

Being Your Own Worst Enemy, and Learning to Win: A Model of Government Learning in Responding to Insurgency

D. Scott Bennett
Department of Political Science
The Pennsylvania State University
318 Pond Lab
University Park, PA 16802
(814)865-6566
sbennett@psu.edu

May, 2008

Work very much in progress – please do not cite without permission of the author

Abstract

Bennett (2008a) introduced an agent-based simulation of insurgency government which modeled the interactions of civilians, insurgents, and soldiers (government forces). That simulation allowed insurgents to attack government (military) targets, and in turn be attacked and captured or killed by government forces. When government forces inflicted collateral damage, it risked creating new insurgents. The paper suggested that accuracy in avoiding collateral damage was more important for defeating insurgencies quickly than being effective at capturing a targeted insurgent in any one effort. This paper expands the initial model by exploring what happens when governments can “learn” and increase their accuracy and effectiveness rates in the midst of a conflict. The results suggest that there is a range of circumstances under which learning to be more accurate and effective can help governments to defeat insurgencies. However, the lengths of insurgencies where governments “start dumb” but then learn are dramatically longer than situations where governments are “smart” in the first place. That is, learning is no substitute for appropriate initial activity. In addition, there is a range of conditions (starting conditions for government accuracy and effectiveness, and magnitude of learning) where learning may not be enough. Even if the final learned strategy would have been enough to defeat a nascent insurgency, if an insurgency has become well established under an inferior strategy, the strategy may not be enough. If we equate learning and strategy change with the U.S. “surge” in Iraq, we learn that surges may sometimes work, may sometimes fail, and even a final strategy that would be good enough to win if used early may be inadequate if used too late.

Introduction

Bennett (2008a) introduced an agent-based simulation of the early stages of an insurgency against a government. The simulation modeled the interactions of civilians, insurgents, and soldiers (government forces), allowing actions by insurgents who could attack government (military) targets, and in turn be attacked and captured or killed by government forces. Civilians constituted an audience for insurgent-soldier interactions, and could join or be deterred from joining the insurgency. The initial simulation focused on the conditions under which soldiers' actions would lead to the growth or defeat of the insurgency. In particular, it focused on combinations of *accuracy* and *effectiveness* and how they affected the duration of nascent insurgencies. Accuracy reflected the ability of soldiers to avoid collateral damage. Specifically, after an attack on an insurgent, each nearby civilian was injured with probability 1-accuracy. Effectiveness reflected the ability of soldiers to capture or kill insurgents; soldiers removed targeted insurgents with a probability equal to the soldier's level of *effectiveness*. When government forces inflicted collateral damage, it risked creating new insurgents. When governments "missed" targeted insurgents, the insurgents remained to carry out further attacks on government targets. The paper suggested that accuracy in avoiding collateral damage was more important for defeating insurgencies quickly than being effective at capturing a targeted insurgent in any one effort. An extension to the model (Bennett 2008b) added peaceful recruitment to the initial simulation, allowing both insurgents and governments to recruit civilians in their favor. The core results on accuracy and effectiveness remained in the extended model

In that initial simulation, the conditions of government counterattacks vs. insurgents were static. That is, soldiers' levels of effectiveness and accuracy were fixed through the simulation. Even with fixed levels of effectiveness and accuracy, interesting dynamics that resulted in the spread and defeat of insurgencies emerged. But an interesting question emerges from considering whether effectiveness and accuracy might *change* over time, and what effects this might have. Accuracy and effectiveness might change either as a result of learning by individual soldiers on the battlefield, or as a result of new training or strategy change "at the top." If soldiers can learn to be more accurate or effective over time, what difference does it make? Can learning "fix" the mistakes of early inaccuracy or ineffective soldiers? Or once insurgencies become well-established as a result of poor government strategy early on, is learning useless? And, more likely, under what conditions can learning help, and when does it not matter? This paper introduces and explores a new version of the simulation that allows soldiers to learn (improve accuracy and effectiveness) either in one shot, or incrementally over time. With this simulation we can explore how this particular aspect of learning affects the course of insurgency.

Below, I start by summarizing the initial insurgency model, and then proceed to its expansion to allow learning.

The Initial Model

The simulation assumes two core types of agents: civilians, and soldiers. Soldiers represent the government, or potentially troops of a supporting (or occupying) power. No distinction for simulation purposes is made between government soldiers and international soldiers; all soldiers are taken as representatives of a group to which the insurgents are opposed. Civilians represent the mass of a population. Individual civilians are characterized by

a degree of anger at the government, a degree of fear of the government, and a propensity to use violence under the right circumstances. An *insurgent* is a civilian who is willing to engage in violence against the government, and specifically against a nearby soldier. Insurgents in turn may be *latent* insurgents who are willing to engage in violence, or *active* insurgents who have actually done so.

Growth of insurgency occurs in the model when ordinary civilians turn into insurgents, which occurs under particular circumstances. The transition decision rule for the civilian agents is that if a civilian is more angry than afraid, and if the civilian's anger passes an individual threshold propensity to use violence, then this civilian becomes a *latent insurgent*. This is a civilian who is angry enough, and is not so afraid of the government, that they will attack a government target (a soldier) if given the opportunity. If a latent insurgent is within a defined range of a soldier and is given an opportunity, then the latent insurgent will conduct an attack and become an *active insurgent*.

Any attack by an insurgent is assumed to expose a previously-hidden insurgent, and allows a targeted soldier to respond to the attack by targeting the insurgent with a *counterattack* of their own. The intent of such counterattacks is to remove (capture or kill) the insurgent so that the insurgent is unable to make further attacks, and reduce the size of the insurgency. Such counterattacks may also deter other civilians from becoming insurgents by increasing their level of fear.

It is the government's assumed goal to remove all of the latent insurgents from the world, thereby defeating the insurgency. But a soldier's counterattack can sometimes backfire. The simulation models effects on both the targeted insurgent and surrounding civilians that result from counterattacks. First, a counterattack may remove (capture or kill) an insurgent. But this is not certain, as even effective militaries are not always able to capture or kill the insurgents that they pursue. Soldiers remove targeted insurgents with a probability equal to the soldier's level of *effectiveness*. Second, a soldier's counterattack may affect the anger and fear levels of surrounding civilians. Militaries run the risk of inflicting "collateral damage" or surrounding homes or individuals when they engage in anti-insurgent activity. The probability that a soldier's counterattack inflicts collateral damage on each civilian in the neighborhood of the targeted insurgent is parameterized by the soldier's level of *accuracy* (each nearby civilian is injured with probability $1 - \text{accuracy}$, and multiple civilians may be hurt in the soldier's counterattack). When any civilian is hurt, his/her level of fear of the government increases (a deterrent effect), but their level of anger also increases *in proportion to the total number of civilians injured* by the soldier's counterattack. Counterattacks against insurgents can thus backfire because they can create more insurgents than are removed if they increase the level of anger of a significant number of civilian agents.

With these features in place, the dynamic of the simulation is one where an insurgent attack triggers a response that may remove an insurgent, but also runs the risk of creating more latent insurgents. The relative risk of these occurrences, and the variety of interesting dynamics that can result from the model, emerge from the combinations of values on these critical parameters that influence the course of insurgencies. The main analytic focus of version 1 of the insurgency model was on tradeoffs and effect of the accuracy and effectiveness of government forces (denoted "soldiers"). That simulation indicated that it was very easy for counterattacks against insurgents to be self-defeating, and illustrated that an insurgency could grow rapidly and become widespread purely as a result of government actions against insurgents. It further suggested that a higher level of *accuracy* among government forces (avoiding civilian injury and the resulting widespread anger increases) was more important in defeating an insurgency quickly than was a higher level of *effectiveness* at capturing insurgents

in any given counterattack, particularly if we hypothesize a real-world tradeoff between the two concepts.

The course of the simulation is that in each time tick of the simulation, one latent insurgent is randomly selected to attack a soldier in range of the insurgent. Following this attack, the soldier responds with a counterattack. With probability equal to *effectiveness*, the insurgent is removed from the simulation. Then, with probability equal to *1-accuracy*, each civilian in the neighborhood surrounding the insurgent is injured. The level of *fear* of each injured civilian in the neighborhood is increased, and the level of *anger* of each injured civilian is increased by an amount proportional to the total number of civilians injured. If anger and fear levels change appropriately, either latent insurgents in the neighborhood may become non-latent insurgents, or non-insurgents may become latent insurgents. In combination with the (possible) removal of the initial insurgent, this may lead to a net increase or decrease in the total number of latent insurgents. The simulation continues until all of the insurgents in range of a soldier that could be targeted are removed from the simulation, or the simulation hits a specified tick cutoff (typically simulated at 2000 or 5000 ticks, indicating a self-sustaining insurgency, which is typically quite obvious well before the cutoffs used in the simulations).

The simulation was programmed in Java using the RePast simulation libraries. RePast handles many standard agent-based simulation functions with ease, such as drawing and displaying maps and displays, keeping track of agents and their status, and recording the history of multiple simulation runs. Code was developed for the soldier and civilian agents, to specify the agents' decision rules, and to specify the sequence of actions. A single iteration of the model can take from a few seconds (if the insurgency ends with 50 to 100 ticks), or a few minutes (if the insurgency becomes self-sustaining). Code for the simulation is available from the author (eventually on the author's website).

The Learning Model

Version 1 of the insurgency model (Bennett 2008a) produced very interesting and suggestive results about the dynamics of insurgency. It intentionally focused on insurgent actions and government reactions to show how insurgency may be abetted by poorly-planned (i.e. inaccurate) government responses. Version 2 of the model (Bennett 2008b) added peaceful recruitment options to the model. But both versions assume that government soldiers had fixed levels of accuracy and effectiveness. This may not always be the case. In reality, soldiers gain experience, military and political leadership changes, and governments and militaries react to what they see occurring on the battlefield. Any of these may lead to changes in levels of accuracy and effectiveness. I will refer to such changes in accuracy and effectiveness levels as "learning." But here I have a broad conception of learning, which might include real cognitive learning among (human) soldiers, or might include central direction about weaponry or strategy that has the effect of changing these levels. Also, learning generally has a positive connotation – we learn in order to be better at some task. And in most circumstances described and analyzed below, it is true that changes and learning will be positive given the military goals of defeating insurgency. But the simulation will actually allow for "learning" to be both positive and negative. For instance, a military strategy might be imposed on soldiers that leads to better outcomes (supposing, say, that it increases accuracy), or it could lead to worse outcomes (if it deemphasized accuracy, for instance). For purposes of this paper, I will fit any of these changes under the rubric of "learning."

More specifically, we could conceive of several reasons for levels of accuracy and effectiveness in soldiers to change in a counterinsurgency situation.

- Individual soldiers or units may learn to be better at counterinsurgency warfare. They may do a better job in targeting, or learn to be sensitive to local customs and venues, decreasing collateral damage in all of its forms (and so increasing accuracy). Or they may become better at picking their target out in a crowd, or following a target to the point where the target may be captured (increasing effectiveness). Perhaps soldiers in a counterinsurgency environment learn to spot booby traps or dangerous situations so that they do not need to respond with overwhelming violence, again reducing collateral damage and increasing accuracy. Such learning may occur regardless of any central direction.
- When they realize they are fighting a sustained insurgency, militaries may engage in increased counterinsurgency training. Such training and the emphasis it places on intelligence, combined military/political/economic engagement, and a “light hand” is rather different from more common traditional military attrition or movement style strategies. Additional training about general urban warfare, for instance, might increase accuracy and/or effectiveness. Unit-level training emphasis on factors such as gathering local intelligence, or more training in local language and custom, similarly works to improve these factors. Such general training has increased in the U.S. military during the course of its involvement in Iraq, for instance, as illustrated by the rewriting and reissuing of the U.S. Army/ Marine Corps Counterinsurgency Field Manual in late 2006.
- Strategy choices made at top military and political levels may affect levels of accuracy and effectiveness. If, as in the U.S. case in Iraq, a conscious decision is made to use new counterinsurgency strategy that focuses on avoiding collateral damage, accuracy may improve across the board among soldiers. But apart from specific counterinsurgency doctrines, rules of engagement and an emphasis on deterrence strategies (heavy handed, capture insurgents at all costs) vs. more cooperative “hearts and minds” strategies (lighter, careful to avoid damage, careful to avoid increasing local resentments) may be made in the absence of a clear counterinsurgency doctrine. These choices need not always be favorable in the context of the model – a state might, for instance, decide to move towards a scorched-earth type policy where maximum destruction and effectiveness is emphasized over accuracy.
- Technology and weaponry may be developed or distributed with particular profiles that affect accuracy and/or effectiveness. For instance, we could imagine that developing better guidance systems for “smart” bombs might increase accuracy and effectiveness. Reducing the size of bombs might increase accuracy (avoidance of civilian fatalities), but deployment of a new large bomb lower accuracy at the same time it increases effectiveness. The replacement of inaccurate gun sights on some weapons with better ones, or replacement of hand-held weapons with old technology with a new generation that is easier to use, might increase effectiveness at no cost to accuracy. Here, deployment choices, or rules of engagement as to when to use very destructive weapons, might involve tradeoffs between accuracy and effectiveness, or might be able to increase or decrease both, depending on the specific choices made.

In the simulation, learning (change) occurs in two parameters, accuracy and effectiveness. This version of the simulation (version 3) builds on version 2 of the simulation which added recruitment to the base model, but in all of the results and sequences of actions I present here, all recruitment options are turned off. The degree of learning can be varied in three ways:

- When does learning occur?
- Does learning occur once, or several times?
- What is the magnitude of learning when it does occur?

The first core addition to the model is a specification of when learning occurs in the model (if it occurs at all). A tick when initial (and perhaps only) learning occurs is specified in the `learningInitialTime` parameter of the model. Learning may occur only once, or multiple times; the parameter `learnHowManyTimes` determines how many times learning will occur. If it occurs more than once, then learning occurs (by construction) at regular intervals, with the interval specified by the parameter `learningInterval`.

The magnitude of learning when it occurs is determined by two parameters. First, a value for magnitude is given in the parameter `learningIncrement`. Depending on the setting on a second parameter, `learnAsPercentage`, the value on `learningIncrement` may mean either of two things. If `learnAsPercentage` is false, then the learning increment is taken as an absolute increase (or decrease if < 0) to accuracy or effectiveness. That is, if initial accuracy was 0.6 and the learning increment was 0.1, then the resulting accuracy would be 0.7. Accuracy and effectiveness are bounded at 0 and 1, and so if the change would increase (decrease) the value above (below) 1 (0), then the value is set to 1 (0). On the other hand, if `learnAsPercentage` is true, then the learning increment is taken as to represent a proportional shift towards 0 (if $\text{increment} < 0$) or 1 (if $\text{increment} > 0$). So for example, if initial accuracy was 0.6 and the learning increment was 0.1, then in this case the resulting accuracy would be 0.66. This formulation automatically bounds the result to between 0 and 1.

With these parameters, we can simulate any of several different learning scenarios. One would be a situation of a single major learning event, perhaps a policy change, completion of a new training regime, or imposition of a new military doctrine that systematically changes accuracy and effectiveness levels. In this case we would set the model to learn once, at a specified initial learning time, and with the specified magnitude of change (presumably a significant change such as a 0.10 or 0.20 improvement, representing a 10% or 20% gain). Another scenario might be a simulation of ongoing incremental learning, perhaps simulating soldiers developing better counterinsurgency skills as they gain experience. In this case, after specifying an initial learning time, we would specify a large number of occurrences of learning with very small improvements each time. For instance, we might say that learning occurs every 10 ticks at a magnitude of 0.001, suggesting that every 10 ticks accuracy/effectiveness would increase by 1%. Or we might specify something in between these extremes that might represent the gradual rollout of new policies, for instance by specifying 5 learning episodes of 0.05 improvement each, and spaced 100 ticks apart. The combinations of parameters are infinite, but particular sets of values could be envisioned as representing different scenarios.

Initially, and in the simulation results I present below, learning occurs in both parameters – accuracy and effectiveness – simultaneously and in parallel. Thus, a 0.1 learning value will increase both accuracy and effectiveness by this amount. This simplification initially excludes the possibility of illustrating tradeoffs in learning – for instance, a circumstance where a military learns to be more accurate but only at the expense of effectiveness – but it reduces the number of parameters needed to simulate learning, and shrinks the evaluation space considerably (since values fall between -1 and +1 on one scale, rather than two; if we vary the

learning increment by 0.1, then there are only be 21 values in this setup, rather than $21 * 21$ if accuracy and effectiveness learning occurred independently). In addition, I will limit learning to positive changes in accuracy and effectiveness. Certainly it is theoretically possible for “negative learning” to occur, and militaries might implement counter-productive doctrines. For the initial questions I have asked, however, such as whether and how much learning can help, those scenarios do not fit.

In more detail, the sequence of actions in the new simulation (v3.0, with learning enabled but recruitment options turned off) is as follows.

1. Model initialized. Soldiers and civilians are distributed randomly around the hypothetical landscape. Individual civilian agent anger, fear, violence thresholds are generated from distributions. Civilians with appropriate levels of anger, fear are set as “latent” insurgents. Display is updated.
2. If this is a tick in which learning is to occur (either it is the first tick of learning, or the initial learning tick plus the learning interval if repeated learning is occurring), global accuracy and effectiveness parameters for soldiers are updated. These levels are then propagated to all individual soldiers.
3. One latent or active insurgent is selected at random. If within range of a soldier, the insurgent attacks that soldier, and is designated an active insurgent. If multiple soldiers are in range, one is selected at random. Display is updated to show attack and activation of insurgent.
4. The soldier counterattacks.
 - a. With probability= $effectiveness$, the insurgent is captured/killed by the counterattack. In expanded versions of the simulation, a replacement civilian is randomly placed in the space to replace the removed insurgent.
 - b. With probability= $(1-accuracy)$, each nearby civilian (within interaction range) is hurt by the counterattack. If a civilian is hurt, his/her anger and fear increase following formulas above, and using the total number of civilians hurt.
 - c. Display updated as needed to remove attacking insurgent and change civilian display colors.
5. Simulation ends if a) no latent insurgents remain, b) latent insurgents remain, but none are within interaction range of any soldier, or c) simulation hits fixed tick count (set at 2000 or 5000 in various runs).
6. Repeat from step 2.

The model parameters and their default values in the initial version (version 1) of the simulation are listed in Table 1. The additional parameters introduced in the new learning simulation (version 3) are listed in Table 2.

Table 1: Summary of Key Model Parameters, Initial Model

Parameter	Scope	Definition	Range	Default Value
Space Size	Global	Dimensions of map grid	1+	50
Interaction	Global	Size of neighborhood for	1 to size	3

Range		insurgents looking to attack, and for collateral damage range	of map	
Effectiveness	Global / All Soldiers	Probability of removing an insurgent during a soldier's counterattack	0.0 to 1.0	0.8
Accuracy	Global / All Soldiers	1-accuracy equals the probability of injuring any civilian within interaction range during a soldier's counterattack	0.0 to 1.0	0.8
Sensitivity	Global / All Soldiers	Probability of a soldier responding to an attack. Always 1.0 for this simulation (included for later expansion).	1.0	1.0
Anger	Individual / Each Civilian	Level of this civilian's anger at the government	0.0 to 1.0	Normally distributed, mean 0.25, std. dev. 0.125.
Fear	Individual / Each Civilian	Level of this civilian's fear of government response to insurgent act	0.0 to 1.0	Normally distributed, mean 0.5, std. dev. 0.25.
Violence Threshold	Individual / Each Civilian	Level that anger must pass for civilian to be willing to engage in violent insurgency	0.0 to 1.0	Normally distributed, mean 0.5, std. dev. 0.25.
Fear Increment	Global	Amount that fear increases in each civilian injured by a soldier counterattack	0+	0.10
Anger Increment per Injury		* Superseded in recruitment model		

Table 2: Summary of Key Additional Model Parameters, Learning Model

Parameter	Scope	Definition	Range	Default Value
soldierLearningOn	Global	Boolean turning soldier learning on	True/False	True
learningInitialTime	Global	Tick where initial learning occurs	0+	300*
learningIncrement	Global	Amount of learning that will occur each time it occurs; depending on setting on "learnAsPercentage" this is either an absolute value, or a	-1.0 to 1.0	0.10*

		percentage of the distance toward 0 or 1 accuracy/effectiveness		
learnHowManyTimes	Global	How many times learning will occur. For continuous learning, set to a value larger than the stopping time of the simulation (e.g. > 5000)	0+	1
learningInterval	Global	Frequency at which repeated learning will occur, if learnHowManyTimes is > 1. E.g. if set to 1 or 100, learning occurs every tick or every 100 th tick (respectively) after the initial learning tick.	Integer > 0	100
learnAsPercentage	Global	If true, then the learningIncrement is taken to represent the % of distance to 0 (if increment<0) or 1 (if increment>0) that learning moves accuracy and effectiveness. If false, then learningIncrement represents an actual change.	True/False	False

* This is one of the key parameters examined in the initial simulations

Simulations

I now turn to a discussion of the simulation results. Analysis is performed primarily by examining graphs of the relationships between government strategy and insurgent strategy (strategy being the probability of recruiting vs. attacking when given a choice), and between the timing and magnitude of learning.

Dependent Variables

Each run of the simulation produces an entire history of interactions between an insurgency and the government, and continues either until the insurgency is defeated (no latent insurgents remain in the world), or 2000 ticks are reached (loosely indicating a self-sustaining, or at least long-running, insurgency). For this paper, over 4,000,000 simulations were run, consisting of multiple runs given combinations of input parameters selected to cover the parameter space. Since it is difficult to analyze a “history” in some general sense, and since we are ultimately interested in what the simulations tell us about the success of the insurgency under the specified conditions, I focus on 3 indicators of success.

- Length. Insurgents want a long insurgency (that is, an insurgency that remains undefeated for a long time), while governments want to quickly defeat insurgencies.

- Peak number of insurgents. Insurgents want and need to gain support in the population if they hope to defeat an incumbent government. The simulation does not attempt to model the collapse of a government or its replacement by an insurgent-led group. The simulation does track the size of the insurgency, though. Presumably, the larger the insurgency becomes, the more likely it is to be successful. Hence, we can analyze the peak number of (active or latent) insurgents as an indicator of the probability of success of the insurgency. Insurgents want a large peak strength, while governments would prefer to defeat an insurgency before it becomes large, and so government success is indicated by a lower peak insurgent count.
- Speed of growth. Since insurgencies are especially vulnerable early in their lifespan, insurgencies would prefer to gain widespread support quickly rather than slowly, all other things equal. A final measure of the success of an insurgency is how quickly it can gain the support of the population. I will examine the time it takes for the insurgency to gain 25% support from the population in the form of latent insurgents. Insurgents want the shortest possible time to reaching this level, while governments want to delay this time.

In particular with each outcome variable, we want to know how outcomes differ when soldiers learn, vs. when they do not. We would expect that learning affects the results, but exactly how and to what magnitude remains to be seen.

Anticipated Patterns

In this study, I am particularly interested in the timing of learning and its magnitude. It would be sensible to anticipate several possible patterns that we might observe:

- We would anticipate that the earlier learning occurs, the more effect it will have.
- We would expect that the larger the increment of learning, the more effect learning will have.
- To the extent that repeated learning scenarios are likely accompanied by smaller learning increments, we would expect that repeated learning with the same initial learning time would have less dramatic immediate effects than a single larger increment of learning equal to the sum of smaller increments.
- However, the comparison of repeated learning that ultimately produces a larger cumulative change in accuracy and effectiveness to one-shot learning with a smaller increment is less clear.
- We might think that a large enough amount of learning could overcome any initially bad scenario for a military. However it is possible that the degree of needed learning is too high as to be realistic, and this will need to be explored.
-

Actual Patterns

1. Overall Pattern

An analysis of the simulations taken as a whole initially suggests

----- Figures 1, 2, 3 about here -----

The general finding that XXX is actually somewhat surprising at first glance.

2. Interactions: Context

As a first cut at identifying those circumstances where violence by insurgents is optimal, we begin by varying levels of government accuracy and effectiveness individually. XXX

-- Figure 6 about here --

3. Strategic Interactions

The first two dependent variables XXX in Figure 6 demonstrated only interactions with “context.” But XXX

XXX Other combinations of parameters can produce similar relationships, but this appears to be the case only within the range of relatively low soldier accuracy and effectiveness. Within a range where these levels are low, we can produce similar looking relationships even if we tradeoff between parameters, so that (for instance) we could raise the number of recruitable civilians if we lower effectiveness and still see similar strategic interactions. But if accuracy and effectiveness are high, then the simulations are dominated by relatively quick victories by governments, without interaction.

Conclusions

This paper has focused on development of a simulation that allows the exploration of learning in the accuracy and effectiveness of government military forces in counterattacking insurgents.

In addition to determining how the timing and magnitude of learning would interact in changing outcomes, the analysis looked to determine whether learning could always ultimately produce government victory in a counterinsurgent campaign, or whether there were circumstances where even extensive learning could be “too little, too late.”

While it appeared that XYZ, ...

These results give us some interesting insights into real world circumstances of learning in counterinsurgency.

Bibliography

- Bennett, D. Scott. 2008a. "Governments, Civilians, and the Evolution of Insurgency: Modeling the Early Dynamics of Insurgencies." Manuscript.
- Bennett, D. Scott. 2008b. "Recruiting Your Way to Victory: Varying Strategies in Insurgent/Counterinsurgent Warfare." Paper presented at the 2008 Midwest Political Science Association Annual Meeting, Chicago, IL, April 5, 2008.
- United States Department of the Army. 2007. *The U.S. Army/Marine Corps Counterinsurgency Field Manual*. Chicago: University of Chicago.