Drone Wars: The Terrorists Strike Back

Cyrus Mohammadian[*], Fouad Pervez[**] & Yu-Ming Liou[**]

[*]School of International Relations, University of Southern California
[**]Department of Government, Georgetown University

March 2016
ISA 2016 Annual Convention
Atlanta, Georgia

Abstract

As the use of drones has increased dramatically over the past decade, so too has the debate over its effectiveness as a counterinsurgency strategy. Some argue that drones are efficient ways to target terrorists with minimal loss of life. Others claim that anger over civilian casualties from drone strikes, coupled with concerns over violations of state sovereignty, serve to radicalize populations and motivate those targeted to retaliate in kind. One perspective predicts violence to decrease in response to drone strikes, while the other expects it to increase. We test these two propositions using a newly released geocoded dataset of drone strikes in Pakistan. The results of our analysis suggests that drone strikes actually significantly increase the probability of localized terrorist related violence in Pakistan at the substate level, particularly soon after the strike takes place.

Keywords: UAVs; Drone Strikes; Counterterrorism; National Security; Violence; International Conflict; Security

1 Introduction

There is much contention over the effectiveness of the use of Unmanned Aerial Vehicles (UAVs), commonly referred to as drones, in conflict. Drone strikes, often carried out by operators far away from the zone of impact, have been increasingly used, particularly by the United States in Pakistan and Yemen, to target suspected terrorists. The assumption is that these strikes, by eliminating terrorist leaders and serving as a deterrent to those not hit, will reduce terrorist activity. UAVs are expensive, but carry minimal human costs, at least from the perspective of the targeting country, as no soldiers are put at risk in potentially dangerous operations. However, these strikes have resulted in civilian casualties[1], and while the number of civilians killed or injured by UAVs is debatable, there is little doubt that considerable collateral damage has been incurred. This, along with questions of legality and sovereignty, have fueled tremendous public outrage at the US drone policy. Many have even called into question the policy's effectiveness as a counterinsurgency tactic. This leads to an obvious question: do drone strikes increase or decrease violence?

However, testing whether drone strikes mitigate or agitate militant violence is challenging. The main issue is the availability of data. Much of the information about drone strikes are classified - we don't know which group or individuals are being targeted. We only know if a strike occurred, and, in some cases, an accurate number of victims, including civilians killed. This presents a serious problem: if we are unaware which group is being targeted, it becomes difficult to know which group's responses are retaliation for the strike, or, conversely, which group's lack of activity after a strike can be attributed to it. Beyond that, there is another problem. Even if we knew which groups were targeted by strikes, there is a high likelihood that other groups will use strikes as motivation to commit acts of political violence for their own purposes. There is a competition for recruits, and it is not necessary that the group

[1]The number of civilian casualties in drone strikes is a very contested topic. Two of the most common sources, the Bureau of Investigative Journalism, and the New America Foundation, have very different numbers. The Obama administration has argued in the past that practically no civilians have been harmed, though this is based largely on an extremely liberal classification of what constitutes a suspected terrorist.

targeted will be the one using the strike in a retaliatory act. As such, it becomes extremely difficult to link acts of political violence with drone strikes with the kind of precision needed to properly evaluate whether strikes themselves are effective or not.

For this reason, we instead argue that it is more effective to examine if there is a change in militant activity in a particular region following a drone strike. This is not quite the same as examining whether the strike was effective or not, but rather examining the heterogeneity in the response to the strike in a particular area. We test competing hypotheses on drone effectiveness, namely an attrition hypothesis that suggests that drone strikes will lead to a decrease in militant activity, and an agitation hypothesis that suggests that drone strikes trigger reactions among militant groups and lead to an increase in political violence. Using geocoded data, we find evidence for the latter, as there is an increase in terrorist acts in the region a drone strike occurred. We believe this suggests that militant groups have political incentives to commit acts of terrorism, as a signal to both international and domestic actors, as well as a consequence of inter-group and intra-group power vacuums resulting from the drone strike. The rise in terrorist activity occurs soon after the strike, and is followed by no clear relationship in future time periods, suggesting that the window for this signaling is maximized near the time of the strike itself.

The rest of this paper is as follows. We first discuss relevant literature on drone strikes, targeted assassinations, and signaling. We then lay out our theory. Following that, we describe our methodology and present our findings. Finally, we conclude with our thoughts on what our results mean for future research on drone strikes, as well as the policy implications of our findings.

2 Literature Review

The use of drones or UAVs to conduct targeted killings by the United States in its prosecution of the Global War on Terror (GWOT) has expanded dramatically over the past

decade. Analysts and policymakers have extolled the virtues of drone strikes in disrupting terrorist networks without putting American lives at risk (see *e.g.*, Byman 2013). Similarly, public opinion research has shown that the American public is highly supportive of the use of drones in the GWOT (Kreps and Wallace 2015). Nonetheless, the empirical literature on the efficacy of drone strikes is quite limited and what studies do exist have thus far produced contradictory results.

On the one hand, Johnston and Sarbahi (2015) find that US drone strikes in the frontier regions of Pakistan lead to a decline in violence in surrounding areas and more specifically, a decline in selective targeting of tribal elders by terrorists. Likewise, Byman (2013) finds that Hamas attacks against Israel lose efficacy following drone strikes against terrorist leaders, arguing that this is a result of the removal of the most effective and competent members of terrorist groups. Similarly, Jordan (2014) argues, based on a case study of the US drone campaign against Al Qaeda Central, that drone strikes work to disrupt hierarchically organized groups like Al Qaeda by destroying many of its most competent operatives and destroying key material resources.

Conversely, Boyle (2013) theorizes that assessments of the efficacy of drone strikes focus too narrowly on successes in killing individual terrorists and rarely take full account of the broader costs of drone campaigns, which may be more difficult to measure, but can be both substantial and long-lasting. These include, he argues, weakening of local governments (*e.g.,* in Pakistan or Yemen) due to perceptions of incompetence, weakness. and illegitimacy following US drone strikes, as well as growing anti-American sentiment, and increased militancy in areas in which drone strikes have been concentrated. Empirically, drone skeptics have pointed to the apparent ineffectiveness of the US drone campaign on Al Qaeda's propaganda output[2] (Smith and Walsh, 2013).

Finally, work by Jaeger and Siddique (2011) finds mixed evidence of drone campaign

[2]Smith and Walsh (2013) argue that propaganda may be a more appropriate measure of terrorist "productivity" because of issues with data collection in event-based terrorism databases, difficulties with choosing the appropriate time window for coding terrorist "responses" to drone strikes, and open-source data availability on terrorist activities.

efficacy in a study of the effect of strikes on Al Qaeda and Taliban violence in Afghanistan and Pakistan from 2007-2010. The lack of consensus in this literature echoes similar mixed findings in a closely related literature examining the effect of targeted assassinations (sometimes called "leadership decapitation") on terrorist and violent criminal groups (*e.g.,* drug cartels). For example, while some authors report evidence that targeted killings are effective at reducing terrorist attacks (Morag 2005; Jaeger and Paserman 2009; Neumann, Evans, and Pantucci 2011; Price 2012; Johnston 2012), others find that leadership decapitation has either mixed or conditional effects on terrorist groups, depending on their structure (Mannes 2008; Jordan 2009, 2010), that it results in an increase in attack volume, but a decrease in average attack efficacy (Wilner 2010), has no statistically significant effect (Hafez and Hatfield; Hepworth 2014), or can in fact, lead to increases in terrorist attacks (Kaplan Mintz & Mishal 2006; Kaplan Mintz, Mishal & Samban 2005).

Theoretically, drone optimists have tended to emphasize two causal mechanisms in arguing for the efficacy of drone campaigns by the US and Israel. First, attrition of capabilities: whether through sheer volume of strikes or through the removal of key leaders and technicians (see, *e.g.*, Jaeger and Paserman, Price 2012; Johnston 2012; Johnston and Sarbahi 2015). Relatedly, other authors have suggested that this mechanism may depend upon the command or organizational structure of terrorist groups (Mannes 2008; Jordan 2009. 2010). Second, deterrence of future attacks that might draw drone strikes in response (Jaeger and Siddique 2011). Conversely, drone skeptics have suggested that post-strike attacks may increase as terrorist groups seek vengeance (Kaplan et al. 2005, 2006; Jaeger and Siddique 2011) or as a result of increasing terrorist recruitment stemming from drone-strike related anger among civilians (Boyle 2013).

However, a related literature offers a further potential mechanism: signaling strength and commitment by terrorist groups in the wake of a drone strike. An extensive theoretical literature on terrorism and signaling argues that terrorists use attacks to signal both capabilities and resolve both to governments and potential recruits and supporters (Overgaard

1994; Hoffman and McCormick 2004; Kydd and Walter 2006). An extension of this logic to post-drone strike terrorist behavior is straightforward: groups that have been targeted by drone strikes are incentivized to signal their undiminished strength and resolve to both foreign and domestic enemies (US and local governments), competing terrorist groups, and local civilians, upon whom they depend for resources, recruits, intelligence, and compliance. Likewise, even groups not targeted by drone strikes may seize upon the opportunity presented by a drone strike on a competing group to seek additional "market share" by demonstrating that they are both unhurt and undeterred through increased violent activity.

3 Theory and Hypotheses

3.1 Theoretical Framework

We test two competing theories regarding the effectiveness of drone strikes. The underlying justification for the US drone policy is that they reduce terrorism through a battle of attrition. As militants are successfully targeted, we should see a decline in terrorism. This is consistent with the literature that focuses on leadership decapitation and how it reduces groups' opportunities and capabilities to engage in violence. If this theory is correct, our empirical tests should reveal a decrease in the likelihood of terrorist violence in the vicinity of the targeted regions. We label this a theory of attrition.

Alternatively, an agitation theory holds that the drone policy does not come without its costs and unforeseen consequences, primary of which is heightened hostilities. There are several reasons: strikes may heighten general grievances that lead to retaliatory violence, they may push groups towards retaliation to signal resilience and strength, they may lead to greater group decentralization that results in more, not less, violence, or they may increase inter-group hostilities in the midst of a regional power vacuum. While it is difficult to differentiate between these mechanisms, our empirical tests should reveal an increase in the likelihood of terrorist violence in the vicinity of the targeted regions if any one of these

mechanisms hold.

Both theories assume the same trio of primary actors, the US military in Pakistan, militant groups targeted by drone strikes in the FATA/Frontier region, and rival militant groups in the same or neighboring region. They also share a pair of passive actors - civilians who are either targeted purposefully by militants or inadvertently by US drone strikes and Pakistani government interests that bear the brunt of militant reprisals (i.e. military bases, military personal, government infrastructure, local police, etc.). In both theories, the US government has a clear motive in launching drone strikes in Pakistan: undermining militant organizational capacity while incurring as few military casualties as possible.

The distribution of incentives for militant groups and their rivals is less clear, however. The theory of attrition assumes that militant groups successfully targeted by the strikes respond *not* by capitalizing on an opportunity to mobilize sympathizers against Pakistani interests, but by reducing their activities in the region either to avoid incurring further attention or as a consequence of reduced capacity to continue engagement. Similarly, rather than taking advantage of the resulting power vacuum, this narrative assumes rival militants will actually learn from the misfortune of their competitors and reduce their activities in response.

In contrast, the theory of agitation posits a much more delicate network of ties between rival groups and assumes a far more complex set of rules govern their interactions between one another and the Pakistani government. In this narrative, it is assumed rival groups are as concerned with their fight against US and Pakistani government interests as they are with the share of local power they command in the regions they reside. Their interests are to increase their share of local power. Drone strikes not only change the distribution of militant incentives in the regions they target, successful ones alter the distribution of capabilities as well. Changes in the distribution of capabilities and incentives increases hostilities between rival groups as well as all militant groups and US/Pakistani interests in a number of ways. As drone strikes increase, they reduce or eliminate the organizational capabilities of targeted

groups, thereby increasing the pool of militants viable for recruitment by rivals unscathed by the strikes. These strikes also increase the demand for retaliation from members of their support base (civilian or otherwise). When targeted groups cannot respond due to degradation in organizational capacity, the theory of attrition anticipates rivals groups, who share similar goals and ideology, to meet this demand by taking up the retaliatory mantle as a means of attracting new recruits at the expense of targeted groups. Thus, as the demand for retaliatory actions increases and the supply of militants available to recruit increases, the incentives to engage in violence should also increase (i.e. the costs of engaging in violence are reduced because its payoffs are higher). We label this supply and demand of militant violence the "marketplace of militant extremism". Finally, a successful strike might eliminate leaders who would be more strategic about political violence and hold back more zealous group members. Thus, where the theory of attrition assumes losses in organizational capacity reduce violence, the theory of agitation allows for the possibility that groups whose centralized command structures are degraded by successful decapitation strikes, have fewer constrains on the uncoordinated violent actions of lone wolves. An animal is most dangerous when it is wounded.

3.2 Anticipated Effects

Militant groups targeted by drone strikes are in a precarious position. On the one hand, a successful strike most certainly weakens their organizational capacity to some degree. According to the attrition narrative, these organizations have lost critical leaders and other important human and/or material capital in the strike, which results in the degradation of their capacity to engage in subsequent violence. Even unsuccessful strikes may depress militancy; out of fear of future reprisals, targeted groups may be less willing to engage in violence. If the theory of attrition holds, These groups might also be influenced by drone strikes in neighboring regions. A "neighborhood effect" might exist for two reasons: one, militant groups' zones of engagement may overlap multiple districts,

and, two, the close network of ties between rival militant groups also transcends district boundaries, which may result in coordinated actions or learning between different groups in close proximity. Therefore, if drone strikes damage a group's capabilities or reduce their incentives to engage in violence, these effect may be observable in neighboring regions as well.

H1$_a$: *Drone strikes decrease the likelihood of terrorist violence in the districts they strike.*

H1$_b$: *Drone strikes decrease the likelihood of terrorist violence in neighboring districts.*

On the other hand, drone strikes, successful or otherwise, may also provide militants with countervailing incentives to engage in *more*, not less, violence. According to the theory of agitation, these groups have political incentives to signal their resilience and strength to their American and Pakistani government adversaries, as well as other militant groups that they compete with over local recruits. In this narrative, being weakened, or perceived as weakened, by a drone strike renders a targeted group vulnerable to losing recruits to their competitors. Showing strength through retaliatory attacks following drone strikes is one way for militant groups to signal that they are still in the game. Groups hit by drone strikes might focus their retaliatory efforts on other regional actors as well. For example, they could assume that the intelligence to strike at them came from rivals in nearby villages, towns, or cities. Moreover, successful leadership decapitations may also result in power vacuums that promote intra-group violence in the marketplace of militant extremism. Furthermore, if a strike inflicts real leadership damage, it can result in a more decentralized and undisciplined command structure. This could lead to problems of group coordination and may result in less severe but more frequent attacks. Finally, this signal of strength is at least as much directed towards rival groups as it is toward the US and Pakistani government interests in the region, primarily as a show of strength but perhaps also as an expression of their grievances. All of these mechanisms predict increases in the likelihood of terrorism following drone strikes.

Thus, if the agitation hypothesis is correct, we see multiple possible mechanisms that

result in a rise in terrorism following a drone strike. In one, a group hit by a drone strike commits terrorism after the strike to signal its strength, to address grievances, or as a result of greater group decentralization or intra-group competition. In the other, groups not hit by drone strikes commit terrorism to gain recruits and resources from the robust marketplace for militant extremism, possibly in an effort to fill a regional power vacuum if the targeted group was actually weakened by the strike. Given the limitations of the data on drone strikes, particularly due to the classified nature of much of the program, it is impossible to differentiate these two mechanisms. However, they both lead to the same result: political violence should increase following a drone strike, as all region-specific domestic actors have political incentives to act. As noted earlier, there may be a neighborhood effect. In this instance, if a drone strike does not debilitate a group, or creates a power vacuum in the general region, this could lead to a positive neighborhood effect, whereby a drone strike in one district leads to an increase in terrorism in a neighboring district.

$H2_a$: *Drone strikes increase the likelihood of terrorist violence in the districts they strike.*

$H2_b$: *Drone strikes increase the likelihood of terrorist violence in neighboring districts.*

Thus far, our hypotheses have focused on the direction of the effect (terrorism) and not on its timing. Here too, we adopt an agnostic view. If drone strikes do degrade group capacity in a way that results in observable decreases in violence, then violence very well may decrease immediately as groups scramble to reorganize. Whether there are any lingering effects is another question. If the command structure of these groups is fluid enough (i.e. replaceable), they could conceivably recoup their capacity quick enough such that any depression in the level of violence following the strikes would be fleeting at best. Similarly, if drone strikes are shown to increase the probability of terrorist related violence the logic could support both immediate and latent effects. On the one hand, if the drone strikes result in heightened grievances that then manifest in increased levels of

violence, there is no reason to suspect the effect of those grievances should wane over a short period of time. Therefore, we should observe both immediate and latent effects. On the other hand, if groups are retaliating as a means of signaling their strength or in a window of time in which they have distinct political incentives to act, we should expect an immediate surge in violence following attacks without an accompanying latent effect.

H3$_a$ *The impact of drone strikes on the likelihood of terrorist violence should be observable in the time period immediately following the strike.*

H3$_b$ *There is a time lag in the observable impact of drone strikes on the likelihood of terrorist violence.*

4 Research Design and Data

4.1 Methodological Approach

Five distinct but related methodological challenges we face guide our empirical strategy. First, we must choose a geographic unit of analysis whose scale is most appropriate for modeling the hypothesized spatial relationship between drone strikes and terror. In the spatial sciences literature this phenomenon is known as the(modifiable areal unit problem (MAUP) (Iqbal and Starr 2014). Here it becomes useful to consider the options ???-Pakistan???s Federally Administered Tribal Areas (FATA), like the rest of the country, is divided into districts (seven tribal agencies and six frontier regions), which are furthered divided into forty-one tehsils[3] and six frontier regions (the boundaries of the frontier regions are consistent across both divisions). The centroid to furthest distance of the smallest tehsil in the region is less than 3 km long, while the centroid to furthest distance of the largest tehsil is about 50 km, with an average of about 15 km. In contrast, the largest district in the Fata region has a

[3]The tehsil (known as "taluko" in Sindh province is the second of Pakistan's three administrative levels–between the District and the "Union council".

centroid to furthest distance of about 170 km, its smallest is about 25 km, and the average is about 60 km.

We argue that the substantially shorter radii characteristic of the tehsils problematize inference because, although we expect terrorists groups to retaliate when attacked at their villages, we do not anticipate this blowback to occur in regions that are close to the residential dwellings where these individuals are targeted. Put more simply, we do not expect revenge seeking terrorists to retaliate by bombing their own villages or targets too close to home. Therefore, we argue that to model the relationship between drone strikes and terrorist incidents, our spatial unit of analysis must be large enough to allow for a reasonable distance between targets and their anticipated zones of retaliation but small enough as to only capture incidents that can reasonably be tied to retaliatory actors. We feel safe in assuming that twice the radial distance of the district average, about 120 km, is a reasonable distance that individuals could be expected to travel within our temporal unit of aggregation, which we turn to next.[4]

The second methodological hurdle we face is the issue of endogeneity, or reverse causality. One could sensibly argue that a positive association, for example, between drone strikes and terrorist incidents could be as much the result of retaliation as it is the result of US forces simply striking where the heat is. In other words, terrorism and drone strikes might be associated with one another not because terrorists increase their engagement following strikes but because the US forces actually anticipate terrorism and, therefore, time their attacks during periods when terrorism tends to spike. We address this issue to a degree by

[4]One way to address the issue of the relationship between target and retaliation zones without adopting the larger district unit of analysis is to include neighboring drone strikes in the model. This way, if drone strikes do increase the likelihood its victims retaliate in tehsils away from where they are targeted, the coefficient in front of the variable that captures neighboring drone strikes could identify such an effect. We, however, contend that given the small scale and the massive range in the centroid to furthest distance across the different tehsils, the choice of tehsils as our spatial unit of analysis is still problematic. The small scale may require higher order spatial lags that may overcompensate in larger tehsils yet continue to be under-encompassing for the smaller ones. At any rate, our appendix includes a robustness check run using tehsils as our spatial unit. The results are in line with our stated expectations. We also present results of a model that includes spatial lags of drone strikes at the district unit of analysis and find no evidence to justify the use of an even larger observational unit.

selecting weeks as our temporal unit of analysis, which we argue provides our research design with a quasi-random treatment assignment. We suspect the decision to strike a target in one week as opposed to the preceding or proceeding week is dictated by a number of factors unrelated to the likelihood of terrorist attacks, such as weather patterns and other logistical issues.[5]

While the behavior of terrorist groups may be effected in one direction or another (or not at all) in the immediate aftermath of a drone strike, that does not preclude percolating effects that manifest in subsequent time periods, which brings us to our third methodological obstacle–temporal effects. Do terrorists groups respond immediately in order to demonstrate resolve and capacity? Or do drone strikes degrade their capacity to respond? Or is such a pacifying period simply downtime these organizations use to recoup, recover, and retaliate later? In order to address these questions, we use a series of higher order temporal lags to model the effect of drone strikes on of the likelihood of terrorism over time.

Fourth, our dependent variable, terrorism, demonstrates considerable spatial autocorrelation that must be accounted for. Specifically, we add a first order spatial lag of terrorism weighted by shared border length.[6] The spatial autocorrelation in the dependent variable is likely the result of at least three processes. One, terrorists target regions with similar attributes (and those attributes tend to cluster geographically) Second, terrorism in one district my actually increase the probability of terrorism in neighboring districts through a diffusion process. Third, and most theoretically interesting to us, is that drone strikes may actually increase terrorism in neighboring districts as well because the zone of retaliation may extend beyond the boundaries of the district targeted by drones. Although, we already discussed that our spatial unit of analysis is on average large enough that, assuming drone strikes do

[5]For example, in a letter dated November 24, 2010, an Al Qaeda commander reassures Osama Bin Laden that his wife won't travel until a "convenient cloudy day.??? (See http://www.mei.edu/content/news/fear-surveillance-haunted-bin-laden.) These considerations are echoed in a NASA statement on the problems of flying UAVs in certain frigid conditions. Available at: https://spinoff.nasa.gov/Spinoff2010/ps_2.html. Also see Johnston and Sarbahi (2015) for a similar application with drone strikes in Pakistan.

[6]Here we include incidents of terrorism in neighboring regions outside of FATA as well.

affect terrorist behavior, we could reasonably expect to observe coterminous changes in drone strikes and terrorism, it is not unreasonable to fathom that a district???s level of terrorism may be affected by drone strikes in neighboring districts. If we believe this to be true, then we must account for the effect of neighboring strikes. However, drone strikes are geographically clustered–strikes in one region tend to correlate with strikes in neighboring regions. As a result, including both sets of temporal lags may lead to problems of multicollinearity and could result in inflated standard errors and reduced significance levels. We believe that an indicator of terrorism in neighboring districts (spatial lag model) acts as an effective proxy for these strikes. Alternatively, we estimate a model that drops the spatial lag of the dependent variable in favor of a model that includes the spatial lag of the explanatory variable instead (spatial Durbin model).

Fifth and most importantly, because the data is organized hierarchically, our estimation technique must account for the nested structure of panel data. Fortunately, since we are only interested in the effect of drone strikes on terrorism, a variable that varies across time, we use unit fixed effects to eliminate all unit-level time-invariant unobserved heterogeneity driven by factors such as terrain, distance to capital and borders, level of development, and many other geographic and sociopolitical variables. We also suspect seasonal effects to influence the variation in levels of terrorism across time (changing seasons might affect dividends from the sale of opium by terrorist groups for example) and, therefore, include week fixed effects as well. Finally, we use robust standard errors and cluster them by district.

4.2 Estimation Technique

We employ two different models but our primary approach uses a binary spatiotemporal autoregressive (STAR) model with district and week fixed-effects.[7] In addition, we also estimate a spatial Durbin model with spatial lags of the explanatory variable (EV) –drone

[7]For more on the estimation of spatiotemporal binary outcome models in political science see Franzese, Hays, and Schaeffer (2010)

strikes replacing the spatial lags of the dependent variable –terrorism.[8] The STAR model can be written in matrix notation as:

$$y = \rho \mathbf{W} y + \phi \mathbf{V} y + \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\alpha} + \mathbf{r} + \boldsymbol{\epsilon} \tag{1}$$

where y, the binary dependent variable of a terrorist attack, is an $NT \times 1$ vector of cross sections by time periods. The parameter ρ and ϕ are the spatial and temporal autoregressive coefficients respectively, which makes $\mathbf{W}y$ (where $\mathbf{W}$ is an $NT \times NT$ spatial-weights matrix) the spatial lag of the dependent variable and $\mathbf{V}y$ (where V is an $NT \times NT$ matrix) its temporal lag. $\mathbf{X}$ is an $NT \times NT$ matrix of our lower order (*i.e.*, non-spatial) covariates (drone strikes in consecutive time periods) and the parameter $\boldsymbol{\beta}$ is a vector of its coefficients. Finally, $\boldsymbol{\alpha}$ is an $N \times 1$ vector of district fixed effects, $\mathbf{r}$ is a $T \times 1$ vector of weekly fixed effects, and $\boldsymbol{\epsilon}$ is the error term. Similarly, the Durbin model can be written as:

$$y = \mathbf{W}\mathbf{X}\theta + \mathbf{X}\boldsymbol{\beta} + \phi \mathbf{V} y + \boldsymbol{\alpha} + \mathbf{r} + \boldsymbol{\epsilon} \tag{2}$$

where the parameter $\boldsymbol{\theta}$ is the higher order (*i.e.*, spatial) coefficient of the covariates and $\mathbf{WX}$ (where $\mathbf{W}$ is an $NT \times NT$ spatial-weights matrix) is its spatial lag. Note in contrast to the STAR model, the Durbin model drops the spatial lag of the dependent variable, terrorism, in favor of a spatial lag of the explanatory variables, drone strikes.

4.3 Data and Variables

The data for our dependent variable (DV), terrorist attacks, is drawn from the Global Terrorism Database (GTD). Due to certain changes in the GTD's coding procedures over time, we restrict our analysis to the time period to 2005 and 2013.[9] The dataset geocodes

[8] We include district and week fixed-effects in this model. We have also estimated each of these models using tehsil-weeks as unit of analysis. These results are available upon request.

[9] Changes were made to the coding procedure in 2007 and again in 2012. We run our models with controls for these time periods (2005-2007, post 2007-2011, and 2012) and we also run a restricted model between 2007-2011 only; our results remain robust and are available upon request. For more on these coding inconsistencies see https://www.start.umd.edu/gtd/using-gtd/.

terrorist attacks to the nearest village, town, or city center. We aggregate attacks over each district for each week and convert these totals to a binary measure. The binary contiguity matrix W conditions the spatial lag of the model. In the STAR model, this results in a variable that records whether at least one terrorist event took place in any neighboring district.

We obtained data for our explanatory variable (EV), drone strikes, from the Bureau of Investigative Journalism, a new dataset that records the precise location of American drone strikes in Pakistan between 2004 and 2014.[10] To our knowledge, this is the first analysis of drone strikes and terrorism to use this recently released dataset. We subset the data to the years between 2005 and 2012 (see footnote 9 on page 16). We aggregate these over our unit of analysis, which results in a total count of drone strikes over each district in each week during the sample period. We also employ a fourth order temporal lag of our EV. In other words, we include a time lag of drone strikes for the four weeks prior to our observation, one for each week for the STAR model. For the Durbin model, we drop the spatial lag of the DV and, instead, include the spatial lag of drone strikes (our EV) for each of its four temporal lags. This spatiotemporal lag captures the aggregate number of drone strikes experienced by a district's neighbors across four consecutive weeks prior to each observation. Finally, we include district and weekly fixed effects to remove any unobserved unit and seasonal heterogeneity. Figure 1 shows the spatial distribution of terrorist attacks and drone strikes for the entirety of the sample.

5 Results

Table 1 shows the results of our regression analysis. We find no support for the theory of attrition but do find evidence for the theory of agitation. In each model specification, we

[10]There is contention over the counts of civilian casualties in the different drone strike datasets, but since we are not looking at that, this should not pose a problem. For purposes of robustness, we compared the Bureau of Investigative Journalism's drone data to the New America Foundation's drone data, another common source. We found over a 98% correlation between the dates in these two datasets.

Figure 1: Terrorist Attacks and Drone Strikes in Pakistan's FATAH Region

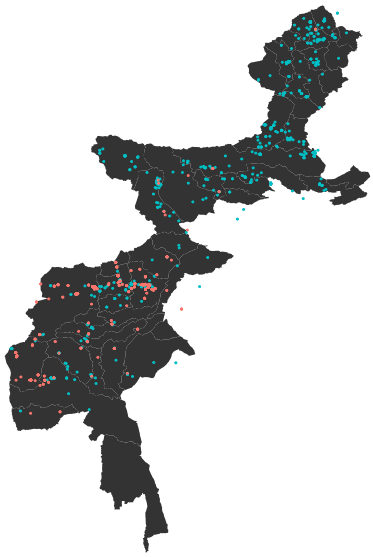

Note: Drone strikes are in red and terrorist attacks in green.

control for temporal and spatial autocorrelation, both of which appear significant regardless of model of specification. We conclude that that district level terrorism is influenced by a district's past levels of violence as well as the levels of violence experienced by neighboring districts.

The main variable of interest is the drone strikes, shown by the UAV variable. We are interested in examining whether spikes in terrorist attacks are preceeded by drone strikes. The attrition hypothesis holds that they should be negatively correlated. However, the evidence does not bear this theory out —there are no statistically significant negative effects of drone strikes on terrorist acts. The agitation hypothesis holds that drone strikes should be positively correlated to terrorist actions, particularly soon after the drone strike. We do find evidence of this, as the period t-2 is statistically significant and positively signed. This suggests that retaliatory violence takes between 8-14 days to emerge.

Figure 2 shows the substantive effect of a drone strike on acts of terrorism. It shows an increase in the predicted probability of militant attacks as a function of increasing drone

Figure 2: Impact of Drone Strikes on Acts of Terrorism, with 95% CIs

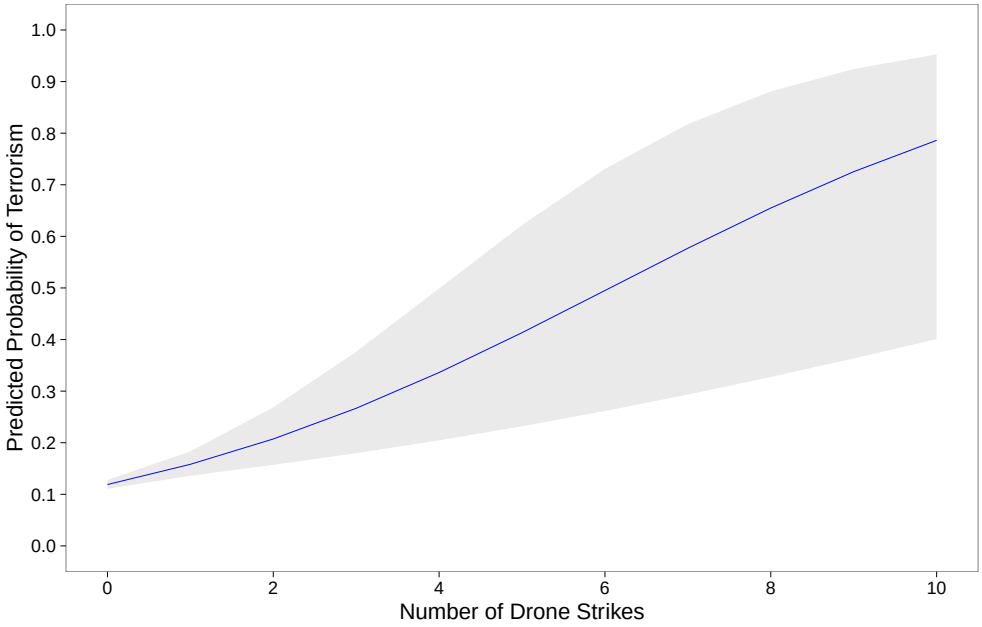

Note: this is for drone strikes occurring in period t-2, or 8-14 days prior to the act of terrorism. All other variables in the model are held at their means.

strikes in the preceding 8-14 days. This effect is quite large, as the predicted probability of a terrorist action increases from just above 10% to over 70% as the drone strikes increase. This highlights that our statistically significant finding for the agitation hypothesis is also quite substantively large.

We see no spatial effect in our results. None of the time lags for drone strikes in neighboring districts had a statistically significant relationship with acts of terrorism. This suggests that contiguity is not a factor - districts may simply be too large for there to be any spillover effects, either negative or positive, of drone strikes in neighboring districts on political violence.[11]

[11] We ran additional time lags (out to t-8), but none of those were significant, nor did they change our main finding. After examining AIC values, we found that the models fit best with four time lags, so for that reason, we only ran models out to t-4 for purposes of presentation in this paper.

Table 1: Regression Results

	Dependent variable:		
	Terrorist Incident		
	(STAR 1)	(STAR 2)	(Durbin)
$\text{Terrorism}_{t\text{-}1}$	0.967^{***}	0.967^{***}	0.970^{***}
	(0.089)	(0.089)	(0.089)
$\text{Neighborhood Terror}_{t\text{-}1}$	1.040^{***}	1.053^{***}	1.054^{***}
	(0.164)	(0.164)	(0.164)
$\text{UAV}_{t\text{-}1}$	0.015	-0.046	
	(0.077)	(0.087)	
$\text{UAV}_{t\text{-}2}$		0.185^{**}	
		(0.084)	
$\text{UAV}_{t\text{-}3}$		0.011	
		(0.085)	
$\text{UAV}_{t\text{-}4}$		0.054	
		(0.086)	
$\text{Neighborhood UAV}_{t\text{-}1}$			-0.147
			(0.197)
$\text{Neighborhood UAV}_{t\text{-}2}$			0.024
			(0.185)
$\text{Neighborhood UAV}_{t\text{-}3}$			-0.259
			(0.199)
$\text{Neighborhood UAV}_{t\text{-}4}$			0.171
			(0.179)
Constant	-1.453^{***}	-1.461^{***}	-1.445^{***}
	(0.319)	(0.320)	(0.319)
Observations	6,188	6,188	6,188
Log Likelihood	$-1{,}873.320$	$-1{,}869.726$	$-1{,}871.642$
AIC	3,882.640	3,881.453	3,885.284
Note:			*p<0.1; **p<0.05; ***p<0.01

6 Conclusion

Based on data on US drone strikes in Pakistan from 2005-2013, we find evidence for no systematic long-term effect of drone strikes on local terrorist activity, as well as evidence for a short-term increase in local terrorist attacks following US drone strikes. The implications for both theory and policy of these findings is clear: drone strikes do not appear to effectively reduce the incidence of terrorism in Pakistan, accounting for temporal and geographic effects. In fact. contrary to the expectations of the attrition mechanism proposed by drone optimists, violence not only does not decrease, but temporarily increases in response to drone strikes. This is particularly striking given that any likely selection bias in US policymakers' choice of targets should bias *against* these findings. If, as seems reasonable, US decision-makers are targeting the most violent terrorist groups and individuals for drone strikes, these strikes should be particularly effective at decreasing violence, rather than ineffective or counter-productive.

Instead, this suggests that drone strikes may be either provoking the most violent groups without meaningfully debilitating them or, by stoking civilian resentment and/or undermining local government legitimacy and trust, increasing support for militants in the local population. Both of these implications should be disturbing for US policymakers seeking to both disable existing militant groups and to "defund" the market for Islamic extremism. At the very least, policymakers and analysts should be careful of claims proclaiming the drone campaign to be an unqualified success in the Global War on Terror.

References

Boyle, Michael J. 2013. "The costs and consequences of drone warfare." *International Affairs* 89(1):1–29.

Byman, Daniel. 2013. "Why drones work." *Foreign Affairs* 92(4):32–43.

Franceze, Robert J., Jude Hays and Lena M. Schaffer. 2010. "Spatial, Temporal, and Spatiotemporal Autoregressive Probit Models of Binary Outcomes: Estimation, Interpretation, and Presentation". *APSA 2010 Annual Meeting Paper*.

Hafez, Mohammed M and Joseph M Hatfield. 2006."Do Targeted Assassinations Work? A Multivariate Analysis of Israel's Controversial Tactic during Al-Aqsa Uprising 1." *Studies in Conflict and Terrorism* 29(4):359–382.

Hepworth, Daniel P. 2014. "Terrorist retaliation? An analysis of terrorist attacks following the targeted killing of top-tier al Qaeda leadership." *Journal of Policing, Intelligence and Counter Terrorism* 9(1):1–18.

Iqbal, Zaryab and Harvey Starr. 2014. "Introduction: Spaces and Places: Geopolitics in an Era of Globalization." *International Studies Review* 17.1 (2015): 1-5.

Jacobson, Daniel and Edward H Kaplan. 2007. "Suicide bombings and targeted killings in (counter-) terror games." *Journal of Conflict Resolution* 51(5):772–792.

Jaeger, David A and M Daniele Paserman. 2007. "The Shape of Things to Come? Assessing the Effectiveness of Suicide Attacks and Targeted Killings.".

Johnston, Patrick B. 2012. "Does decapitation work? Assessing the effectiveness of leadership targeting in counterinsurgency campaigns." *International Security* 36(4):47–79.

Johnston, Patrick B and Anoop K Sarbahi. 2016. "The Impact of US Drone Strikes on Terrorism in Pakistan." *International Studies Quarterly* p. sqv004.

Jordan, Javier. 2014. "The Effectiveness of the Drone Campaign against Al Qaeda Central: A Case Study." *Journal of Strategic Studies* 37(1):4–29.

Jordan, Jenna. 2009. "When heads roll: Assessing the effectiveness of leadership decapitation." *Security Studies* 18(4):719–755.

Kaplan, Edward H, Alex Mintz and Shaul Mishal. 2006. "Tactical prevention of suicide bombings in Israel." *Interfaces* 36(6):553–561.

Kaplan, Edward H, Alex Mintz, Shaul Mishal and Claudio Samban. 2005. "What happened to suicide bombings in Israel? Insights from a terror stock model." *Studies in Conflict and Terrorism* 28(3):225–235.

Kreps, Sarah and Geoffrey Wallace. N.d. "The Marketplace of Ideas at War: Issues, Elites, and Public Support for American Drone Strikes." . Forthcoming.

Mannes, Aaron. 2008. "Testing the Snake Head Strategy: Does Killing or Capturing its Leaders Reduce a Terrorist Group's Activity?" *The Journal of International Policy Solutions* 9(40):40–49.

Morag, Nadav. 2005. "Measuring success in coping with terrorism: the Israeli case." *Studies in Conflict and Terrorism* 28(4):307–320.

Neumann, Peter, Ryan Evans and Raffaello Pantucci. 2011. "Locating Al Qaeda's center of gravity: The role of middle managers." *Studies in Conflict and Terrorism* 34(11):825–842.

Price, Bryan C. 2012. "Targeting top terrorists: How leadership decapitation contributes to counterterrorism." *International Security* 36(4):9–46.

Smith, Megan and James Igoe Walsh. 2013. "Do drone strikes degrade Al Qaeda? Evidence from propaganda output." *Terrorism and Political Violence* 25(2):311–327.

Walsh, James Igoe. 2015. "Precision weapons, civilian casualties, and support for the use of force." *Political Psychology* 36(5):507–523.

Wilner, Alex S. 2010. "Targeted killings in Afghanistan: Measuring coercion and deterrence in counterterrorism and counterinsurgency." *Studies in Conflict and Terrorism* 33(4):307– 329.