional support from the Air Force Office of Scientific Research (Lyal; Grant #FA9550-09-1-0314) and the National Science Foundation (Imai; Grant SES-0849715) is also acknowledged. This research was approved by Yale’s Human Subjects Committee under IRB protocol #1006006952.

[†]Postdoctoral Fellow, Department of Politics, Princeton University, Princeton NJ 08544. Email: hirose@princeton.edu, URL: <http://scholar.princeton.edu/hirose>

[‡]Professor, Department of Politics, Princeton University, Princeton NJ 08544. Phone: 609-258-6601, Email: kimai@princeton.edu, URL: <http://imai.princeton.edu>

[§]Associate Professor, Department of Political Science, Yale University, New Haven, CT 06520. Phone: 203-432-5264, Email: jason.lyall@yale.edu, URL: <http://www.jasonlyall.com>

Are civilian attitudes a useful predictor of patterns of violence in civil wars? The past decade has witnessed the renewal of a debate over the importance of winning civilian “hearts and minds” in counterinsurgency wars such as Afghanistan and Iraq. Billions of dollars have been spent in these wars by militaries and development agencies alike to persuade civilians — through the provision of services and economic assistance — to support the central government, thereby reducing insurgent violence Bank (2012); Sambanis, Schulhofer-Wohl, and Shayo (2012); Beath, Christia, and Enikolopov (2011); Berman, Shapiro, and Felter (2011). Indeed, such efforts have been elevated to the level of official U.S. Army doctrine U.S. Army (2007).

Despite influential studies in psychology Paluck and Green (2009); Wallace et al. (2005); Kraus (1995); Ajzen and Fishbein (1977), economics Krueger and Malečková (2009); Chaudhuri and Holbrook (2001); Alesina, Baqir, and Easterly (1999), and political science Campbell, Green, and Layman (2011); Habyarimana et al. (2009) demonstrating the importance of attitudes for explaining and predicting behavior, much of the theoretical literature on civil war violence remains skeptical of the link between attitudes and behavior. Leading accounts, for example, emphasize the futility of measuring wartime attitudes accurately Kalyvas (2006, 2012); Leites and Wolf (1970) and, in some cases, dismiss entirely the need for contextual knowledge such as the “hearts and minds” of local residents Johnson et al. (2011).

Skepticism is clearly warranted. The danger of social desirability bias DeMaio (1984), where respondents hide their preferences while publicly adopting those of the interviewer, is rife in war zones. Even if they can be measured accurately, attitudes may be unreliable guides to future action since they may only reflect expedient decisions by civilians seeking to ensure their safety or continued economic assistance.

As a result, the small but growing literature on predicting violence has proceeded without incorporating any contextual measures such as relative support for the combatants among civilians Meier (2013); Montgomery, Hollenbach, and Ward (2012); Goldstone et al. (2010); Bohorquez et al. (2009). In one recent study Zammit-Mangion et al. (2012), for example, the authors draw

on Wikileaks' Afghan War Diary (AWD) and its 75,000 military logs for 2004–2009 to predict the number of insurgent attacks in 2010 at the weekly provincial level in Afghanistan. Their dynamic point process modeling of these data confirms the conventional wisdom that prior violence is an important predictor of future attacks in areas with significant levels of insurgent violence. And yet like others this study remains silent about the relationship between civilian attitudes and violence.

We adopt a different approach by directly exploring whether additional predictive gains can be made by incorporating survey data on civilian attitudes into forecasting models. Our approach unfolds over several steps. We first conduct a survey experiment in 204 villages in Pashtun-dominated provinces of Afghanistan — the heart of the Taliban insurgency — to measure combatant support using an indirect questioning methodology. We then utilize declassified insurgent attack data to demonstrate that a robust association exists between favorable attitudes toward the International Security Assistance Force (ISAF) and the timing and location of insurgent violence around these sampled villages. We then extend our analysis to 14,606 non-surveyed villages to validate our in-sample findings. We find that the addition of our measure of relative support for ISAF improves out-of-sample predictive performance by 20–30% over models that only include prior insurgent attacks. Moreover, we are able to make predictions over variable spatial and temporal windows at the village-level, a degree of granularity not approached by existing efforts to predict wartime violence.

1 Measuring Support

We measure support for ISAF and the Taliban using a survey of 2,754 respondents from 204 villages in five Pashtun-dominated provinces of Afghanistan (Logar, Kunar, Uruzgan, Helmand, and Khost). The survey was conducted in January-February 2011. These villages are located in some of the most violent areas of Afghanistan. Figure 1 illustrates the location of the originally sampled villages and the out-of-sample villages.

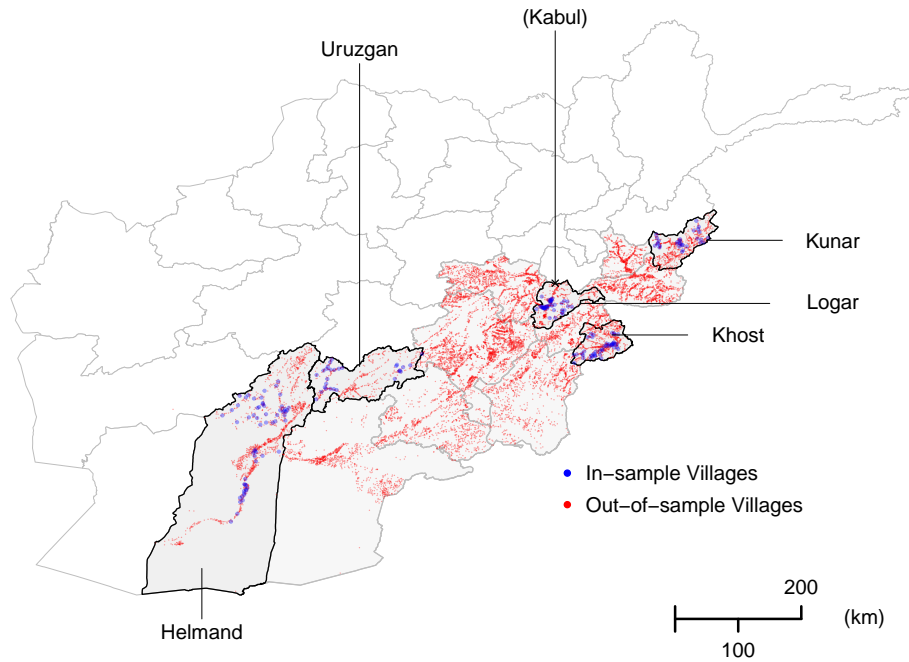


Figure 1: Location of 204 in-sample and 14,606 out-of-sample villages in Pashtun-dominated provinces of Afghanistan.

The sample was constructed using a multistage sampling method where first provinces, and then districts, villages, households, and finally individuals were randomly selected. Owing to security considerations and cultural constraints, all respondents were males aged 16 years and older. We obtained an 89% response rate; further survey design details are provided in Lyall, Blair, and Imai (2013).

Given the sensitivity of measuring support for armed combatants during wartime, we employ a battery of four endorsement experiments, a form of indirect questioning, to estimate support for ISAF and the Taliban. This indirect approach minimizes social desirability bias and item non-response when asking questions about sensitive issues Bullock, Imai, and Shapiro (2011); Blair et al. (2013). 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 when approached by enumerators. Our refusal rate was about 5%, a difference we ascribe to our indirect questioning method.

The mechanics of an endorsement experiment are straightforward. We first divide a sample of respondents randomly into groups. In the “control” group, respondents are asked to rate the level of their support for a particular policy. For those in the “treatment” group, the identical question is asked except that the policy is said to be endorsed by an actor of interest. The idea is to 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. 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 — within each sampled village.

Typically, multiple policies in the same policy domain are selected so that the measurement does not rely on a single instrument. Statistical power is also increased by analyzing them together. In our survey experiment, we employ four questions concerning 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 SI. Elsewhere, we provide detailed justifications for the choice of these policy questions Lyall, Blair, and Imai (2013).

A Bayesian hierarchical factor analytic model is used to pool the responses to these four questions together, creating an estimate of individual-level support for ISAF and the Taliban. We then aggregate individual-level support values to create an estimate of village-level support for each combatant. The exact model we use is described in the SI, and the details of this statistical methodology are described in Bullock, Imai, and Shapiro (2011). 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, and Lyall

(Forthcoming).

The resulting measures of village-level support are numerical estimates for each combatant. The support level for the Taliban ranges from -1.15 to 1.37 while that for ISAF ranges from -1.77 to 0.42 with positive (negative) values indicating support for (opposition against) the combatant. These results imply that Afghans are mostly opposed to ISAF; attitudes towards the Taliban are more mixed. 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 the following 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 -2.93 to 0.92 , suggesting that Afghans are far more supportive of the Taliban than ISAF. In the SI, we graph the spatial distribution and density of relative ISAF support in the surveyed villages.

Measuring Violence. We measure insurgent violence using declassified event data from ISAF’s Combined Information Data Network Exchange (CIDNE). These data record the date, location and nature of insurgent attacks against ISAF forces and installations throughout Afghanistan. Distinct from Wikileaks’ Afghan War Diary, these data represent the 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, we are not drawing on the full universe of insurgent violence. We use data from 10 months before and after our January-February 2011 survey for our prediction models; a total of 52,032 attacks were recorded over this period.

CIDNE tracks at least 14 discrete types of insurgent attacks that are relevant for our purposes here. We aggregate these types into two broad categories. First, we constructed an “improvised explosive device” (IED) category that includes 12,861 recorded IED and mine explosions, IEDs that were found and cleared by ISAF forces, and threatened and suspected IED emplacements.

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 using new electronic jammers and other means of unravelling the social networks that facilitate IED emplacement.¹ Second, we created a “Non-IED Attack” category that includes 39,171 insurgent attacks. These attacks include small arms fire, indirect fire (e.g., mortars) and rocket fire against ISAF forces and installations.²

In our analysis, these two categories are operationalized as count variables that record the number of relevant events in specified temporal windows before and after the survey’s fielding 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.

2 Analysis of In-sample Villages

We begin our analysis with the original 204 sampled villages. For simplicity, a linear regression model is used to estimate the association between relative support for ISAF and subsequent insurgent violence. We include a count variable recording the number of insurgent attacks prior to the survey because past violence levels are known to be the single best predictor of future violence Zammit-Mangion et al. (2012); Montgomery, Hollenbach, and Ward (2012); Bohorquez et al. (2009). We therefore examine how our measure of relative ISAF support is associated with *future* insurgent violence even after adjusting for *past* attacks. We also explored the possibility of a non-linear relationship between violence and attitudes using generalized additive models but concluded that the simple linear model captures most of the systematic variation.

As an illustrative example, we present our model’s results using a temporal window of five months pre- and post-survey. We use a 15km radius around each sampled village to calculate the

¹Specific CIDNE categories are (1) IED Explosion, (2) IED Found and Cleared, (3) IED Threat, (4) IED Cache/Find, (5) Hoax IED, (6) False IED, (7) Premature IED Detonation, (8) Suspected IED, (9) Mine Strike, and (10) Mine Found and Cleared.

²Specific CIDNE categories are: (11) Direct Fire, (12) Attack, (13) Raid/Ambush, and (14) Indirect Fire.

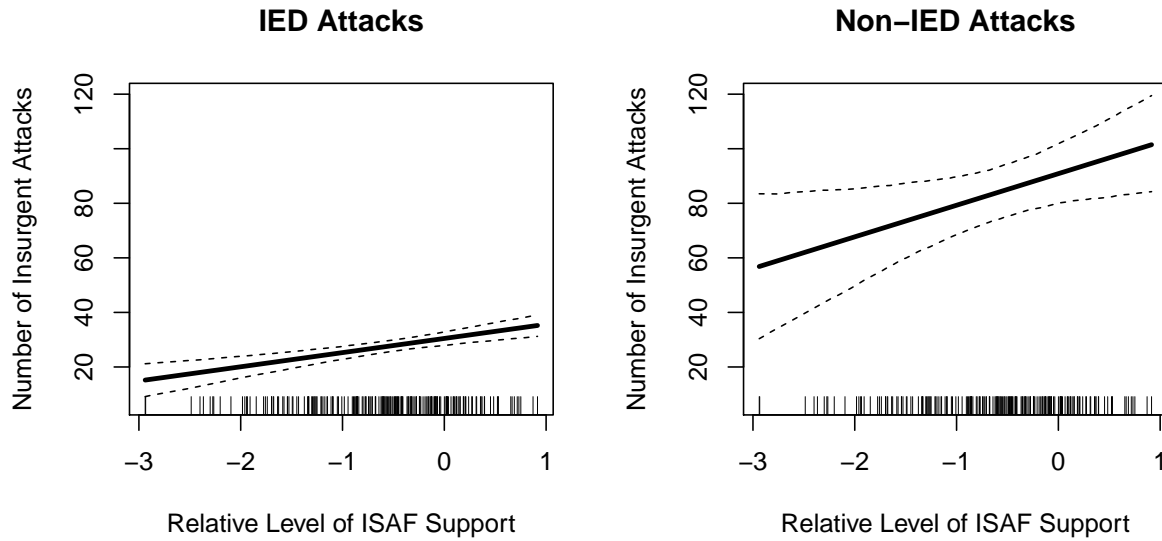


Figure 2: Positive association between number of future insurgent attacks and level of support for International Security Assistance Force (ISAF) relative to Taliban. The plots present statistically significant association between the number of insurgent attacks that have occurred within 15km of its perimeter during the five months after the fielding of our 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 15km around each village) during the five months prior to our survey. The results are based on the linear regression model estimated separately for each of the two violence categories where the number of future insurgent attacks is regressed on the relative level of ISAF support and the number of past insurgent attacks. The dashed lines represent 95% confidence intervals based on robust standard errors.

number of insurgent attacks. Note that each village has a slightly different start and end date given the order they were surveyed in. This results in different temporal windows (of the same size) across villages. Figure 2 demonstrates that a strong positive association exists between relative support for ISAF and future insurgent attacks even after accounting for prior violence. This is particularly true for IED attacks. A village that has modest relative support for ISAF (equivalent to a .5 value) is predicted to have an additional 16 IED attacks on average over the next five months (with a 95% confidence interval of [9, 23] using a heteroskedasticity-consistent standard error) when compared to a village strongly opposed to ISAF (equivalent to a -2.5 value).

This positive association between relative support for ISAF and future insurgent attacks does not depend on our choice of temporal and spatial windows. We repeat the analysis by varying the

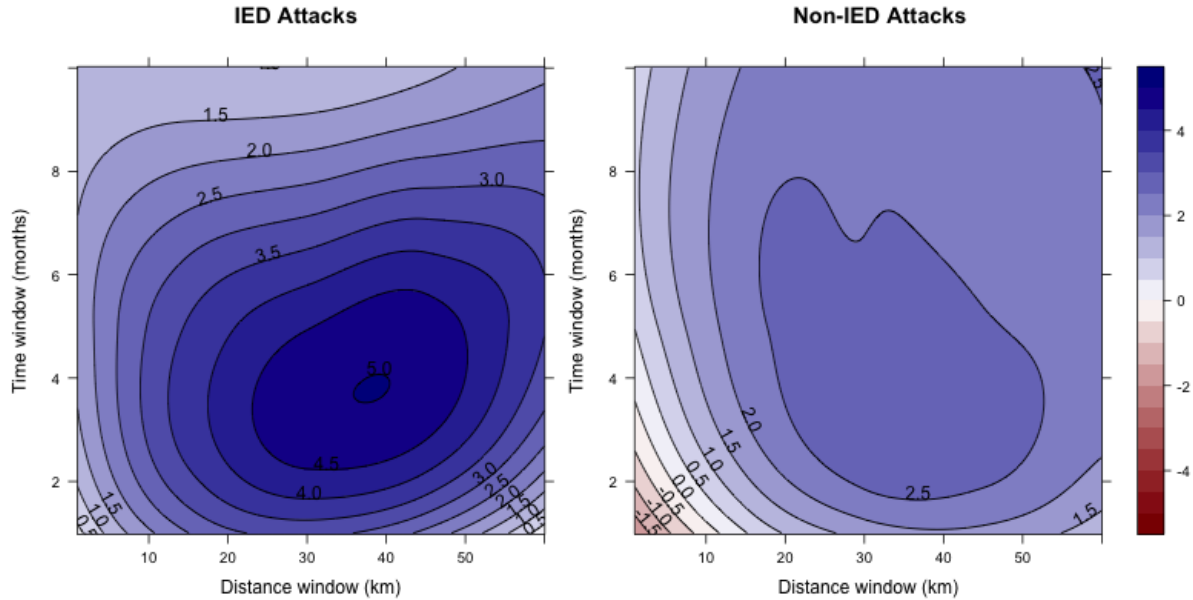


Figure 3: Robustness of positive association between the number of future insurgent attacks and the relative support for International Security Assistance Force (ISAF). The linear regression models which produced the results displayed in Figure 2 are repeatedly estimated using a wide range of time and distance windows (from 1 to 10 months before/after the survey and from 1 to 60km of each village’s perimeter). The contour plots show the resulting t statistics of the estimated coefficient of relative ISAF support, corresponding to its marginal effect on the number of future insurgent attacks, while adjusting for the prior level of insurgent violence. Robust standard errors are used.

temporal window from one to ten months and changing the radius around the surveyed village from 1km to 60km. Figure 3 presents contour plots of the t -statistics for the estimated coefficient of the relative ISAF support measure. We continue to observe a positive and statistically significant association between ISAF support and insurgent attacks while controlling for prior insurgent attacks. Moreover, this relationship is robust to modeling assumptions. In the SI, we present additional results based on matching Ho et al. (2007) where villages are first paired according to prior violence and then the pairwise difference in future violence is regressed on the pairwise differences in both relative ISAF support and prior violence.

As Figure 3 illustrates, we observe important variation in the model’s ability to predict IED and non-IED attacks. The positive relationship between pro-ISAF sentiment and insurgent IED emplacement is strongest at the 4 month interval and about 40 kilometers around the surveyed

village. The t -statistics for non-IED attacks, on the other hand, are relatively consistent throughout the 2–8 month post-survey time window and reach their highest values in the 20–45 kilometer mark. Comparing t -values across attack types, it is apparent that our model is especially well-suited to predicting IED emplacement, with peak t -values approaching 5.0 compared with a still sizable 2.5 for non-IED attacks.

These findings suggest that the Taliban may be using pro-ISAF sentiment as a cue to guide their targeting decisions. These patterns of violence are consistent with claims that the Taliban attack ISAF as a means of demonstrating to their supporters that they retain the ability to harm the counterinsurgent Lyall (2013). Since ISAF troop concentrations are generally highest in and around pro-ISAF villages, the Taliban appear to view these areas as optimum locations for imposing costs on ISAF. In addition, these attacks send a message to ISAF supporters in and around these villages: continuing to support the counterinsurgent is a risky proposition since the Taliban clearly retains the ability to reach, and harm, not just ISAF but its supporters even in these pro-ISAF locations.

It is also noteworthy that the model's predictive power is best within a relatively localized space, approximately 20–45 km, around the sampled villages. This reflects in part the density of ISAF troop concentration around relatively pro-ISAF villages. It is also consistent with the prevailing conception of Taliban recruitment dynamics as largely local and decentralized in nature. With most insurgents recruited from, and operating near, their home villages, it becomes more difficult for the Taliban to coordinate attacks across large geographic areas. As a consequence, we observe that our predictive power is greatest within a fairly constrained geographic area, one that encompasses the pro-ISAF villages and immediate neighbors but rarely scales up to encompass a district-sized area. These results underscore the need to adopt predictive models and data that are close to the action rather than aggregated at some convenient, but ultimately too coarse, subnational administrative unit such as a district or province.

3 Analysis of Out-of-sample Villages

To further investigate the predictive power of civilian attitudes, we extend our analysis beyond the original sample to 14,606 non-surveyed villages located in thirteen Pashtun-dominated provinces. We first predict relative ISAF support for these out-of-sample locations using village-level covariates and then forecast the number of insurgent attacks at these predicted support levels. As before, our aim is to examine whether these predicted support levels improve forecasting performance of future insurgent violence even while controlling for prior patterns of insurgent attacks.

The out-of-sample prediction proceeds in two steps. First, we estimate the relative ISAF support level for each of the non-surveyed villages. We do so using data from the original sample of villages and regress the estimated levels of relative support for ISAF on available village- and district-level characteristics. These 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, and whether the district bordered Pakistan. Employing the resulting regression equation and these same covariates for non-surveyed locations, we derive values for relative levels of ISAF support. Second, 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. With these rescaled estimates of pro-ISAF support, we use the same regression model detailed above to predict insurgent attacks.

We assess the accuracy of our out-of-sample prediction by comparing our forecast with the actual level of insurgent attacks. Note that this analysis utilizes violence data from the original surveyed villages *alone* to forecast insurgent attacks for out-of-sample locations. Data from non-surveyed villages are only used for validation purposes, guarding against over-fitting. To measure the degree to which political attitudes improve forecasting performance, we compute the *mean absolute forecasting error* (MAFE) for two models: (a) one with prior attacks alone as the predictor of future violence and (b) one with prior violence and the estimated support level for

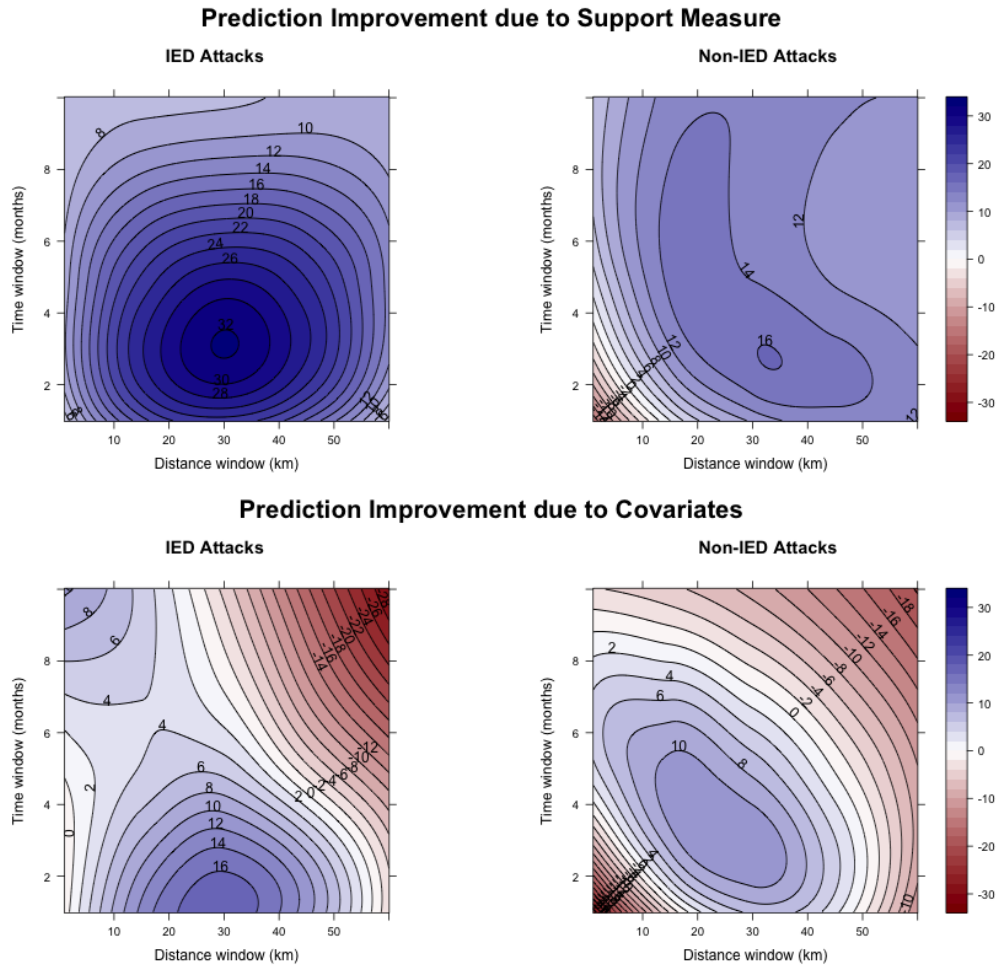


Figure 4: Improvement rate of adding the estimated ISAF support level to the model with the prior insurgent violence level variable alone. The contour plots show $(MAFE_2 - MAFE_1)/MAFE_1 \times 100\%$, where $MAFE_1$ and $MAFE_2$ in the upper two panels are the mean absolute forecasting errors from the models with and without the predicted ISAF support level variable. In the lower two panels, $MAFE_1$ and $MAFE_2$ correspond to the models with and without the village-and district-level covariates.

ISAF as an additional predictor.³ We then compute the percent improvement obtained by adding the estimated relative ISAF support level to the model with only prior insurgent attacks.⁴

We present our result in the upper panel of Figure 4. Similar to Figure 3, we examine our forecasting performance across variable temporal and spatial windows using contour plots for each category of insurgent violence. The inclusion of our measure of relative support for ISAF

³This 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.

⁴Specifically, 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.

improves predictions of the location and timing of IED explosions by up to 32%, a substantial improvement over the base model. This improvement is especially apparent in areas within 30km of the village's center in the 3 months following the survey. Our model also improves predictions of the location and timing of non-IED attacks by up to 16%, again with the greatest improvement occurring near the 30km distance mark and consistently in the 3 month temporal window. 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 rather than our measure of relative ISAF support? To address this concern, we plot the percentage improvement attributable to these covariates in the lower panel of Figure 4.⁵ These covariates in fact add little predictive power to the model. For IED attacks, there is a modest improvement in predictive performance, though the magnitude of the improvement is much less than the model including our measure of ISAF support. For non-IED attacks, the inclusion of these covariates actually worsens the predictive power of the model.

Our out-of-sample analysis confirms many of the insights about Taliban violence uncovered by the in-sample analysis. For both categories of violence, we observe that our predictive power is quite localized, usually peaking about 30km around a village. This is consistent again with a targeting strategy that uses a local population's relative pro-ISAF leanings to determine the location of Taliban attacks. Given this highly selective nature of Taliban violence, it is not surprising that the improvement in prediction performance drops once we move beyond the village and its immediate neighbors.

These findings are robust to different modeling assumptions. In the SI, we present corresponding results for a model that interacts prior attacks with relative ISAF support. If anything, these results are slightly stronger than those outlined in Figure 4: we record $\geq 28\%$ and $\geq 22\%$ improvement in predictive power for IED and non-IED attacks, respectively. Our results are also robust to alternative measures of prediction performance. In the SI, we offer results based on the

⁵The base model includes prior violence as the predictor, while the expanded model includes prior violence and the village- and district-level covariates used to estimate support levels for the out-of-sample villages.

root mean squared forecasting error (RMSFE).⁶ 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 non-IED attacks. As before, adding the 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 predictive accuracy of the base model for both types of attacks.

4 Conclusion

Although the link between civilian attitudes and civil war violence remains contested among scholars and practitioners, we used a survey experiment in Afghanistan to demonstrate that measures of political attitudes can be used to substantially improve the predictive performance of forecasting models of village-level insurgent violence. Our measures of relative ISAF support not only improve prediction of both IED and non-IED attacks by 20–30% while accounting for prior insurgent violence but do so at a far more fine-grained level of analysis than previously attempted. This approach could be extended to a host of other sensitive attitudes — including interethnic relations, perceptions of government legitimacy, and corruption — and associated wartime behavior. It might also inform the basis of an early warning/early response (EW/ER) system that could predict future flash points in a conflict. More generally, the approach underscores the importance of tailored surveys that draw on indirect questioning and multistage random sampling to create a representative sample of respondents and locations. In turn, this sample can be extended to out-of-sample predictions that leverage a few hundred carefully chosen locations into tens of thousands of non-surveyed villages, minimizing both cost and potential harm to enumerators and respondents alike.

⁶This measure is formally defined as $\sqrt{\sum_{i=1}^N (Y_i - \hat{Y}_i)^2 / N}$.

References

- Ajzen, Icek, and Martin Fishbein. 1977. "Attitude-Behavior Relations: A Theoretical Analysis and Review of Empirical Research." *Psychological Bulletin* 84 (5): 888-918.
- Alesina, Alberto, Reza Baqir, and William Easterly. 1999. "Public Goods and Ethnic Divisions." *Quarterly Journal of Economics* 114: 1243-84.
- Bank, World. 2012. *World Development Report 2011*. Washington, D.C.: World Bank.
- Beath, Andrew, Fotini Christia, and Ruben Enikolopov. 2011. "Winning Hearts and Minds? Evidence from a Field Experiment in Afghanistan." MIT Political Science Working Paper No.2011-14.
- Berman, Eli, Jacob Shapiro, and Joseph Felter. 2011. "Can Hearts and Minds Be Bought? The Economics of Counterinsurgency in Iraq." *Journal of Political Economy* 119: 766-819.
- Blair, Graeme, Christine Fair, Neil Malhotra, and Jacob Shapiro. 2013. "Poverty and Support for Militant Politics: Evidence from Pakistan." *American Journal of Political Science* 57 (1): 30-48.
- Blair, Graeme, Kosuke Imai, and Jason Lyall. Forthcoming. "Comparing and Combining List and Endorsement Experiments: Evidence from Afghanistan." *American Journal of Political Science*.
- Bohorquez, Juan, Sean Gourley, Alexander Dixon, and Michael Spagat. 2009. "Common ecology quantifies human insurgency." *Nature* 462: 911-914.
- Bullock, Will, Kosuke Imai, and Jacob Shapiro. 2011. "Statistical Analysis of Endorsement Experiments: Measuring Support for Militant Groups in Pakistan." *Political Analysis* 19 (4): 363-384.

- Campbell, David, John Green, and Geoffrey Layman. 2011. "The Party Faithful: Partisan Images, Candidate Religion, and the Electoral Impact of Party Identification." *American Journal of Political Science* 55: 42-58.
- Chaudhuri, Arjun, and Morris Holbrook. 2001. "The Chain of Effects from Brand Trust and Brand Affect to Brand Performance: The Role of Brand Loyalty." *Journal of Marketing* 65: 81-93.
- DeMaio, T. J. 1984. "Social Desirability and Survey Measurement: A Review." In *Surveying Subjective Phenomena*, ed. C.G. and E. Martia Turner. Vol. 2. New York: Russell Sage Foundation.
- Goldstone, Jack, Robert Bates, David Epstein, Ted Robert Gurr, Lustik Michael, Monty Marshall, Jay Ulfelder, and Mark Woodward. 2010. "A Global Model for Forecasting Political Instability." *American Journal of Political Science* 54: 190-208.
- Habyarimana, James, Macartan Humphreys, Daniel Posner, and Jeremy Weinstein. 2009. *Coethnicity: Diversity and the Dilemmas of Collective Action*. New York: Russell Sage Foundation.
- Ho, Daniel, Kosuke Imai, Gary King, and Elizabeth Stuart. 2007. "Matching as Nonparametric Preprocessing for Reducing Model Dependence in Parametric Causal Inference." *Political Analysis*, 15: 199-236.
- iCasualties.org. 2013. "IED Fatalities."
- Johnson, Neil, Spencer Carran, Joel Botner, Kyle Fontaine, Nathan Laxague, Philip Nuetzel, Jessica Turnley, and Brian Tivnan. 2011. "Pattern in Escalations in Insurgent and Terrorist Activity." *Science* 333: 81-84.
- Kalyvas, Stathis. 2006. *The Logic of Violence in Civil War*. Cambridge: Cambridge University Press.

- Kalyvas, Stathis. 2012. "Micro-Level Studies of Violence in Civil War: Refining and Extending the Control-Collaboration Model." *Terrorism and Political Violence* 24: 658-668.
- Kraus, Stephen. 1995. "Attitudes and the Prediction of Behavior: A Meta-Analysis of the Empirical Literature." *Personality and Social Psychology Bulletin* 21 (1): 58-75.
- Krueger, Alan, and Jitka Malečková. 2009. "Attitudes and Action: Public Opinion and the Occurrence of International Terrorism." *Science* 325: 1534-36.
- Leites, Nathan, and Charles Wolf. 1970. *Rebellion and Authority: An Analytic Essay on Insurgent Conflicts*. Chicago: Markham Publishing Company.
- Lyall, Jason. 2013. "Airpower and Coercion in Counterinsurgency Wars: Evidence from Afghanistan." *Unpublished Manuscript, Yale University*.
- Lyall, Jason, Graeme Blair, and Kosuke Imai. 2013. "Explaining Support for Combatants in Wartime: A Survey Experiment in Afghanistan." *American Political Science Review* 107 (November): 679-705.
- Meier, Patrick. 2013. "Crisis Mapping in Areas of Limited Statehood." In *Information and Communication Technologies in Areas of Limited Statehood*, ed. Steven Livingston and Gregor Walter-Drop. Oxford: Oxford University Press.
- Montgomery, Jacob, Florian Hollenbach, and Michael Ward. 2012. "Improving Predictions Using Ensemble Bayesian Model Averaging." *Political Analysis* 20: 271-291.
- Paluck, Elizabeth, and Donald Green. 2009. "Prejudice Reduction: What Works? A Critical Look at Evidence from the Field and the Laboratory." *Annual Review of Psychology* 60: 339-367.
- Sambanis, Nicholas, Jonah Schulhofer-Wohl, and Moses Shayo. 2012. "Parochialism as a Central Challenge in Counterinsurgency." *Science* 336: 805-808.

Shiraito, Yuki, and Kosuke Imai. 2012. “endorse: R Package for Analyzing Endorsement Experiments.”.

U.S. Army. 2007. *U.S. Army Field Manual No. 3-24*. Chicago: University of Chicago Press.

Wallace, David, Renè Paulson, Charles Lord, and Charles Bond Jr. 2005. “Which Behaviors Do Attitudes Predict? Meta-Analyzing the Effects of Social Pressure and Perceived Difficulty.” *Review of General Psychology* 93 (3): 214-227.

Zammit-Mangion, Andrew, Michael Dewar, Kadiramanathan, and Guido Sanguinetti. 2012. “Point process modeling of the Afghan War Diary.” *Proceedings of the National Academy of Sciences* 109.

Supporting Information

The Endorsement Experiment

Our endorsement experiment uses four questions regarding domestic policy reform in order to estimate support levels for the Taliban and ISAF. The exact question wording is reproduced below.

Prison Reform

- **CONTROL CONDITION:** A recent proposal calls for the sweeping reform of the Afghan prison system, including the construction of new prisons in every district to help alleviate overcrowding in existing facilities. Though expensive, new programs for inmates would also be offered, and new judges and prosecutors would be trained. How strongly would you support this policy?
- **TREATMENT CONDITION:** A recent proposal by foreign forces [or the Taliban] calls for the sweeping reform of the Afghan prison system, including the construction of new prisons in every district to help alleviate overcrowding in existing facilities. Though expensive, new programs for inmates would also be offered, and new judges and prosecutors would be trained. How strongly would you support this policy?

Direct Election

- **CONTROL CONDITION:** It has recently been proposed to allow Afghans to vote in direct elections when selecting leaders for district councils. Provided for under Electoral Law, these direct elections would increase the transparency of local government as well as its responsiveness to the needs and priorities of the Afghan people. It would also permit local people to actively participate in local administration through voting and by advancing their own candidacy for office in these district councils. How strongly would you support this policy?
- **TREATMENT CONDITION:** It has recently been proposed by foreign forces [or the Taliban] to allow Afghans to vote in direct elections when selecting leaders for district councils. Provided for under Electoral Law, these direct elections would increase the transparency of local government as well as its responsiveness to the needs and priorities of the Afghan people. It would also permit local people to actively participate in local administration through voting and by advancing their own candidacy for office in these district councils. How strongly would you support this policy?

Independent Election Commission

- **CONTROL CONDITION:** A recent proposal calls for the strengthening of the Independent Election Commission (IEC). The Commission has a number of important functions, including monitoring presidential and parliamentary elections for fraud and verifying the identity of candidates for political office. Strengthening the IEC will increase the expense of elections and may delay the announcement of official winners but may also prevent corruption and election day problems. How do you feel about this proposal?
- **TREATMENT CONDITION:** A recent proposal by foreign forces [or the Taliban] calls for the strengthening of the Independent Election Commission (IEC). The Commission has a number of important functions, including monitoring presidential and parliamentary elections for fraud and verifying the identity of candidates for political office. Strengthening the IEC will increase the expense of elections and may delay the announcement of official winners but may also prevent corruption and election day problems. How do you feel about this proposal?

Anti-Corruption Reform

- **CONTROL CONDITION:** It has recently been proposed that the new Office of Oversight for Anti- Corruption, which leads investigations into corruption among government and military officials, be strengthened. Specifically, the Offices staff should be increased and its ability to investigate suspected corruption at the highest levels, including among senior officials, should be improved by allowing the Office to collect its own information about suspected wrong-doing. How do you feel about this policy?
- **TREATMENT CONDITION:** It has recently been proposed by foreign forces [or the Taliban] that the new Office of Oversight for Anti- Corruption, which leads investigations into corruption among government and military officials, be strengthened. Specifically, the Offices staff should be increased and its ability to investigate suspected corruption at the highest levels, including among senior officials, should be improved by allowing the Office to collect its own information about suspected wrong-doing. How do you feel about this policy?

The Statistical Model for the Endorsement Experiment

Following Bullock, Imai, and Shapiro (2011), we use a statistical model to estimate support levels for ISAF and the Taliban by efficiently combining the responses to multiple endorsement experiment questions. To do so, we model each respondent's answer to a policy question as a

function of his or her support for the endorser as well as policy preference. Specifically, we apply the following Bayesian ordered probit factor analytic model:

$$\Pr(Y_{ij} \leq l \mid T_i = k) = \Phi(\alpha_{jl} - \beta_j(x_i + s_{ijk})), \quad (1)$$

where $Y_{ij} \in \{1, 2, 3, 4, 5\}$ represents respondent i 's answer to the j th policy question (1 = Strongly agree, 2 = Agree, 3 = Indifferent, 4 = Disagree, and 5 = Strongly disagree) and respondent i 's status regarding the randomized treatment assignment is denoted as $T_i \in \{0, 1, 2\}$ (0 = Control, 1 = ISAF, and 2 = Taliban). The latent variable s_{ijk} measures respondent i 's support level for endorser k in policy j with a greater value of s_{ijk} indicating a higher level of support. For identification, s_{ij0} is fixed at zero. Finally, the latent variable x_i represents the degree to which respondent i is in favor of policy reform in general. The ‘‘item difficulty’’ parameter α_{jl} measures the popularity of the j th policy reform independent of the endorser, while the ‘‘discrimination’’ parameter β_j expresses the degree to which the reform proposal differentiates pro- and anti-reform respondents. We assume $\alpha \sim \mathcal{N}(0, 25)$ and $\beta \sim \mathcal{TN}(0, 25)$ as the priors.

We model the individual-level support s_{ijk} and ideal point x_i using a hierarchical modeling technique with village-level random effect parameters $\lambda_{village[i]}$ and $\delta_{village[i]}$ as follows,

$$s_{ijk} \sim \mathcal{N}(\lambda_{village[i]} + Z_i^\top \lambda_k^Z, \omega_k^2) \quad (2)$$

$$x_i \sim \mathcal{N}(\delta_{village[i]} + Z_i^\top \delta^Z, 1) \quad (3)$$

where Z_i represents the set of individual-level covariates. As the priors, we assume $\lambda \sim \mathcal{N}(0, \psi^2)$, $\delta \sim \mathcal{N}(0, \sigma^2)$, and $\psi^2, \sigma^2, \omega^2 \sim Inv - \chi^2(5, 2)$,

We use an R package `endorse` developed by Shiraito and Imai (2012) to fit this model. The convergence is monitored by running multiple Markov chains with over-dispersed starting values. Using the posterior simulation draws, we compute each respondent's average support level for each endorser across the four policy areas, and then further aggregate it to village-level support by averaging the individual-level estimates.

Additional Figures

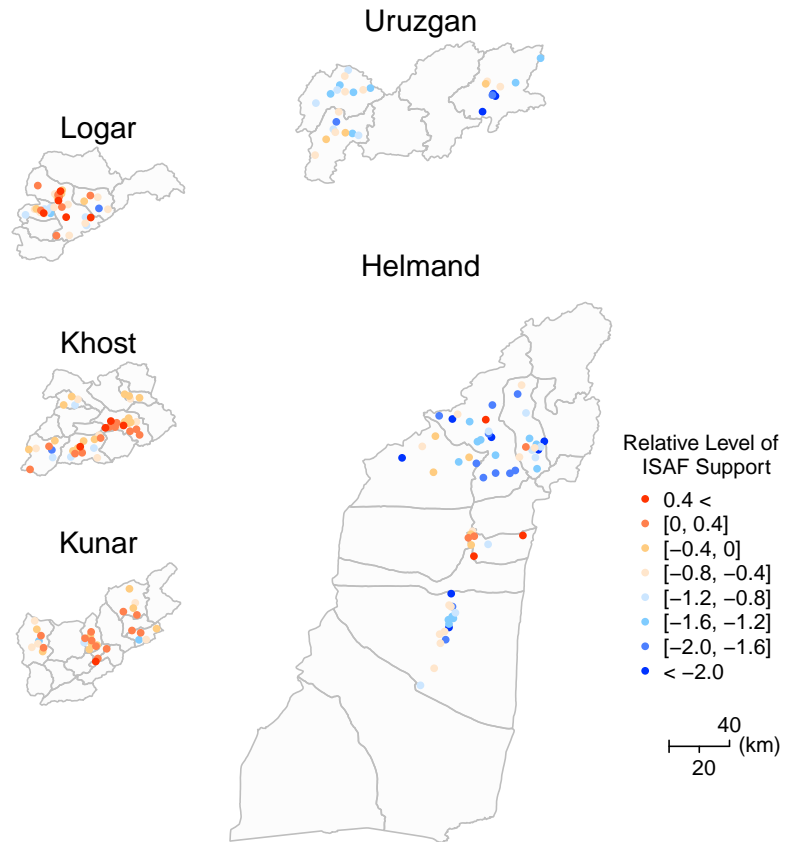


Figure 5: Spatial distribution of relative ISAF support in surveyed villages. The maps show five Pashtun-dominated provinces. The lines within the provinces represent district borders.




Figure 6: Density of relative ISAF support in surveyed villages.




Figure 7: Results of matching analysis for association between the number of future insurgent attacks and the relative support for ISAF. Villages are first paired according to prior violence and then the pairwise difference in future violence is regressed on the pairwise differences in both relative ISAF support and prior violence. The contour plots show the resulting t statistics of the estimated coefficient of relative ISAF support.




Figure 8: Improvement rate using an interaction term between ISAF support levels and prior violence. The contour plots show $(MAFE_2 - MAFE_1)/MAFE_1 \times 100\%$, where $MAFE_1$ and $MAFE_2$ are the mean absolute forecasting errors from the models with and without the predicted ISAF support level variable.

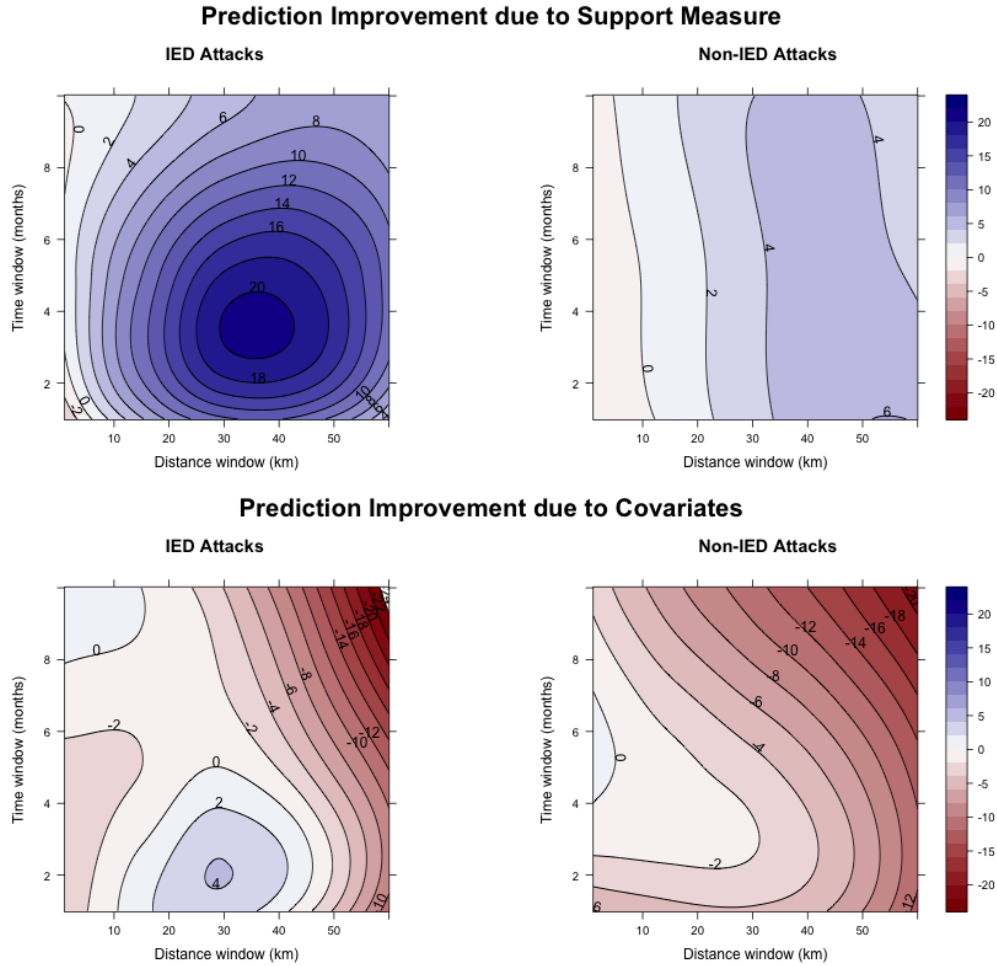


Figure 9: Improvement rate based on the root mean squared forecasting error (RMSFE). The contour plots show $(\text{RMSFE}_2 - \text{RMSFE}_1) / \text{RMSFE}_1 \times 100\%$, where RMSFE_1 and RMSFE_2 in the upper two panels are the root mean squared errors from the models with and without the predicted ISAF support level variable. In the lower two panels, RMSFE_1 and RMSFE_2 correspond to the models with and without the village-and district-level covariates.