Countering Radicalization: An Agent-Based Model of Precision Strike Campaigns

Brendan Greenley and Dave Ohls

April 6, 2014

The emergence and expansion of United States drone operations in Pakistan, Yemen, and beyond has made the effectiveness of precision strike campaigns a hotly disputed topic at the moral, legal, and strategic levels. Debates have emerged over the number of civilian deaths associated with such strikes, their standing within domestic and international law, the perception of such campaigns by the local population, the possibility such strikes are destabilizing to the local political environment, and, most centrally, if campaigns actually accomplish their goal of improving the security environment by reducing the number of radicals and insurgents that threaten the state launching them. Some point to the elimination of specific high-level Al Qaeda leaders as proof these campaigns work, while others assert that in the place of fallen leaders new ones emerge, backed by many more radicalized insurgents. Contradicting studies, data, popular arguments, and congressional testimony highlight the challenge of analyzing the United States' secretive and often highly politicized drone campaign.

Prominent scholars, military strategists, and government officials have hailed the program as a crucial weapon in the war on terror. While cautioning that their use should be limited and run by the military (rather than the CIA), General Colin Powell asserted that "drones are a very, very effective weapon and we will continue to use them... going after the high-value targets that pose a real, immediate threat to us."[1] It is a (perhaps rare) national security issue where high-ranking members of the Obama administration and of the former Bush administration largely agree, both arguing that drone tactics are essential parts of broader counterterrorism strategy and have helped combat the threat posed by Al Qaeda. However, several equally authoritative military strategists have come out publicly against drone strikes. A former top adviser to General Petraeus asserted pointedly that the strikes create more militants than they kill.[2] In 2012, the

[1] Hunt, "Powell Says Military Not CIA Should Direct Drones" *Bloomberg Television* 24 May 2013

[2] McManus, "U.S. Drone Attacks in Pakistan 'Backfiring', Congress Told" *Los Angeles Times* 3 May

former Pakistan station chief for the Central Intelligence Agency (and previous head of the CIA's counterterrorism center) said that the criteria for strikes had become too broad and that "the unintended consequences of our actions are going to outweigh the intended consequences."[3] General Stanley McChrystal urges caution by highlighting the way in which such strikes are perceived, saying "to the United States, a drone strike seems to have very little risk and very little pain. At the receiving end, it feels like war."[4]

Can precision strike campaigns using Unmanned Aerial Vehicles (UAVs) be successful tools of counterterrorism? What is the balance between effectively targeting direct threats and avoiding resentment in response to strikes that radicalizes a population? What factors determine campaign outcomes, and what are the most effective strategic choices states can make to maximize efficiency and minimize civilian casualties? This paper seeks to explore these questions using a computational simulation-based approach to understand the dynamics of precision strike campaigns.

1 The Debate over Drones

This policy debate has been reflected in an emerging discussion in the scholarly literature. Proponents argue that drone strikes are the most effective way to accomplish their objectives. They point out that ignoring extremist leaders is dangerous, arresting them is logistically infeasible, and neutralizing them through other lethal means would incur far greater costs and do far greater damage.[5] In a 2013 paper, Patrick Johnson and Anoop Sarbahi conclude that drone strikes are indeed associated with less violence in the areas where they are launched implying at least some degree of success (though some of the effect may be attributed to militants simply moving elsewhere).[6] By employing selective violence for counterterrorist efforts in a relatively precise manner, drones are able to gather intelligence, remove members (including leaders) of organizations, force frequent relocation and reliance on less effective means of communication, and damage infrastructure.[7] They do this at relatively low cost—both financially and in human terms since pilot's lives are not put at stake. Plaw and Fricker question the objections levied by drone opponents against the tactic at large, though caution that the program should be narrowly defined and not expanded to include low-level targets.[8] Plaw, Fricker, and

2009

[3] Harris, "Drone Attacks Create Terrorist Safe Havens, Warns Former CIA Official" *The Guardian* 5 June 2012

[4] Byers, "McChrystal on Drones: A Covert Fix for a Complex Problem" *Politico* 15 February 2013

[5] Byman 2006, 2013

[6] Johnson and Sarbahi 2013

[7] Walsh 2013

[8] Plaw and Fricker 2012

Williams conclude strongly that, despite some civilian costs and popular resentment, drone strikes "have generally been effective and precise and probably the most humane self-defense option available to U.S. officials."[9]

Others have questioned this (and many of the same scholars have tempered their positions with an acknowledgement of the costs that come along with the benefits). These scholars argue that drone campaigns do more harm than good, create resentment and hostility toward the striking power among individuals who previously held no ill will, and serve as a rallying point to increase the legitimacy of and support for terrorist organizations.[10] Patrick Bergen and Katherine Tiedemann argue that while the strikes have affected the operations and morale of insurgents, the number of retaliatory attacks and suicide attacks (and support for them) has increased as the number of drone strikes increases, and that the regions where strikes occur in Pakistan are a major source of support for attacks elsewhere, such as in Afghanistan.[11] Minimizing violence in the Federally Administered Tribal Areas (FATA) is no sign that the strikes are effective, with one article saying that radicalized individuals, even if they are constrained in their ability to retaliate locally, simply move into Afghanistan to target United States, NATO, and Afghan security forces.[12] The same report suggests that while the United States has eliminated the more visible threats to its homeland—such as Osama bin Laden—the collateral damage has created many more low-level combatants and "fuels instability and escalates violent retaliation against convenient targets."[13] While engaging in direct studies measuring blowback and radicalization of views in a politically unstable and dangerous area such as FATA is difficult, related research has shown that violence against a local community does fortify support for more polarized, extreme, and hardline views. This effect is seen clearly in work linking high levels of terrorist attacks with support for right-wing parties in Israel.[14] Although addressing a different setting (and one in which the relative power dynamics of violent actor and population are reversed), these findings suggest that violence and casualties in a local community will reinforce support for actors who call for a strong response against the aggressor.

This skeptical line of inquiry is complemented by research arguing that the United States drone campaigns fail to (and are inherently unable to) follow through with key counterinsurgency tenets. Drones cannot ensure population security and are unlikely to win support among locals, factors which have been identified as essential for long-

[9] Plaw, Fricker, and Williams 2011
[10] Cronin 2013; Walsh 2013
[11] Bergen and Tiedemann 2011
[12] Hudson, Owens, and Flannes 2011
[13] Hudson, Owens, and Flannes 2011
[14] Berrebi and Klor 2006, 2008

term counterinsurgency success.[15] Whatever degree of tactical success or failure may be, drones by their nature are unable to protect civilians or enhance the authority of the local government.[16] Scholars also point out that measuring efficiency and body counts—be they militant or civilian—is only part of the picture in assessing drone's overall effectiveness; attention must be paid to their influence on recruitment and local government authority.[17]

1.1 Repression and Decapitation Strategies

To understand the strategic implications of precision strike campaigns, it is necessary to first consider how they can be used as a tactic contributing to broader counterterrorism efforts through counterinsurgency campaigns. Cronin outlines six different processes by which terrorist groups end—decapitation, negotiation, success, failure, repression, and reorientation.[18]. Of these, negotiation, decapitation, and repression are the direct result of strategies employed by the target state, and precision strikes from Unmanned Aerial Vehicles may be used as part of the latter two.

Repression, understood as the "use of overwhelming, indiscriminate, or disproportionate force," has long been a strategy employed to counter threats to the state.[19] It can be a natural and instinctual response—particularly at the level of the mass public—to respond aggressively to terrorist attacks in order to reduce the organization's ability to do further damage, if not destroy the organization entirely. Decapitation, defined as "the removal by arrest or assassination of the top leaders or operational leaders of a group," outlines a much more narrowly focused campaign.[20] The potential advantages to such an approach are clear, particularly for groups with hierarchical structures or which are heavily dependent on a small number of leaders for recruitment, organization, and mission execution. By neutralizing these leaders, the potential target state is able to limit or eliminate its ability to do harm, and to lead to the decline of the organization. Repression is more comprehensive, but decapitation is likely to generate fewer civilian casualties and less blowback.

There has been extensive empirical scholarship, though little consensus, on the influence of targeted killing campaigns (beyond drones specifically) on terrorist organizations. Even on the most fundamental question—whether it can be effective at bringing down a terrorist group—scholars disagree. Hafez and Hatfield argue that in theory decapitation strikes may have a deterrent effect on other actors, may incite a backlash against the

[15] Boyle 2010; Matulich 2012
[16] Walsh 2013
[17] Boyle 2013
[18] Cronin 2009
[19] Cronin 2009
[20] Cronin 2009

striking actor, may disrupt the operations of the terrorist group, or may make it impossible to carry on without its leadership. In practice, however, they argue that the empirical record suggests it has none of these effects.[21]. A similarly skeptical perspective comes from the work of Jordan, who finds that across 298 incidents of decapitation strikes there was no increased (and sometimes even decreased) likelihood of organizational collapse, and no consistent effect on organizational ability to carry out attacks.[22] However, Cronin argues that decapitation has been effective in some cases but not others, citing the arrest of the Shining Path's Abimael Guzmán in 1992 and the killing of Abu Sayyaf leader Abdurajak Abubakar Janjalani in 1998 as crippling their organizations, while aggressive targeting of Chechen leaders by the Russian government has only served to broaden the conflict.[23]. Mannes also finds mixed evidence, noting several of the same motivating examples as Cronin, then analyzing 60 instances of leader neutralization and 21 comparable cases where leaders remained in place.[24] He shows that decapitation has no effect on the number of fatalities in following years, but that there is a slight decline in the number of instances of attacks in some cases.

However, other scholars disagree and have a more positive perspective on the effectiveness of targeting top leadership. Case analysis of 35 leader eliminations within 19 terrorist groups lead Langdon, Serapu, and Wells to conclude that there is no evidence that decapitation strikes lead to martyrdom and increased radicalization, and there is some reason to believe that groups are slightly more likely to fail or disband after the leader has been eliminated.[25] Price argues even more strongly that decapitation can be successful, particularly over the long term. Using a hazard analysis of 207 terrorist groups' existence over time, he shows that losing a leader makes groups three to four times more likely to end at any given point.[26]

Importantly, decapitation strategies can take two distinct forms: arrest and assassination. These may be perceived differently by other members of the terrorist group or the broader population. As Cronin highlights, legal capture bring some advantages. Killed leaders are often perceived as martyrs, and a focal point for organizational rallying; captured leaders are demystified by the clear expression of their limited power relative to the state.[27] This argument is supported in part by Price, who finds that both tactics are effective but capture moreso (and capture followed by subsequent execution even

[21] Hafez and Hatfield 2006
[22] Jordan 2009
[23] Cronin 2009
[24] Mannes 2008
[25] Langdon, Sarapu, and Wells 2004
[26] Price 2012
[27] Cronin 2013

moreso).[28] Langdon, Sarapu, and Wells, however find that among their cases arrests are much less likely to result in de-radicalization or disbanding of the terrorist group than killings of leaders.[29] While not disputing that targeted killings bring drawbacks, Byman notes that in many cases there are no other options: when the terrorist leaders are located in inhospitable environments and surrounded by sympathizers capture is at best highly costly and risky and at worst essentially impossible.[30]

Finally, the literature has examined the key factors for success of decapitation strategies. There is general consensus that younger, smaller, and more hierarchical organizations with logistically-involved leaders are more vulnerable to this approach, for obvious reasons.[31] There is notable disagreement, however, on the ideological type of terrorist organization that is most susceptible to the removal of its leaders. Mannes and Langdon, Sarapu, and Wells argue that this strategy is less effective against religiously or spiritually-motivated groups, since there is a greater likelihood that blowback will occur from remaining members rallying around the unifying ideology.[32] Price disagrees, finding evidence that religious groups are significantly more vulnerable to organizational death in the wake of losing a leader, suggesting that this may be a result of the role such leaders play in framing and interpretation of causes. Cronin highlights perhaps the most important, but most difficult to measure, criteria: the likely effects of the leader's removal on those who actively or passively support the campaign.[33] If followers will be deterred, discouraged, or disoriented, targeted strikes can work; if followers will be motivated and potential sympathizers will be radicalized, decapitation will be counter-productive. This is a logical, but nearly unobservable ex ante.

1.2 Drone Data

A central challenge of applying this past knowledge to understanding the impact of current drone campaigns is the paucity of data on accuracy, impact, collateral damage, and responses. Different definitions of what constitutes a targeted killing strategy, the lack of clear metrics for identifying success, limitations on gathering empirical evidence in hostile areas, and the diversity of circumstances between cases has made generalization and confident knowledge difficult.[34] Many of the arguments against drone strikes espoused by policymakers and scholars hinge on the notion of there being excessive civilian deaths,

[28] Price 2012
[29] Langdon, Serapu, and Wells 2004
[30] Byman 2013
[31] Cronin 2009, 2013; Price 2012; Langdon, Sarapu, and Wells 2004; Jordan 2009
[32] Mannes 2008; Langdon, Sarapu, and Wells 2004
[33] Cronin 2013
[34] Carvin 2012

which leads to blowback, and that drone strikes thus create more radicalized individuals than they kill. Inaccurate campaigns with high civilian casualties and low mission success are likely to cause more problems through increased local anger and resentment than they solve; extremely precise campaigns that solely and reliably do harm only to their intended targets is more likely to be effective at preventing and deterring extremist activity.

However, the data available on militant and civilian casualties for the United States' recent use of drones are murky at best. The New America Foundation, a non-partisan think-tank, hosts one of the most widely cited, generally respected, and continuously updated data set of drone strikes. They estimate that 78-81% of casualties involved from 2004-2012 were militants and that the number of drone strikes in Pakistan peaked in 2010 and has steadily decreased since.[35] These data have been used in articles by many scholars and in many major publications, including *The New Yorker, International Affairs, Foreign Affairs, and The Wall Street Journal*[36] However, these figures are challenged by some as being too high and by others as being too low. Their reliance on news reports without having a local presence makes the numbers at best an educated guess with significant potential for noise and bias—a challenge shared by all other data sources. Many argue there is a large discrepancy between official numbers reported in the news and then cataloged in datasets and the real civilian toll, due to the secrecy of the programs and strikes, the inability of monitoring agencies to maintain a presence in the regions where strikes are occurring, and the questionable definition of what entails a civilian casualty versus that of a militant.

Five-month moving averages of the high- and low-estimates of civilian drone fatalities from the two most comprehensive data sources, the New America Foundation and the Bureau of Investigative Journalism, are shown in Figure 1. This shows how far off the numbers can be, even between two relatively middle-ground sources and even within a single sources plausible range. The low-end estimate from the Bureau of Investigative Journalism is usually higher than the high-end estimate from the New America Foundation, and generally BIJs numbers are often three to six times that of NAF. The proportion of civilian deaths as a share of total deaths also changes over time, as shown in Figure 2. Using again using five-month moving averages and BIJ data, there appears to be a general downward trend in the degree of collateral damage, perhaps suggesting improved targeting or strike technology, and possibly bringing different implications for the effectiveness of campaigns.

Figure 3 illustrates the difficulty estimating empirically the influence of drone strikes

[35] New America Foundation 2013

[36] Mayer, "The Predator War" *The New Yorker*, 26 October 2009; Entous, Gorman, and Perez "U.S. Unease Over Drone Strikes" *Wall Street Journal*, 26 September 26 2012; McCrisken 2011; Bergen and Tiedemann 2011

Figure 1: Civilian Drone Fatalities

December 2012)
5-month moving average

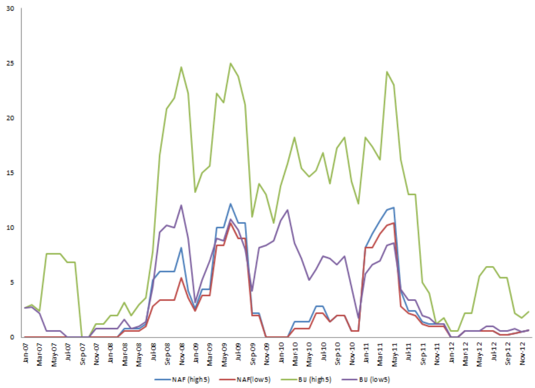

Figure 2: Civilian Death Percentage

- December 2012)
5-month moving average

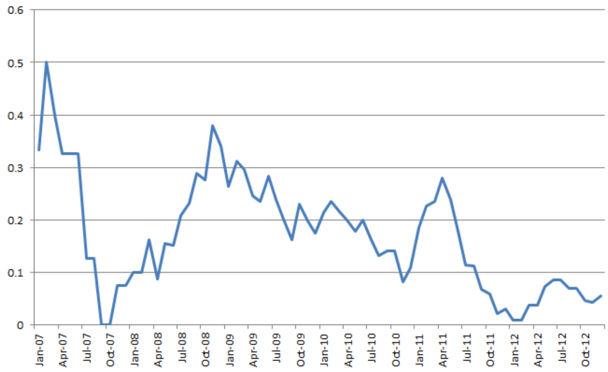

on subsequent terrorist activity. It plots the coefficient estimate and 95% confidence interval of zero-inflated negative binomial models using monthly drone casualties to predict terrorist fatalities, drawing on thirteen different estimates of drone casualties.[37] Red points indicate counts of militant deaths, blue points counts of civilian deaths, and black points counts of total deaths. None of these estimates is statistically significantly different from zero and, more importantly, they vary considerably—some positive, some negative, some with quite wide error bands and others quite narrow. Although this is a quick and crude estimation strategy, it illustrates the point that results depend heavily on data source.

Figure 3: Influence of Drone Fatalities on Terror Fatalities

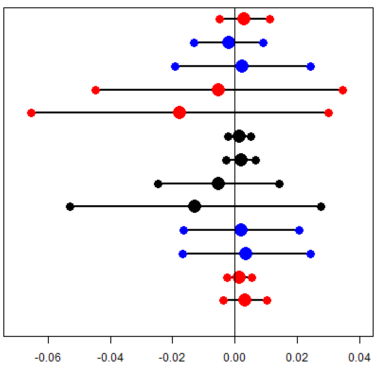

(Zero-Inflated Negative Binomial Coefficients and Confidence Intervals)

[37]Data New America Foundation, Bureau of Investigative Journalism, Long War Journal, Global Terrorism Database.

One the side of drone-pessimists, the Human Rights Clinic at Columbia Law School claims that a careful review of the news reports and other source data show drastically higher civilian casualties—between 72 and 155 in the year 2011 compared to the New America Foundation's estimate of between 3 and 9. In-depth examinations of selected cases in Yemen and Pakistan by Human Rights Watch and Amnesty International, respectively, draw a similar conclusion.[38] They suggest that local evidence shows that civilian casualties are well above those acknowledged by the government or recorded in most news-based data sources.

Yet there are others still who argue in the opposite direction, that the number of civilian deaths are exaggerated. Farhat Taj, a native of Pakistan, replies to New America's "The Year of the Drone" with her article "The Year of the Drone Misinformation," cautioning against accepting notions of widespread civilian deaths.[39] Testimony by officials involved with the drone campaign has also cast a picture of very effective and accurate strikes, with CIA Chief John Brennan (then President Obama's top counterterrorism adviser) describing in June 2011 that in nearly a year's worth of strikes "there ha[d]n't been a single collateral death because of the exceptional proficiency, precision of the capabilities we've been able to develop."[40]. This estimate is likely unrealistically low, and influenced by the administration's classification of militants to include any military-age male within a strike zone unless intelligence after the fact explicitly exonerates them (which is not often a high priority).[41] The political incentives and resulting guilty-unless-proven-innocent methodology for categorizing deaths makes official reports of civilian casualties unreliable. However, recent research by Plaw, Fricker, and Williams suggests that government figures may be roughly accurate, and that in particular improved targeting and intelligence have reduced civilian casualties as the campaigns have gone on.[42]. Along similar lines, in late October 2013 Pakistan's Ministry of Defense revised downward its earlier estimate of civilian casualties, saying that since 2008 only 67 of 2,227 drone deaths (3%) had been civilians—bringing them in line with CIA figures.[43]

However, even if consensus were reached on civilian death tolls, this does not in itself answer the question of the effectiveness of drone strike campaigns. There is a complicated interplay between advantages and limitations of targeting threats at the risk of doing collateral damage and angering a population. Much of the disagreement in the literature is a result of conflicting opinions of the net effect of this interplay. Nearly all scholars

[38] Amnesty International 2013; Human Rights Watch 2013
[39] Taj 2010
[40] Zenko 2012
[41] Becker and Shane, "Secret 'Kill List' Tests Obama's Principles" *The New York Times*, 29 May 2012
[42] Plaw, Fricker, and Williams 2011
[43] Walsh, "In a Surprise, Pakistan Says Fewer Civilians Died by Drones" *The New York Times*, 30 October 2013

believe that judicious, cautious, and rare strikes accomplish some objectives and that frequent, indiscriminate, and careless attacks generate some hostility. The question is in the balance, and it is this balancing trade-off effect that we explore. Further, although prior research has uniformly considered the importance of top organizational leaders, it has used different understandings of who constitutes a leader, and has not paid careful attention to the effects of defining this variable narrowly or broadly. This project also seeks to explore the strategic implications for effectiveness of this choice.

2 Modeling Precision Strike Campaigns

This project employs a computational simulation approach using agent-based modeling to explore the dynamics of precision-strike campaigns aimed at countering radicalization. This follows a long line of scholarship using agent-based models to understand insurgencies, radicalization, and terrorist attacks.[44] However, none of this work looks at the effectiveness of precision or drone strikes, which has quickly become a central topic for academics and policymakers alike in understanding the empirical patterns of counterinsurgency. This project complements past research by introducing that element.

A simulation approach is ideally suited for the data-poor environment surrounding this topic. It allows for repeated trials under different assumptions, and can observe how the tradeoffs play out under different scenarios or strategies. Outcomes such as civilian casualties, long-term average levels of radicalization, the count of high-level terrorist leaders, and other parameters can be measured precisely. This allows us to see how drone strikes affect a local population, when terrorist networks are crippled and when blowback is maximized, and the relative effectiveness of repression versus decapitation campaigns.

2.1 Model

The model uses agent-based procedures and is coded in NetLogo, an open-source, free, and multi-platform programming environment with an easily customizable graphical user interface (GUI) that allows easy modification of variables and collection of data.[45] Agent-based modeling allows for the simulation of complex systems and relationships by populating a world with individual agents who interact with each other based on a given set of instructions provided by the model. For our model, this leads to the influence of agents to be dictated not by a large differential equation, but by a more simple set of

[44] See, e.g., Doran 2005; Stauffer and Sahimi 2006; Bhavnani, Miodownik, and Nart 2008; Cioffi-Revilla and Rouleau 2010

[45] Wilensky 1999

Figure 4: Model Interface

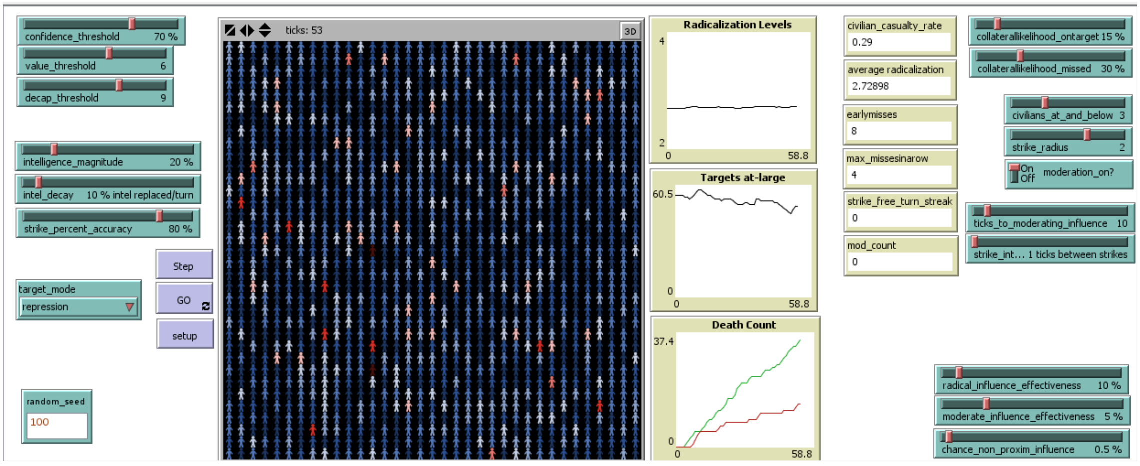

instructions that uses situational factors such as an agent's proximity to radical actors to determine whether or not it will be radicalized that turn or not.

The model comprises five major procedures: the initial setup of the environment, and the influence, targeting, strike, and post-strike algorithms. The model repeatedly calls upon these algorithms cyclically until the time cutoff is met.

Sliders and switches populate the GUI (Figure 4), allowing a user to change a variety of values, including the type of campaign (decapitation or repression), target selection value and confidence threshold, strike percent accuracy, magnitude of intelligence, collateral of both on-target and missed strikes, the effectiveness of radicals in influencing agents in their neighborhood, the effectiveness of moderates in influencing agents in their neighborhood, the effectiveness of extreme radicals in influencing agents throughout the world, the radius in which strikes may hit, the radius in which neighborhood effects occur, and the frequency of strikes. Once the desired variables are set, the user clicks the setup button to populate the world. Every square in a 35 by 35 Cartesian plane is then given an agent with a randomly assigned radicalization value along a Poisson distribution curve with $\lambda = 2.5$, truncated to have a maximum possible radicalization value of 10 (highly-radical) and a minimum of 1 (pacifist). The radicalization level of each agent is visualized by the model with pacifist agents being dark blue, highly radical agents being dark red, and a gradient of color for agents between the two extremes. Once the model is started, time (measured in 'ticks') passes, with actions happening at specified time increments.

With every tick, agents have a chance to influence each other. Extreme radicals (radicalization of 9 or 10) have the ability to influence one agent anywhere in the world, with E% probability set by the extremist influence effectiveness variable. This captures the degree to which terrorist leaders may have networked connections across space. Agents with a radicalization level greater than or equal to 6 have the ability to influence nearby agents to become more extreme, with R% probability set by the radical influence effectiveness variable. An agent can influence a maximum number of agents per turn equal to its radicalization level minus 5 (so more radical agents have more potential for influence), and cannot influence those that are already more radical than itself; nor can it influence those with a radicalization level of 1 who are considered pacifist and immune to radicalization persuasion. Finally, agents with a radicalization level of 1 can influence nearby agents to become more moderate, with M% probability set by the moderate influence effectiveness variable. They cannot persuade agents with radicalization 9 or 10, but may be able to pacify others slightly. This moderating influence can only occur only if there has been (variable length) period with no strikes.

Each tick, the precision campaign applies the targeting algorithm. Intelligence generates a list of agents equal to the intelligence magnitude N% times the total number of individuals in the population. This observation produces a signal of the agent's true radicalization, distributed according to a truncated Poisson distribution with shape parameter equal to the true value and adjusted to only include values between 1 and 10. The power observes this signal and applies Bayesian updating to adjust their beliefs about the likelihood the agent takes on any given radicalization threshold. For agents which have never before been observed, prior beliefs are based the distribution of agents in the population, and for agents which have previously been observed the best estimates are retained and updated on each subsequent observation. Thus every time the campaign seeks a target, a small percentage of the world's agents are known by intelligence, with their radicalization values estimated according to an imperfect signal, while the large majority of actors in the world remain unknown.

The value threshold V and confidence threshold C% determine whether any observed agents are considered valid targets. V sets the level at or above which the power wishes to strike, and C% sets the degree of certainty the power must have that a given agent is at or above that level. If intelligence discovers no agent whose radicalization level is greater than or equal to the target selection threshold with the necessary level of confidence, there is no strike. Otherwise, the procedure chooses the agent with the highest expected radicalization value and passes its location to the strike algorithm.

Each time a target is passed to the strike algorithm, a random number between 0 and 1 is generated. If the number is less than the strike accuracy variable, a strike hits its target accurately and the intended target is killed. Otherwise, the strike misses, and kills an agent residing in another square within the strike radius. Strike accuracy includes the chances of failure associated not only with malfunction in the payload or targeting equipment used in a strike, but also with failures in intelligence gathering. For

Figure 5: Targeting Diagram

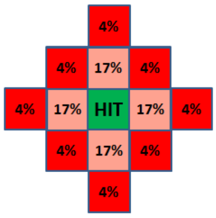

a strike radius of 2 (using von Neumann neighborhoods because the world is composed of squares, see Figure 5), miss probabilities are distributed based on a discrete adaptation of the normal distribution as follows: a 68% chance to hit one of the four squares (17% chance each) one away from the intended target, and a 32% chance to hit one of the eight squares two away from the intended target (4% chance each). Once a strike hits, the model checks to see if collateral damage occurs. The model has two collateral likelihood variables, one for misses and one for accurate hits, as it is assumed that if a strike misses it may have a higher chance to kill other civilians than a strike that hits its target. Like the targeting algorithm, a random number between 0 and 1 is generated, and if it is less than the given collateral likelihood variable, collateral damage occurs. The probability distribution for where collateral occurs is identical to a missed strike distribution, though the reference square is not the original strike target, but where the strike actually hit (thus different for misses).

For each civilian death, all agents within a two-step neighborhood of the now vacant square have their radicalization levels increased by one (this value is still capped at 10). Witnessing a civilian death is the only way an agent with a radicalization level of 1 can become more radical and susceptible to influence by other agents. For each militant death, only agents in the neighborhood that are already moderately radical increase their radicalization level. All empty squares are then repopulated with new agents, whose radicalization values are randomly drawn from the same Poisson distribution used during the initial world setup. The intelligence loses track of D% of known agents, according to the intelligence decay parameter, and replaces them with a new random draw of the same size.

2.2 Experimental Manipulation

To explore the dynamic patterns of this system, we manipulate two sets of variables at different levels. First, we allow the two main variables capturing the striking power's capabilities to take on different values. Intelligence magnitude—the percentage of agents in the world that the power is able to observe each tick—is set at 5%, 10%, 15%, or 20%. Strike accuracy—the likelihood with which a strike hits the intended agent—is set to 60%, 70%, and 80%. We expect higher values of these parameters to be associated with more successful campaigns. The greater the government's ability to accurately find and target intended extremists, the greater the frequency of eliminating major threats and the lesser the frequency of civilian casualties.

Second, we vary the parameters determining which agents are considered viable targets. This is determined by two variables. The value threshold–the level at or above which the power wishes to strike–is set at 5, 6, 7, 8, 9, and 10. The confidence threshold—the likelihood with which it must believe the target is at or above that value—is set at 40%, 50%, 60%, and 70%.

This defines 288 different parameter combinations, each of which is run 10 times, for a total of 2,880 runs of the model.

3 Results

The results for levels of intelligence and precision are largely as expected. The greater the accuracy of strikes and the greater proportion of the population which is observed each tick, the more successful the campaign will be. Both of these parameters are associated with lower average radicalization levels (Table 1) and lower counts of agents above certain levels remaining (Table 2). Interestingly, accuracy seems to matter far more than intelligence. Increasing the intelligence magnitude fourfold yields only a slight decrease in the average radicalization and number of insurgents remaining at the end of campaigns, while moderate increases in strike precision have great effects. This suggests that missed strikes causing collateral damage, rather than failure to identify appropriate targets, is the largest strategic challenge states face, and that the marginal return from investing resources on precision technology is likely greater than from investing in greater information-gathering mechanisms.

Table 1: Average Radicalization

		Strike Accuracy		
		60%	70%	80%
	5%	2.994	2.864	2.759
Intelligence	10%	2.962	2.834	2.737
Magnitude	15%	2.948	2.823	2.726
	20%	2.936	2.806	2.716

(no strikes: 2.772)

Table 2: Count of Radicals ≥6 / ≥8 Remaining

		Strike Accuracy		
		60%	70%	80%
	5%	183 / 59	119 / 32	79 / 17
Intelligence	10%	157 / 48	96 / 23	63 / 11
Magnitude	15%	142 / 41	88 / 20	58 / 10
	20%	134 / 38	80 / 17	55 / 9

(no strikes: 110 / 23)

Several interesting patterns emerge when looking at the effects of strategic targeting decisions on campaign outcomes. As shown in Figure 6 and Table 3, setting the value threshold too low—at 5 or 6 in the terms of the model—has clearly detrimental effects. More civilians are killed, more radicals are created, average radicalization values are higher, and the density distribution of radicalization outcomes includes a much greater frequency of highly values. However, the results also show that setting the value threshold too high—at 9 or 10 in terms of the model—is also not ideal. Although less dramatically so than at the low end, slightly more undesirable outcomes are observed when only the very top agents are considered eligible for targeting. This suggest an ideal targeting strategy in the middle ground.

Figure 6: Value Threshold

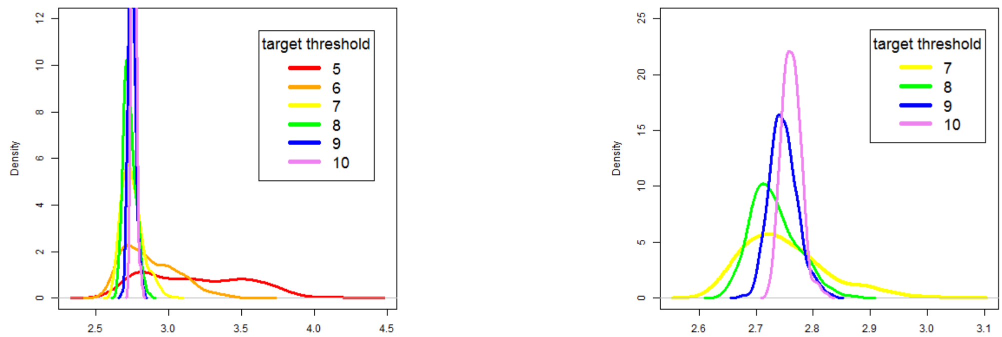

5 through 10 | 7 through 10

The degree of confidence about an agent's extremism that the power requires to consider it a valid target does not appear to matter as much. As shown in Figure 7, the density of radicalization is very similar across different thresholds. To the degree this parameter does have any effect, it appears that higher confidence levels—more cautious strike campaigns—are slightly more effective, as they result in a lower average level of

Table 3: Value Threshold

value threshold	civilians killed	radicals killed	civilian proportion	mean radicalization	radicals ≥6	radicals ≥8	radicals ≥10
5	205.8	684.4	0.23	3.19	221.7	89.2	20.4
6	70.9	218.7	0.24	2.88	78.5	23.2	4.2
7	19.8	58.4	0.25	2.75	62.5	8.4	1.0
8	6.4	19.9	0.25	2.73	71.7	7.5	0.5
9	1.9	5.8	0.25	2.74	90.7	14.4	0.8
10	0.5	1.4	0.26	2.76	102.9	19.7	1.6

radicalization, and fewer civilian deaths.

Figure 7: Confidence Threshold

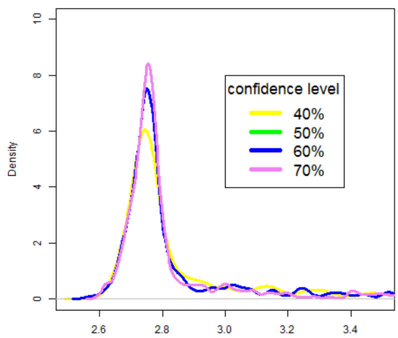

To confirm that the observed patterns are as they appear, we employ linear regression analysis. Because these are experimental data and variables of interest are exogenously set as parameters of the model, there is no concern about non-independence or omitted variable bias, and simple ordinary least squares regression is appropriate. Results of preliminary models are presented in Table 4

Model 1 shows that, as is clear from the chart, greater accuracy and greater intelligence are associated with more successful strike campaigns. For each additional 1% accuracy, the predicted mean radicalization level that results is reduced by 0.011. The marginal effect of increased intelligence is smaller; each additional 1% of agents observed is associated with only 0.003 lower level of radicalization.

The results for confidence and value thresholds are more nuanced. As shown in Model 2, more reticent campaigns are predicted to be more effective, with negative and significant coefficients on both confidence and value thresholds. However, density plots had suggested that there is a limit to how high the value threshold should be set. This intuition is tested in Model 3, which restricts the cases to those where this threshold was

Table 4: Mean Radicalization: Regression Results

	Model 1	Model 2	Model 3 [7-10 only]	Model 4
strike accuracy	-0.011***			
	(4.8e-4)			
intelligence magnitude	-0.003***			
	(7.1e-4)			
confidence threshold		-9.7e-4**	-1.3e-4	-9.7e-4***
		(3.3e-4)	(9.4e-5)	(2.8e-4)
value threshold		-0.072***	0.005***	-0.659***
		(0.002)	(9.4e-4)	(0.019)
value threshold ^{2}				0.038***
				(0.001)
(intercept)	3.68	3.44	2.71	5.52
	(0.04)	(0.02)	(0.01)	(0.07)

set to at least 7. Indeed, within this subsample, there is now a positive and significant relationship between threshold and average radicalization: campaigns which are too hesitant to strike will be less effective. This suggests a non-monotonic relationship, which is estimated in Model 4. The results demonstrate the U-shaped effect of value threshold on average radicalization that results from precision strike campaigns. There is an ideal threshold between the extremes that minimizes the radicalization of a population after a strike campaign.

4 Discussion

Agent-based modeling can serve as a useful tool for exploring the pros and cons of different drone strategies under different parameter assumptions. When data is limited, conclusions using assumptions about empirical facts can vary widely based on which set of assumptions are chosen, leading to incoherent or unclear implications. Simulation is not a replacement for real-world data gathering and estimation, but rather a complement in two ways. First, it quickly generates comparisons of the effects of different strategic choices under different real-world conditions. This can facilitate understanding, prediction, and policy-making for any given set of estimates and assumptions about the parameters that matter. Second, it allows comparative statics analysis—showing how outcomes change as parameters (exogenous facts about the world or endogenous choices) change. Such analysis shows how factors shifting leads to shifts in results, and do not depend on the exact values those factors take on (which, in a simulation environment, do not directly translate to real-world circumstances).

Substantively, the models suggest that broader repression strategies can be more effective at minimizing average radicalization in a society and reducing the number of high-value targets than decapitation campaigns. Setting targeting thresholds too high allows radicalization to spread unchecked, as extremist actors that pose real, immediate threats are able to remain in existence, carry out missions, and influence other agents. However, broad campaigns come at a cost. There is an inherent and unavoidable trade-off whereby lower threshold values and confidence levels increase collateral damage and civilian deaths, which is undesirable both for its own sake and for its enhancing effect on blowback which results.

The model results also show that although the threshold levels should not be set too high, they also should not be set too low. When targeting is too aggressive, the results can be disastrous. Large numbers of civilians are killed, and the campaign ultimately backfires strategically, leading to a far more radicalized population than if the power had simply taken no action. Precision targeting campaigns must be careful not to relax standards to the point this occurs.

Finally, it is important to acknowledge the questions this approach does not help us address. We neither attempt nor claim to speak to the ethical questions, or to define what constitutes an acceptable level of collateral damage given a certain degree or likelihood of mission success. Similarly, we are agnostic on the legal status of precision campaigns at both the domestic and international level. The model generated here simply poses the strategic questions: what are the dynamics of the tradeoff between eliminating threats and engendering hostility, and how do these dynamics change under varying circumstances and tactical choices. However, these strategic questions do have implications for the broader questions; knowing the factors that influence outcomes in drone campaigns is an essential first step for debating the social, legal, or moral issues.

References

Amnesty International. 2013. *'Will I Be Next?' US Drone Strikes in Pakistan.* Available from: http://www.amnestyusa.org/sites/default/files/asa330132013en.pdf.

Bergen, Peter and Katherine Tiedemann. 2011. Washington's Phantom War: The Effects of the US Drone Programs in Pakistan. *Foreign Affairs* 90(4): 12-18.

Berrebi, Claude and Esteban F. Klor. 2006. On Terrorism and Electoral Outcomes: Theory and Evidence from the Israeli-Palestinian Conflict. *Journal of Conflict Resolution* 90(6): 899-926.

——-. 2008. Are Voters Sensitive to Terrorism? Direct Evidence from the Israeli Electorate. *American Political Science Review* 102(3): 279-301.

Bhavnani, Ravi, Dan Miodownik, and Jonas Nart. 2008. REsCape: An Agent-Based Framework for Modeling Resources, Ethnicity, and Conflict. *Journal of Articficial Societies and Social Simulation* 11(2). Available from: http://jasss.soc.surrey.ac.uk/11/2/7.html.

Boyle, Michael J. 2010. Do Counterterrorism and Counterinsurgency Go Together? *International Affairs* 86(2): 333-353.

Boyle, Michael J. 2013. The Costs and Consequences of Drone Warfare. *International Affairs* 89(1): 1-29.

Byman, Daniel. 2006. Do Targeted Killings Work? *Foreign Affairs* 85(2): 95-111.

——-. 2013. Why Drones Work: The Case for Washington's Weapon of Choice. *Foreign Affairs* 92(4): 32-43.

Carvin, Stephanie. 2012. The Trouble with Targeted Killing. *International Affairs* 21(3): 529-555.

Cioffi-Revilla, Claudio and Mark Rouleau. 2010. MASON RebeLand: An Agent-Based Model of Politics, Environment, and Insurgency. *Intenrational Studies Review* 12(1): 31-52. Available from: http://cs.gmu.edu/∼eclab/projects/mason/publications/rebeland09.pdf.

Cronin, Audrey Kurth. 2006. How Al-Qaida Ends: The Decline and Demise of Terrorist Groups. *International Security* 31(1): 7-48.

——-. 2009. *How Terrorism Ends: Understanding the Decline and Demise of Terrorist Campaigns.* Princeton, NJ: Princeton University Press.

——-. 2013. Why Drones Fail: When Tactics Drive Strategy. *Foreign Affairs* 92(4): 44-54.

Doran, Jim. 2005. Iruba: An Agent-based Model of the Guerrilla War Process. *Representing Social Reality.* Pre-proceedings of the 3rd Conference of the European Social Science Association. Available from: http://cswww.sx.ac.uk/staff/doran/34.PDF.

Grut, Chantal et al. 2012. Counting Drone Strike Deaths. *Human Rights Clinic, Columbia Law School.* Available from: http://web.law.columbia.edu/sites/default/files/microsites/human-rights-institute/COLUMBIACountingDronesFinalNotEmbargo.pdf.

Hafez, Mohammed and Joseph M. Hatfield. 2006. Do Targeted Assassinations Work? A Multivariate Analysis of Israeli Counter-Terrorism Effectiveness during Al-Aqsa Uprising. *Studies of Conflict & Terrorism* 29(4): 359-382.

Human Rights Watch. 2013. *Between A Drone and Al-Qaeda: The Civilian Cost of US*

Targeted Killings in Yemen.
Available from: http://www.hrw.org/sites/default/files/reports/yemen1013_ForUpload.pdf.

Hudson, Leila, Colin S. Owens, and Matt Flannes. 2011. Drone Warfare: Blowback from the New American Way of War. *Middle East Policy* 18(30): 122-132.

Johnson, Patrick J. and Anoop Sarbahi. 2003. The Impact of US Drone Strikes on Terrorism in Pakistan. Unpublished manuscript, RAND Corporation and University of California–Los Angeles.
Available from: http://patrickjohnston.info/materials/drones.pdf.

Jordan, Jenna. 2009. When Heads Roll: Assessing the Effectiveness of Leadership Decapitation. *Security Studies* 18(4): 719-755. doi: 10.1080/09636410903369068.

Langdon, Lisa, Alexander J. Sarapu, and Matthew Wells. 2004. Targeting the Leadership of Terrorist and Insurgent Movements: Historical Lessons for Contemporary Policymakers. *Journal of Public and International Affairs* 15: 59-78.

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

Matulich, Peter. 2012. Why COIN Principles Don't Fly with Drones. *Small Wars Journal* 8(2).

McCrisken, Trevor. 2011. Ten Years on: Obama's War on Terrorism in Rhetoric and Practice. *International Affairs* 87(4): 781-801.

New America Foundation. 2013. The Year of the Drone.
Available from: http://counterterrorism.newamerica.net/drones.

Plaw, Avery and Matthew S. Fricker. 2012. Tracking the Predators: Evaluating the US Drone Campaign in Pakistan. *International Studies Perspectives* 13(4):344-365.

Plaw, Avery, Matthew S. Fricker, and Brian G. Williams. 2011. Practice Makes Perfect? The Changing Civilian Toll of CIA Drone Strikes. *Perspectives on Terrorism* 5(5-6): 51-69.

Price, Bryan C. Targeting Top Terrorists: How Leadership Decapitation Contributes to Counterterrorism. *International Security* 36(4): 9-46.

Stauffer, Dietrich and Muhammad Sahimi. 2006. Discrete Simulation of the Dynamics of Spread of Extreme Opinions in a Society. *Physics A: Statistical Mechanics and its Applications*. 364: 537-543.
Available from: http://econpapers.repec.org/RePEc:eee:phsmap:v:364:y:2006:i:c:p:537-543.

Taj, Farhat. 2010. The Year of the Drone Misinformation. *Small Wars & Insurgencies* 21(3): 529-535. doi:10.1080/09592318.2010.505486.

Walsh, James Igoe. 2013. The Effectiveness of Drone Strikes in Counterinsurgency and Counterterrorism Campaigns. *United States Army War College Press*: Carlisle, PA.

Wilensky, Uri. 1999. NetLogo. Center for Connected Learning and Computer-Based Modeling, Northwestern University. Evanston, IL.
Available from: http://ccl.northwestern.edu/netlogo.

Zenko, Micah. 2012. Politics, Power, and Preventive Action: How Many Civilians Are Killed by U.S. Drones? *Council on Foreign Relations*.
Available from: http://blogs.cfr.org/zenko/2012/06/04/how-many-civilians-are-killed-by-u-s-drones.