



Public response to government alerts saves lives during Russian invasion of Ukraine

David Van Dijke^{a,b,1} , Austin L. Wright^{c,1,2}, and Mark Polyak^b

Edited by David Laitin, Stanford University, Stanford, CA; received December 4, 2022; accepted March 26, 2023

War is the cause of tremendous human suffering. To reduce such harm, governments have developed tools to alert civilians of imminent threats. Whether these systems are effective remains largely unknown. We study the introduction of an innovative smartphone application that notifies civilians of impending military operations developed in coordination with the Ukrainian government after the Russian invasion. We leverage quasi-experimental variation in the timing of more than 3,000 alerts to study civilian sheltering behavior, using high-frequency geolocation pings tied to 17 million mobile devices, 60% of the connected population in Ukraine. We find that, overall, civilians respond sharply to alerts, quickly seeking shelter. These rapid postalert changes in population movement attenuate over time, however, in a manner that cannot be explained by adaptive sheltering behavior or calibration to the signal quality of alerts. Responsiveness is weakest when civilians have been living under an extended state of emergency, consistent with the presence of an alert fatigue effect. Our results suggest that 35 to 45% of observed civilian casualties were avoided because of public responsiveness to the messaging system. Importantly, an additional 8 to 15% of civilian casualties observed during the later periods of the conflict could have been avoided with sustained public responsiveness to government alerts. We provide evidence that increasing civilians' risk salience through targeted government messaging can increase responsiveness, suggesting a potential policy lever for sustaining public engagement during prolonged episodes of conflict.

war | population displacement | humanitarian crisis | public policy | information nudges

Interstate military disputes, especially those that unfold in densely populated urban areas, disrupt civilian life, undermining human welfare and reducing economic activity. In an attempt to reduce civilian casualties and promote freedom of movement, governments often engage in extensive messaging about where and when potential attacks may occur. Governments of 40 countries, representing at least 58% of the global population, have developed these messaging systems to address persistent threats to civilians' physical security. By informing the public during conflict, governments and aligned private actors may prevent harm.

Despite the importance of alert systems and their extensive use in states under conflict, there is little evidence to date on whether and under what conditions these alerts impact public behavior. Although it is intuitive to expect civilians to take immediate, costly action in response to alerts of imminent harm, it is not clear how quickly they respond to these messages, and, indeed, whether they do so at all (1). This evidence gap is largely due to our inability to reliably measure how people's movements shift in the moments following notification of imminent threats. Yet public response to this type of high-frequency, localized messaging remains a first-order concern for public policy (2). In order to minimize harm while enabling continued economic and social activity during conflict, public actors need a mechanism for transmitting information that shapes mobility and enables the public to seek shelter and calibrate their movements with respect to the militarized environment.

Study Overview. We provide credible estimates of immediate behavioral change among the public in response to government alerts about imminent risk. We study these dynamics in Ukraine following the February 2022 invasion by Russian forces. After the incursion of military forces into urban areas, the Ukrainian government coordinated the development of a smartphone application for transmitting public alerts about impending Russian military operations. These messages were then recirculated via a collection of mobile device applications as well as through social media platforms (e.g., Telegram) and traditional air sirens. The content of the messages corresponded to that of a traditional siren, indicating the start and end of an air alert, together with

Significance

War often puts civilians at risk. To minimize casualties, governments can use information to shape how civilians evade harm, by seeking shelter from violence. This study looks at an application-based warning system in Ukraine, using device-by-minute changes in movement to understand postalert behavior. We find that public response to the system strongly reduced civilian casualties, though attenuation in this response over time likely led to avoidable deaths. As such, policy interventions that boost engagement with public messages during conflict may help save lives. We show that civilian shelter response in Ukraine is higher when risks are more salient, which suggests messaging strategies that target civil attitudes to risk may bolster the effectiveness of warning systems.

Author affiliations: ^aDepartment of Economics, University of Michigan, Ann Arbor, MI 48104; ^bRisk Analytics Division, Ipsos Public Affairs, Washington, DC 20006; and ^cHarris School of Public Policy, University of Chicago, Chicago, IL 60637

Author contributions: D.V.D., A.L.W., and M.P. designed research; D.V.D. and A.L.W. performed research; D.V.D. and A.L.W. contributed new reagents/analytic tools; D.V.D. and A.L.W. analyzed data; M.P. contributed data access and feedback; and D.V.D. and A.L.W. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2023 the Author(s). Published by PNAS. This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹D.V.D. and A.L.W. contributed equally to this work.

²To whom correspondence may be addressed. Email: austinlw@uchicago.edu.

This article contains supporting information online at <http://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2220160120/-DCSupplemental>.

Published April 24, 2023.

its location (3). This was the case across platforms, so we treat the multiplatform message as a single notification transmitted through several channels of communication. To quantify the information available to civilians, we compile these messages about locational and temporal risks. We combine the location and timing of these messages with high-frequency, spatially disaggregated data on device mobility. This pairing of messages and mobility enables us to study whether device movement changes discontinuously as air raid alerts are transmitted to civilians and their mobile devices. This quasi-experimental approach provides credible estimates of costly, real-world responses to alerts during conflict.

Relying on estimates from more than 3,000 local, device-by-minute event studies, we document five core findings: i) Civilians, on average, respond sharply to alerts, rapidly increasing their movement patterns as they flee imminent harm; ii) these rapid postalert changes in civilian movement attenuate substantially as the war progresses; iii) this attenuation cannot be explained by adaptations in sheltering such as seeking protection in underground bunkers or using an alternative tactic called the “two wall rule”; iv) public responsiveness attenuates even when civilians are exposed to higher-quality information; v) postalert movement patterns attenuate more rapidly when the local population has been living under an extended “state of alarm,” where they have been exposed to a high duration of recent bombardment alerts. Taken together, these results are consistent with the presence of an alert fatigue effect.

To quantify the consequences of diminished public responsiveness to government messages, we conduct a series of counterfactual exercises to estimate the excess civilian casualties prevented by the alert system as well as those that could have been avoided in the absence of any alert fatigue. Linking the mobility response to a source of information on civilian harm, we find that the overall mobility in response to the alerts significantly reduced the number of civilian casualty events, but diminished responsiveness also led to a large number of plausibly avoidable civilian deaths. In particular, our bounding exercises suggest that 35 to 45% of potential civilian casualties were avoided through the messaging system, although between 8 and 15% of observed civilian casualties could have additionally been avoided if postalert responsiveness had remained the same over time. These figures suggest that government messaging can be a powerful tool to minimize harm during war, but public engagement with these alerts is essential. To that end, we present suggestive evidence of the potential for salient government messaging, over and above basic alerts, to increase public engagement. First, we document that the public shelter response to alerts is nearly twice as large on days when the Ukrainian government sent out special nationwide alerts regarding Russian operations in addition to the usual air alerts. Second, we estimate that devices that are one standard deviation (275 km) closer to the front line at the time of an alert have a shelter response that is nearly twice as large. This gives an indirect estimate of the potential gains in engagement that can be achieved by stimulating civilian risk perceptions through appropriate messaging interventions.

In general, our findings imply, on the one hand, that civil defense alert systems play an important role in protecting civilians from harm. On the other hand, we provide evidence that engagement with these systems declines with repeated exposure and with decreased civilian risk salience, leading to avoidable deaths. This suggests that further optimization of government messaging strategies during conflict can lead to important welfare gains. In particular, open questions remain regarding what types

of messages are most effective at boosting public engagement and sustaining this boost even in the presence of continued, high-frequency signals.

Literature and Contribution. This research addresses a number of prominent, open questions in the social sciences. Prior work has considered how civilians respond to weather shocks and natural disasters (4), localized economic shocks (5), and political crises (6). Other work has considered how civilians’ decisions to flee are shaped by exposure to violence during war (7, 8). Less focus has been paid to how public response during a period of heightened uncertainty can be shaped by an information operation run by a government actor (9). The literature on early warning systems for natural disasters has long recognized the importance of a people-centered approach to the development of these systems (10). Recent research in this literature has studied the public perception of and response to early warning systems through surveys, for example, in the context of earthquakes (11, 12) and floods (13, 14). Some evidence also exists on the public’s psychological response to air raids in the Second World War, based on medical records (15). The public’s immediate behavioral (nonsurvey) response to warning systems, however, is understudied (16, 17), especially in the context of mobile warning message systems (18–20). We fill this gap in the literature by providing credible estimates of the public’s immediate response to air alerts using high-frequency mobile device location data.

Moreover, research on informational nudges typically focuses on low-cost, low-stakes, one-shot settings, where behavioral change may have a marginal effect on worker productivity (21), engagement with public services (22, 23), or engagement with politics (24). We focus on a setting that is repeated many times in a short time span, where acquiring and disseminating information to the public is relatively high cost and where the public’s response to the informational nudge is typically very high stakes. In the cases where research has focused on high-risk settings, this work has focused almost exclusively on downstream behavioral proxies or on a narrow geographic context. Our paper fills this gap by directly studying the behavioral outcome of interest—whether and how much civilians evade danger—in a country-scale, repeated quasi-experimental setting.

This paper also addresses a prominent gap in our understanding of decisions under risk. Prior work has found that risk profiles remain largely stable over time (25), with exogenous shocks typically triggering an increase in preferences for certainty (26). Research on these topics is usually limited by the use of lab-based measures of low-stakes and/or hypothetical decision-making. These measures are also difficult to track over time and may not have much transferability to high-stakes decisions under risk. Our study speaks directly to these gaps in design and setting by studying device-by-minute variation in responsiveness over time, in response to a multitude of high-risk information shocks. Unlike prior work, we find that responsiveness attenuates over time, even after we account for the possibility that this is driven by civilians acquiring higher-quality information (false positive alarms) or substituting risk through alternative channels (sheltering underground or inside the home). Moreover, a gap in responsiveness quickly emerges over time between subjects exposed to longer versus shorter emergency alert duration. This broadly suggests that risk profiles do not demonstrate a tendency toward risk aversion in a high-stakes setting. Instead, our results are consistent with cognitive fatigue generating negative externalities.

Conceptual Framework. We anticipate, all else equal, that civilians will respond to warnings about an imminent threat by seeking shelter. In our primary design, this will involve evading harm through rapid movement. This overall effect may vary over time, as the dynamics of the conflict environment change. There are three plausible mechanisms that could explain why escaping harm through evasion (population movement) might attenuate over time. First, civilians may have quickly adapted to the threat environment, seeking shelter underground—possibly in newly erected or designated underground shelters—rather than fleeing above ground; or sheltering inside of their own homes using the so-called “rule of two walls.” Second, any attenuation in responsiveness could be due to civilians calibrating whether signals from the government are informative (i.e., government messages in a given area may be undermined by a high false positive rate). Third, attenuation could be due to alert fatigue (alternatively, normalization to risk), which is likely accelerated by the time civilians spend under persistent “states of emergency” (longer duration alerts). We evaluate these arguments empirically using a combination of high-frequency alerts and population mobility information, observed over time.

Results

Evading Harm through Movement. We begin by studying whether and how civilians respond to bombardment alerts. This is shown in Fig. 1 for the pooled sample between March and September 2022. Across our various outcomes and across periods, the prealert indicators help to validate that movement patterns were not significantly shifting prior to the threat notification being sent. This suggests that civilians were not acting in anticipation of a future threat prior to the alert, which is evidence in favor of the credibility of our identification strategy. After the alert is sent, we find a large, consistently positive effect of the notification on overall movement as well as speed. Civilians move quickly to avoid the potential military operation that prompted the alert.

Attenuation over Time. We test whether these effects vary heterogeneously across periods.* We have split the conflict into three phases: the first 2 mo of recorded alert activity; the second two months, when Russian forces had settled into certain areas and were engaging in regular activities; and a final phase, when Russian forces had largely reached stasis or were losing ground in certain regions. These are the three response profiles in Fig. 2, where increasing warmth of color corresponds to a later stage of the war. The public was most responsive to alerts in the first two periods of the conflict. However, the postalert response diminished significantly over time, suggesting limited civilian sheltering as the war drags on. This pattern continues to hold when we consider extended postalert windows of 1 and 4 h for the subsets of alerts that last at least that long, in *SI Appendix, Fig. SI-8*.† The pattern also holds across urban and rural areas, as shown in *SI Appendix, Fig. SI-7*, though the relative decline in response is more muted in rural areas.‡

*We discuss sample composition changes and changes in the app's notifications in *SI Appendix*.

†Shelter response appears to persist throughout these extended alarms. There may be a slight reversion to zero at the 3-h mark, though the 4-h estimates are rather noisy due to the small sample size, with 50% of alarms lasting less than an hour.

‡We designate urban areas as those covered by 131 “cities of regional significance”—an administrative division that was replaced in 2020—plus Kyiv and Sebastopol (which had a separate administrative designation). This covers all Ukrainian cities with a population of over 250,000.




Fig. 1. Strong overall public response to bombardment alerts. *Notes:* The figure shows pooled estimates of air alert event studies between March and September 2022. Figure documents changes in distance traveled (sum) across time (minutes relative to treatment) for the 61-min window centered around the minute of the air alert notification. The vertical line separates the preperiod trends from the posttreatment effects. Colored bands depict 95% confidence intervals of a smoothed local linear (loess) regression estimated from the 10 and 31 event study dummies on the *Left* and *Right* sides of minute 0 for each of the underlying event studies (*SI Appendix* for the event study specification).

Adapting to Air Alerts. The attenuation in civilian movement after alerts may also reflect adaptation—increasing use of potentially improvised on-site or nearby bomb shelters or of the so-called “rule of two walls” (pravilo dvokh stin), which substitutes underground sheltering with sheltering in indoor spaces that are separated from the outdoors by at least two walls.

We investigate the first of these alternative hypotheses by leveraging the spatial telemetry of cellphone devices, which includes a device’s altitude and thus allows us to estimate vertical population movement. We anticipate that civilians engaging in postalert flight will take advantage of underground infrastructure to avoid potential bombardment risks. We reproduce the event study specification estimated above, switching the outcome of interest from distance traveled to vertical movement via discernible changes in altitude. These results are shown in Fig. 3. There are significant reductions in altitude postalert for the earliest period of the conflict. This is consistent with civilians seeking shelter underground. However, the estimated effects following the alert attenuate to roughly zero during the later stages of the conflict. This indicates that the public was not adapting to the bombardment risks by substituting spatial flight for sheltering below ground. Instead, our results suggest that civilians were similarly less likely to engage in efforts to avoid bombardment overall after the alerts were circulated.

To assess whether the observed attenuation over time can be explained by civilians increasingly substituting actual bomb shelters for indoor sheltering using the “rule of two walls,” we isolate subsets of devices that were plausibly unable to shelter indoors at the time the alarm went off. In particular, in Fig. 4, we replicate Fig. 2 for devices that were at least 100 m away from home throughout the entire alert window (*Left* panel) or traveling at a speed of at least 0.3 km/h (5 m/min) throughout the alert window. The observed attenuation in shelter response is almost identical to the one observed for the entire sample, suggesting that the decline in response is not driven by devices that increasingly substitute at-home or indoor shelters for bomb shelters. It is, of course, likely that a subset of devices does rely on the rule of two walls, but Fig. 4 suggests that this subset is not increasing substantially over time in a way that could explain the observed attenuation. This is further supported by the fact that

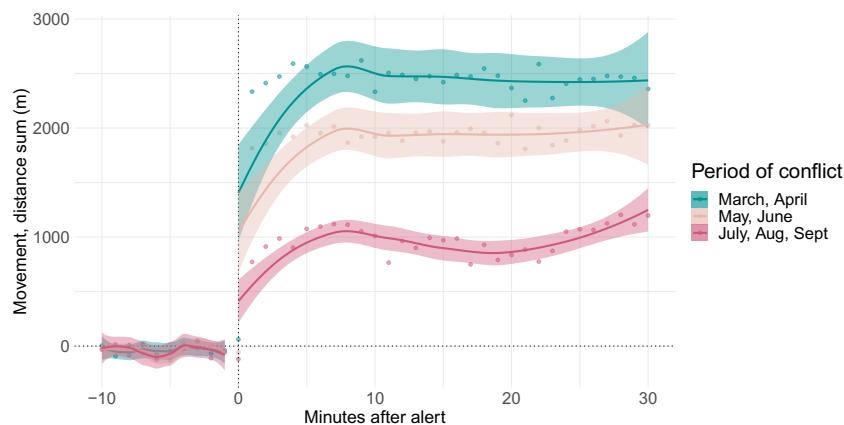


Fig. 2. Strong public response to alerts declines as war progresses. *Notes:* The figure shows pooled estimates of air alert event studies between March and September 2022. The figure shows changes in the distance traveled (sum) across time (minutes relative to treatment) for the 61-min window centered around the minute of the air alert notification. The vertical line separates the preperiod trends from the posttreatment effects. Periods are indicated by various colors. Colored bands depict 95% confidence intervals of a smoothed local linear (loess) regression estimated from the 10 and 31 event study dummies on the *Left* and *Right* sides of minute 0 for each of the underlying event studies (*SI Appendix* for the event study specification).

the Government of Ukraine was already distributing information about the rule of two walls as early as March 16 (27), which suggests that most citizens were likely aware of this option from early on in the war, instead of awareness of the rule spreading gradually among the population. Moreover, in what follows, we document a robust negative correlation between local shelter response and local civilian casualties. This suggests that at-home sheltering is not a perfect substitute for bomb shelters, which would explain why we observe a substantial shelter response in the first period even if citizens were already fully aware of the rule of two walls at the time. Indeed, while bomb shelters are purpose-built for withstanding aerial bombings, the two-wall

rule is likely a less effective sheltering alternative, as confirmed by reports of drone attacks on residential buildings (28) and by Kyiv's emergency management chief in various public media reports (29).

Calibrating Signal Quality. Another central mechanism that could explain attenuation is public calibration of information quality. Messages transmitted by the government vary in quality, as measured by the extent to which an anticipated military operation materializes in a particular location at a specific time during which there was an alert sent to the public. Since conflict events are salient, we anticipate that civilians are aware




Fig. 3. Altitude response suggests decline in movement not due to increased sheltering underground. *Notes:* The figure shows pooled estimates of air alert event studies between March and September 2022. The figure shows changes in altitude (e.g., movement up/down within buildings) across time (minutes relative to treatment) for the 61-minute window centered around the minute of the air alert notification. The vertical line separates the preperiod trends from the posttreatment effects. Periods of conflict are designated using various colors. Effects are shown in level changes. Colored bands depict 95% confidence intervals of a smoothed local linear (loess) regression estimated from the 10 and 31 event study dummies on the *Left* and *Right* sides of minute 0 for each of the underlying event studies (*SI Appendix* for the event study specification).

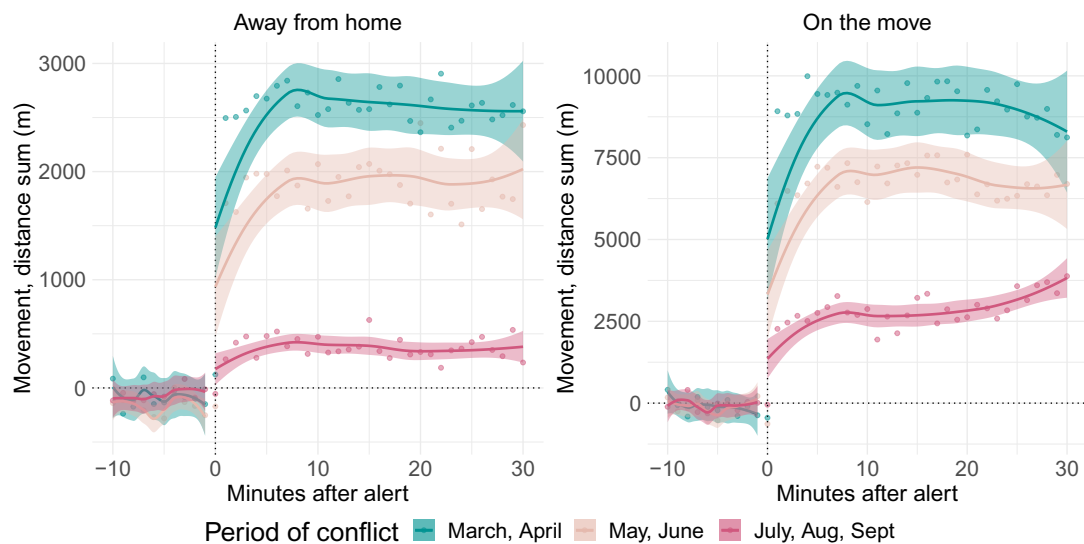


Fig. 4. Outdoors response suggests decline in movement not due to increased use of “two-wall rule”. *Notes:* Figures show pooled estimates of air alert event studies between March and September 2022, for subsamples of devices that are (*Left panel*) at least 100 m away from their estimated home locations during all minutes in the event window; (*Right panel*) “on the move”: Traveling faster than 300 m per hour (5 m per minute) on average during all minutes in the event window. The figure shows changes in distance traveled (sum) across time (minutes relative to treatment) for the 61-min window centered around the minute of the air alert notification. The vertical line separates the preperiod trends from the posttreatment effects. Periods are indicated by various colors. Colored bands depict 95% confidence intervals of a smoothed local linear (loess) regression estimated from the 10 and 31 event study dummies on the *Left* and *Right* sides of minute 0 for each of the underlying event studies (*SI Appendix* for the event study specification).

of and respond to conflict. In another study, we test this conjecture and find a robust association between retrospective self-reports of communal violence exposure and local violent events (as measured through the VIINA platform) during the sample period in Ukraine. If civilians learn over time that signal quality varies dynamically, they may adjust their behavior accordingly. In particular, we would expect to see less attenuation in responsiveness among devices receiving higher-quality signals (lower false positive rates) if the attenuation is driven by devices calibrating with respect to information quality. We test this conjecture using trends in local false-positive rates in 14-, 7-, and 3-d windows prior to the alert of interest. Following our other tests, we split the sample at the median of these false

positive rates and compare the period-by-period change in responsiveness across populations with higher- and lower-quality signals. These results are shown in Fig. 5. We find no evidence of diminished attenuation among populations exposed to higher-quality information. Instead, the trend in attenuation remains similar across high- and low-quality signals. These results are also stable if we use alternative windows (*SI Appendix*, Figs. SI-10 and SI-11).

Alert Fatigue. We evaluate whether the duration of alert exposure influences public responsiveness. To measure alert duration, we take advantage of information about when alerts are active (start and end time) as well as when alerts occur within a




Fig. 5. Attenuation in responsiveness to alerts present even among population exposed to high-quality information about risk. *Notes:* Figures show period-specific estimates of air alert event studies between March and September 2022. *Left* shows changes in distance traveled (sum) across time (minutes relative to treatment) for the 61-min window centered around the minute of the air alert notification. The vertical line separates the preperiod trends from the posttreatment effects. Effects are shown in meters. Samples are split at the median of false alerts using a 14-d bandwidth (additional results demonstrate robustness). Results are shown for March/April. *Middle* shows changes in distance traveled during May/June. *Right* shows changes in distance traveled between July and September. Colored bands depict 95% confidence intervals of a smoothed local linear (loess) regression estimated from the 10 and 31 event study dummies on the *Left* and *Right* sides of minute 0 for each of the underlying event studies (*SI Appendix* for the event study specification).




Fig. 6. Increased exposure to alerts associated with decreased response. *Notes:* Figures show period-specific estimates of air alert event studies between March and September 2022. *Left* shows changes in distance traveled (sum) across time (minutes relative to treatment) for the 61-min window centered around the minute of the air alert notification. The vertical line separates the preperiod trends from the posttreatment effects. Effects are shown in meters. Samples are split at the median of alert exposure (duration of alerts within regions and time periods) calculated as a 21-d moving average. Results are shown for March/April. *Middle* shows changes in distance traveled during May/June. *Right* shows changes in distance traveled between July and September. Colored bands depict 95% confidence intervals of a smoothed local linear (loess) regression estimated from the 10 and 31 event study dummies on the left and right sides of minute 0 for each of the underlying event studies (*SI Appendix* for the event study specification).

moving window of fixed width (21 d). We then split the sample of notification alerts based on whether the alert occurs during a period of high (above median) alarm exposure or not. Subsequently, we visualize how the public responds to bombardment alerts during various conflict periods, allowing our estimated effects of notification to vary with recent trends in the duration of time under alarm. These results are introduced in Fig. 6. Notice in the *Upper Left* panel that high- and low-duration estimates overlap significantly, suggesting that increased movement patterns after alert notifications were consistent early in the conflict. However, as the conflict reaches the later phases, in the *Middle* and *Right* panels, the total movement response declines, and it declines disproportionately for individuals who have been exposed to an intense information pressure to seek shelter. These effects are also consistent if we use cumulative duration of exposure, rather than recent trends in exposure, indicating that alerts with longer lengths may undermine the public's responsiveness to government messages. Taken together, these results suggest that the observed attenuation in shelter response was primarily driven by alert fatigue.

Policy Impact of Notification: Avoided and Avoidable Harm to Civilians. We next estimate how many lives were saved through deployment of the notification system. We also estimate the number of avoidable excess deaths due to nonresponsiveness in later periods.

In order to evaluate these counterfactuals, we rely on several parameters: the mobility–casualty association; the damage multiplier linking observed military activity with casualty counts; the counterfactual change in mobility, which depends on each policy scenario; and the intensity of alert activity during a given period. To recover these parameters, we begin by calculating the mobility–casualty association using the early period, finding that increased postal alert movement significantly reduces subsequent casualties in the alert window. We also gather statistics about mobility across alerts throughout the study period and *SI Appendix* data that allow us to identify civilian casualty counts. We combine these measures to calculate the counterfactual avoided casualties, where the mobility–casualty relationship allows us to calibrate the downstream effects of observed (and counterfactual)

responsiveness to alerts, taking into account the time-varying intensity of the conflict and impacts of changes in weapons technology and lethality. We provide additional technical details in *SI Appendix*.

Our first counterfactual exercise investigates how many potential casualties were avoided given the observed level of mobility triggered by early warning alerts. We estimate that an additional 1,617 civilian casualties were avoided relative to a counterfactual with no postal alert movement. This avoided loss of life represents 45% of observed casualties recorded in our primary casualty data during the sample period ($\sigma_{ce} = .093$).[§] Alternatively, if we adjust the counterfactual to allow for some anticipatory sheltering even in the absence of an alert, equivalent to what is observed during the final period of the study, we estimate that alerts prevented an additional 1,269 civilian casualties (approximately 35% of observed casualties, $\sigma_{ce} = .075$). We believe these to be credible counterfactuals as it is unlikely that civilians could accurately anticipate the exact timing of air alerts without information from the military's aircraft and missile detection systems. These exercises suggest that the public benefit due to government alerts was a counterfactual reduction of civilian casualties between 35% and 45%.

Our second counterfactual exercise investigates how many excess deaths plausibly occurred due to nonresponsiveness in later periods. That is, if responsiveness remained high, how much additional welfare gain could have been achieved using the alert system. If civilians moved, on average, as much as they did during this initial phase, how many civilian casualties could have been avoided? Additional details on the methodology used for calculating excess death are provided in *SI Appendix*.

We present several excess death estimates, which help bound the figures we present.[¶] We first incorporate day-specific measures of postal alert mobility as well as day-specific weights that help us adjust the casualties-per-event calculation. We next adjust our measure of postal alert mobility using trends in movement observed between the first and third periods (but keeping our casualty

[§]Additional details on the calculation of counterfactual variability are presented in *SI Appendix*.

[¶]We illustrate the corresponding trends in excess deaths in *SI Appendix*, Fig. SI-1.

weights fixed). This approach is most similar to the split-sample approach in our main design. We finally combine the trends in movement with trends in event severity, which allows us to smooth out spikes in casualties per event due to a sudden but temporary shift in weapon lethality. The time series of excess deaths is shown in [SI Appendix, Fig. SI-1](#). The first and second approaches bound the estimated number of excess deaths due to the observed alert fatigue between 8% ($\sigma_{ce} = .015$) and 15% ($\sigma_{ce} = .024$). The third approach bounds it at 12% ($\sigma_{ce} = .020$).

Optimizing Messaging During War: Theory and Evidence

In the light of the evidence presented above, what steps could policymakers take to address alert fatigue during an ongoing conflict?

First, theoretical models of persuasion provide insights into ways to optimize messaging (30), including using stylized models to clarify complex dynamics in order to persuade message recipients (31).[#] In the setting we study, this would involve providing information that allows recipients—members of the public—to fit data they have about the world (i.e., their wartime experiences) with present risk (e.g., the specific threat triggering a message from the government). Particularly relevant is prior work on information operations during conflict (9). This work suggests that messages which provide a narrative model can significantly increase welfare-enhancing but costly behaviors among recipients.^{||}

Second, a battery of observational and experimental studies has investigated the effectiveness of various message types and narrative primes on public behavior (35). This work has demonstrated, for example, that highlighting the behaviors of neighbors or members of the respondent's social network can significantly impact attitudinal and behavioral outcomes in high-stakes settings (23, 36, 37). Information treatments that emphasize personal health risks lead to an increased willingness to change health-related behaviors (38). Importantly, these primes shaped behavior across polarized ideological groups. Messages that emphasize communal behaviors and collective incentives also affect behavior (39, 40).

We add to this body of evidence by studying how government messages that aim to increase the salience of air raid risks among the public affect their sheltering response. Although the application studied above helps transmit government alerts by providing a communication channel, it did not, during our sample period, communicate any additional information besides from that conveyed by traditional air sirens: start, end, location, and type of possible air attacks. This suggests that there is room to complement the air alerts with alternative types of messaging that highlight risk salience, such as the cues relating to social context and personal risk studied in the literature.^{**} As it is difficult to obtain experimental estimates in our setting due to practical and ethical constraints, we rely on two observational measures of government messaging and risk salience.

First, we leverage the timing of nationwide alerts sent by the government about highly credible intelligence concerning Russian operations before and during several important periods of public remembrance, most notably Victory Day, the

Day of Mourning and Commemoration of War Victims, and Independence Day. Since information about these threats was broadcast widely during preparations for these special periods, we study whether the combination of local alerts and macrolevel information about the credibility of potential threats significantly enhances public responsiveness. Importantly, these episodes of heightened alert occurred during periods when responsiveness was trending downward, enabling us to study whether a double-barreled message (local alerts combined with a signal of credibility) can motivate sheltering amid increasing alert fatigue. We find that responsiveness to these combined messages, relative to the median of similar alerts during the same time periods, increased by approximately 50% ($\beta = 593.27$, $P < .001$), helping to close the sheltering gap in later periods. This large positive effect on shelter response compared to other days further suggests that there is indeed room for additional types of messaging to bolster response.

Second, we estimate how the shelter response varies with a device's distance to the front line. We expect that citizens perceive the risk the war poses to their safety to be higher the closer they are to the front line and that they accordingly respond more strongly to the threat of air raids.^{††} As such, being closer to the front line should mimic the influence of an effective messaging campaign: increasing risk salience. Building on this intuition, we reestimate the event study in our main specification, allowing the shelter response to vary with a device's distance to the front line, within any given alert region. Using this approach, we find that devices that are farther from the front line at the time an alert is activated respond much less strongly to the alert ([SI Appendix, Fig. SI-14](#)): being one standard deviation (275 km) away from the front line decreases, on average, the response by about 50% compared to being right at the front line. This gap in response persists across periods, although it narrows in the last period, when the response of devices close to the front line also declines. Importantly, the threshold for the activation of an air alert is identical for locations with varying proximity to the frontline within each alert region. Moreover, the associated likelihood of an actual bombing should not be expected to vary substantially within a given region either. This suggests that there is indeed room for policymakers to heighten people's sensitivity to risk, all else equal. Though this reduced-form exercise cannot tell us which kinds of salient messages are effective, it helps to give a sense of the potential welfare gains associated with increasing citizens' risk salience.

Overall, these two exercises suggest that additional government messaging aimed at increasing the salience and credibility of potential threats can increase civilian shelter response over and above the level induced by simple alert messages alone. As such, this type of messaging may be effective in addressing the alert fatigue this study has documented and thus increase the estimated positive welfare impact of early warning alert platforms. Future work could study in more detail which types of messages are most effective in conflict settings, perhaps combining alerts with informative nudges about harm avoidance and various statistics on localized communal harm. It may also be important to consider how to sustain durable public engagement in the presence of high-frequency primes. Prior work, cited above, largely focuses on the impact of one information treatment, or on regular but infrequent messages. The context we study, on the other hand, is characterized by a high volume of messages that may, in a manner independent of otherwise persuasive message

[#]For additional models of signal-based persuasion, refs. 32 and 33.

^{||}Ref. 34 for evidence on the effectiveness of narrative models during conflict.

^{**}Indeed, 2 mo after the end of this study's sample period, Ajax Systems introduced a postalert mobile push notification promoting fundraising efforts for the Ukrainian armed forces (41).

^{††}This is corroborated by qualitative reports from Ukraine, e.g., ref. 42.

content, have dynamic effects on responsiveness over time.^{‡‡} We anticipate that this is an area of significant policy interest that future research might be able to address.

Discussion

We provide credible quantitative evidence of the effectiveness of messaging about imminent risk during conflict. Combining granular data on population movement and high-frequency, localized alerts about imminent threats, we find that civilians respond sharply to these alerts, moving significantly in the minutes following notifications. These rapid shifts in mobility decline as the war continues, consistent with public fatigue. We corroborate these findings by leveraging devices' elevation, speed, and estimated home locations, which indicate that the public did not adjust to a sharp attenuation in observable movement by seeking underground shelter closer by or inside their own homes. Attenuation also could not be explained by calibration to the quality of information. This attenuation in movement, however, is marginally greater among individuals exposed to longer alerts that kept the Ukrainian public in affected areas under protracted states of alarm. Our counterfactual exercises suggest that early warning alerts significantly reduced the overall number of potential civilian casualties (by 45%), though 8 to 15% of observed civilian casualties could have been avoided if public responsiveness remained as high throughout the conflict as it was during the initial phase. Understanding why the public's engagement changed so quickly during a high-stakes conflict has significant implications for future policy interventions as well as prior scholarship.

Materials and Methods

Mobile Device Movement. Anonymized device-level location data are obtained from location data provider Veraset.^{§§} The data consists of "pings," which are timestamped GPS locations shared by the device with a mobile app. Veraset aggregates and cleans such data, obtained from thousands of so-called "Software-Development Kits" (SDK), which are packages of tools that provide the infrastructure for many mobile applications. Location data from the same device but different SDKs can be combined by relying on the anonymized device ID, which is a unique string associated with a particular mobile device and can be changed only through a factory reset.

As a result, the data provide insights into the movement of a substantial share of Ukrainian mobile devices: After cleaning the data and filtering it on the period and regions of interest, we obtain around 500 million unique pings, corresponding to around 17 million unique devices. With a population of 44 million and smartphone penetration of 63%, this corresponds to around 60% of Ukrainian mobile devices.

The cleaning and filtering steps applied are as follows. First, we restrict the sample to pings observed between January 1 and September 30, 2022. To improve data quality, we remove "jumpy" pings, which result from distortions in the GPS signal, by filtering out pings where the device moved from one location to the next at a speed faster than 300 km/h. We also remove pings with a horizontal accuracy (the radius of the margin of error of the device's location) of more than 150 m. Finally, we retain pings in only those regions for which we have air raids data (Below for more details). The sample scale of devices remains consistent with changes in the relocation patterns of Ukrainian citizens during the sample period as reported by the United Nations High Commissioner for Refugees (UNHCR). We find similar distances traveled per device by day throughout the sample period, with a temporary increase during August 2022.

^{‡‡}Ref. 37 find that there is no crowding out when multiple nudges are combined at a given point in time. But it is not clear whether this extrapolates to multiple nudges spread out over time.

^{§§}<https://www.veraset.com/>.

Air Raid Alerts. Data on the start and end times of region-wide air raid alerts were scraped from the Telegram channel of the popular mobile app Air Alert Ukraine by Volodymyr Agafonkin^{¶¶} and published online. The Air Alert app was developed by the Ukrainian software development company stfalcon^{##} and Ukrainian security company Ajax Systems (<https://ajax.systems/>) with support from the Ukrainian Ministry of Digital Transformation. Users can select regions of interest and receive loud alert warnings informing them of the beginning and end of a civil defense alert in the region. Importantly, the app does not collect geolocation data, which means that activation of the alert should not bias the geolocation signals received from the Veraset data. Moreover, it is possible to activate alerts for only one region at a time, so we can reasonably expect users in Ukraine to tune in to the alerts for their region of residence. The app creators claim that the app is the only one of its type that supports critical alerts, where notifications are delivered even when a smartphone is in silent or sleep mode. The type of alert the app delivers falls under civil defense alerts and includes airstrike warnings, chemical attacks, impending technological catastrophes, etc. We retain only alerts classified as airstrike warnings, which results in 3,256 unique alerts over the sample period, which runs from March 15, 2022, to September 31, 2022 [Though air raids started occurring at the beginning of the war (February 24), reliable digital alert data only started coming in after March 15]. Between February and September 2022, the app had been downloaded around 5.3 million times in Ukraine, with over half of those downloads occurring in March 2022 (*SI Appendix, Fig. SI-3*). Not all regions have scraped alert data available: Fig. 7 shows the regions for which there are air alert data.

Violent Event Data. We leverage violent event data from two sources. The primary source is the Violent Incident Information from News Articles (VIINA) (43). The platform tracks violence in Ukraine using a multitude of source streams, including domestic sources in Ukraine and Russia, as well as open-source intelligence reports gathered from social media. The balance of source locations and types is used to establish cross-source coverage, addressing potential sources of reporting bias. Our measures of combat activity and events involving civilian casualties are drawn from the platform's recurrent neural network (RNN) classification of events. To calculate the damage multiplier used in our excess deaths exercise, we rely on information from the Armed Conflict Location and Event Data Project (ACLED), which combines supervised and unsupervised collection and processing techniques (44). Although this source lacks the source coverage and cross-validation of VIINA, ACLED includes information on the estimated number of casualties associated with each recorded event.

Territorial Data. Data on territorial control by Russian and Ukrainian forces are obtained from the VIINA database (43). We back out daily estimates of the front line from these data by buffering each locality with a 5-km radius, dissolving these buffered points into two polygons—one for Russian- and one for Ukrainian-controlled territory—and backing out the estimated front line territory as the area where both polygons intersect.

Geospatial Data. We match the alert regions from the Air Alert app to official administrative regions in Ukraine from 2015 at the oblast (region), raion (district), and city levels (45), the result of which is shown in Fig. 7. Similarly, we intersect the mobile device pings with these administrative regions to deduce which alert a device is exposed to.

Research Design. To assess civilian response to air raid alerts, we estimate the following event study specification for a window of 30 min around each alert,

$$Response_{it} = \sum_{t=t_0-10}^{t_0+30} \delta_t + \varepsilon_{it} \quad [1]$$

where i indicates a unique mobile device; t is a minute of the hour (e.g., 5:00pm to 5:01pm); $Response_{it}$ is a measure of device i 's movement in minute t , where

^{¶¶}<https://agafonkin.com/>.

^{##}<https://stfalcon.com/>.




Fig. 7. Air Alerts by Day and Region. *Note:* The figure depicts the number of distinct air alerts that were broadcast on the Air Alert Ukraine app in a given region on a given day. An alert that spans multiple days is counted on each of the days.

our main measure is the total distance the device moved between subsequent pings; δ_t is a dummy variable for being in minute t ; the sum iterates over all such dummies from 10 min before to 30 min after the alert; and, ε_{it} is an error term. The choice of time window corresponds to the minimum alert duration in the sample, which is 30 min (The longest alert duration is 183 d in Luhansk. The median is 59.78 min and the 99th percentile 279.53 min). The end of the alarm is followed by an additional “all-clear” notification. Results for extended time windows are shown in *SI Appendix, Fig. SI-8*. We allow for the panel to be unbalanced (i.e., not every device has a ping in every minute in the window) but require there to be at least one ping in each minute, or else we drop the corresponding alert. This results in a total of 3,256 estimated versions of Eq. 1. To illustrate the results of our event studies in a digestible format, we plot the central tendency of the 10 min before the alert, the minute of the alert notification, and the subsequent 30 min. We also plot the variability of the underlying estimates with 95% confidence

intervals constructed using a local polynomial regression (loess) across the event study bins (i.e., across time-to-treatment bins). We supplement this approach in Supporting Information with various clustered bootstrapping techniques to account for uncertainty in the underlying estimates themselves (*SI Appendix, Fig. SI-5*). Our setting can be viewed as a meta-analysis with a fixed research design conducted by a single research team. The clustered bootstrap has been shown to deliver valid confidence bounds compared to robust variance estimation approaches for meta-analysis when there is dependency across studies (46). Additional details about split-period and split-sample estimates are also provided in *SI Appendix*.

Data, Materials, and Software Availability. Data and code (with exception of Veraset geolocation data) have been deposited in Harvard Dataverse

(<https://doi.org/10.7910/DVN/YH1GJT>) (47) and Github (https://github.com/Davidvandijcke/ukraine_air_raids) (48). The geolocation data is proprietary and owned by Veraset. Interested researchers can review documentation at this link (<https://bit.ly/3lxyOE1>) and contact Chris Youngblade at Veraset regarding this context or similar movement data.

1. E. Jones, The psychology of protecting the UK public against external threat: COVID-19 and the blitz compared. *Lancet Psychiatry* **7**, 991–996 (2020).
2. J. J. V. Bavel *et al.*, Using social and behavioural science to support COVID-19 pandemic response. *Nat. Hum. Behav.* **4**, 460–471 (2020).
3. Ajax Systems, Jedi update of the air alert app (2022). <https://ajax.systems/blog/zastosunok-povitryana-trivogal/>.
4. K. R. Camarena, S. Claudy, J. Wang, A. L. Wright, Political and environmental risks influence migration and human smuggling across the Mediterranean sea. *PLoS One* **15**, e0236646 (2020).
5. K. R. Camarena, N. Hägerdal, When do displaced persons return? Postwar migration among Christians in mount Lebanon. *Am. J. Polit. Sci.* **64**, 223–239 (2020).
6. S. V. Rozo, J. F. Vargas, Brothers or invaders? How crisis-driven migrants shape voting behavior. *J. Dev. Econ.* **150**, 102636 (2021).
7. M. T. Fullilove, Psychiatric implications of displacement: Contributions from the psychology of place. *Am. J. Psychiatry* **153**, 1516–1523 (1996).
8. X. H. Tai, S. Mehra, J. E. Blumenstock, Mobile phone data reveal the effects of violence on internal displacement in Afghanistan. *Nat. Hum. Behav.* **6**, 624–634 (2022).
9. K. Sonin, A. L. Wright, Information operations increase civilian security cooperation. *Econ. J.* **132**, 1179–1199 (2022).
10. R. Basher, Global early warning systems for natural hazards: systematic and people-centred. *Philos. Trans. R. Soc. A: Math. Phys. Eng. Sci.* **364**, 2167–2182 (2006).
11. K. Nakayachi, J. S. Becker, S. H. Potter, M. Dixon, Residents' reactions to earthquake early warnings in Japan. *Risk Anal.* **39**, 1723–1740 (2019).
12. R. Bossu, F. Finazzi, R. Steed, L. Fallou, I. Bondár, "shaking in 5 seconds!"-performance and user appreciation assessment of the earthquake network smartphone-based public earthquake early warning system. *Seismol. Res. Lett.* **93**, 137–148 (2022).
13. N. E. Alias *et al.*, Community responses on effective flood dissemination warnings-a case study of the december 2014 kelantan flood, Malaysia. *J. Flood Risk Manag.* **13**, e12552 (2020).
14. M. K. Lindell, S. Arlikatti, S. K. Huang, Immediate behavioral response to the June 17, 2013 flash floods in Uttarakhand, North India. *Int. J. Disaster Risk Reduct.* **34**, 129–146 (2019).
15. E. Jones, R. Woolven, B. Durodié, S. Wessely, Civilian morale during the second world war: Responses to air raids re-examined. *Soc. Hist. Med.* **17**, 463–479 (2004).
16. M. L. Tan *et al.*, Understanding the social aspects of earthquake early warning: A literature review. *Front. Commun.* **178**, 8 (2022).
17. E. S. Cochran, A. L. Husker, How low should we go when warning for earthquakes? *Science* **366**, 957–958 (2019).
18. H. Bean *et al.*, The study of mobile public warning messages: A research review and agenda. *Rev. Commun.* **15**, 60–80 (2015).
19. J. M. Gutteling, T. Terpstra, J. H. Kerstholt, Citizens' adaptive or avoiding behavioral response to an emergency message on their mobile phone. *J. Risk Res.* **21**, 1579–1591 (2018).
20. P. Weyrich, A. Scolobig, F. Walther, A. Patt, Do intentions indicate actual behaviour? A comparison between scenario-based experiments and real-time observations of warning response. *J. Contingencies Crisis Manage.* **28**, 240–250 (2020).
21. S. Benartzi *et al.*, Should governments invest more in nudging? *Psychol. Sci.* **28**, 1041–1055 (2017).
22. A. Banerjee *et al.*, "Messages on COVID-19 prevention in India increased symptoms reporting and adherence to preventive behaviors among 25 million recipients with similar effects on non-recipient members of their communities" (Tech. Rep., National Bureau of Economic Research, 2020).

ACKNOWLEDGMENTS. We thank Luda Andriyevska, Christopher Blair, Jane Esberg, Patrick Francois, Scott Gehlbach, Guy Grossman, Florian Gunsilius, Harry Kleyer, Dariia Mykhailishyna, Juan Felipe Riano, Raul Sanchez de la Sierra, Nicholas Sambanis, Jacob Shapiro, Konstantin Sonin, Roya Talibova, Erik Wibbels, and Rebecca Wolfe for feedback. All errors remain our own.

23. M. Hallsworth, J. A. List, R. D. Metcalfe, I. Vlaev, The behavioralist as tax collector: Using natural field experiments to enhance tax compliance. *J. Public Econ.* **148**, 14–31 (2017).
24. T. Dunning, *et al.*, Voter information campaigns and political accountability: Cumulative findings from a preregistered meta-analysis of coordinated trials. *Sci. Adv.* **5**, eaaw2612 (2019).
25. H. Schildberg-Hörisch, Are risk preferences stable? *J. Econ. Perspect.* **32**, 135–54 (2018).
26. M. Callen, M. Isaqzadeh, J. D. Long, C. Sprenger, Violence and risk preference: Experimental evidence from Afghanistan. *Am. Econ. Rev.* **104**, 123–48 (2014).
27. M. Bzikadze, The rule of "two walls": Which place in the apartment is the safest (2022).
28. D. Lapin, O. Voityvych, T. Lister, H. Ritchie, M. Knight, Russian missile strike on apartment building in Dnipro kills 40 people, Ukraine officials say (2023).
29. A. Shepeleva, Ukrainians lack access to bomb shelters (2022).
30. E. Kamenica, M. Gentzkow, Bayesian persuasion. *Am. Econ. Rev.* **101**, 2590–2615 (2011).
31. J. Schwartzstein, A. Sunderam, Using models to persuade. *Am. Econ. Rev.* **111**, 276–323 (2021).
32. M. Gentzkow, E. Kamenica, Costly persuasion. *Am. Econ. Rev.* **104**, 457–62 (2014).
33. M. Gentzkow, E. Kamenica, Competition in persuasion. *Rev. Econ. Stud.* **84**, 300–322 (2016).
34. A. Armand, P. Atwell, J. F. Gomes, The reach of radio: Ending civil conflict through rebel demobilization. *Am. Econ. Rev.* **110**, 1395–1429 (2020).
35. S. DellaVigna, M. Gentzkow, Persuasion: Empirical evidence. *Annu. Rev. Econ.* **2**, 643–669 (2010).
36. I. Ayres, S. Raseman, A. Shih, Evidence from two large field experiments that peer comparison feedback can reduce residential energy usage. *J. Law Econ. Organ.* **29**, 992–1022 (2013).
37. A. Brandon, J. A. List, R. D. Metcalfe, M. K. Price, F. Rundhammer, Testing for crowd out in social nudges: Evidence from a natural field experiment in the market for electricity. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 5293–5298 (2019).
38. M. Motta, S. Sylvester, T. Callaghan, K. Lunz-Trujillo, Encouraging COVID-19 vaccine uptake through effective health communication. *Front. Political Sci.* **3** (2021).
39. P. Argote Tironi *et al.*, Messages that increase COVID-19 vaccine acceptance: Evidence from online experiments in six Latin American countries. *PLoS One* **16**, e0259059 (2021).
40. N. Legate *et al.*, A global experiment on motivating social distancing during the COVID-19 pandemic. *Proc. Natl. Acad. Sci. U.S.A.* **119** (2022).
41. Ajax Systems, The air alert app is now available for Android and iOS: Ajax systems blog (2023). <https://ajax.systems/blog/air-alert-english/>.
42. L. Deresh, The piercing sound of Ukraine's new reality (2022).
43. Y. M. Zhukov, Viina: Violent incident information from news articles on the 2022 Russian invasion of Ukraine. Ann Arbor: University of Michigan, Center for Political Studies (2022). <https://github.com/zhukovyuri/VIINA>. Accessed 24 October 2022.
44. C. Raleigh, A. Linke, H. Hegre, J. Karlsen, Introducing ACLED: An armed conflict location and event dataset: Special data feature. *J. Peace Res.* **47**, 651–660 (2010).
45. R. Hijmans, N. Garcia, J. Weiszorek, GADM: Database of global administrative areas, version 3.6. GADM Maps Data (2018).
46. M. Joshi, J. E. Pustejovsky, S. N. Beretvas, Cluster wild bootstrapping to handle dependent effect sizes in meta-analysis with a small number of studies. *Res. Synth. Methods* **13**, 457–477 (2022).
47. D. Van Dijke, A. L. Wright, M. Polyak, Replication data for: Public response to government alerts saves lives during Russian invasion of Ukraine. Harvard Dataverse. <https://doi.org/10.7910/DVN/YH1GJT>. Deposited 11 April 2023.
48. D. Van Dijke, A. L. Wright, M. Polyak, Davidvandijcke/ukraine_air_raids. Github. https://github.com/Davidvandijcke/ukraine_air_raids. Deposited 11 April 2023.