

Approaching ABM Virtualizations as Complex Games

Ian Lustick, (University of Pennsylvania)
ilustick@sas.upenn.edu

Miguel Garces (Lustick Consulting)

Thomas McCauley (Lustick Consulting)

Patrick O'Mahen (Lustick Consulting)

*Paper prepared for presentation at
The American Political Science Association Annual Meeting,
San Francisco Sept 3-7, 2015*

Abstract: Holland (2000) highlights how the study of board games can yield profound but underexploited insights into complex social phenomena. We build on this insight, treating games as small complex adaptive systems, with emergent properties susceptible to the kinds of analysis appropriate for much larger, highly complex arenas of interaction. Building on this foundation, we use Agent-Based Model (ABM) simulations as “virtual board games” featuring thousands of agents treatable as either players or game elements. Here we do so to study how insurgents use Improvised Explosive Devices (IEDs) and other available modes of attack and evaluate how a player’s counter-strategies can minimize the impact of IEDs. To develop this complex game, we use our Virtual Strategic Analysis and Forecasting Tool (V-SAFT). V-SAFT is a pioneering ABM simulation platform for building theoretically grounded, realistic models that simulate politics and conflict in real-world countries (Reichert, et al., 2014). This paper contributes to theorization of the analytic potential of board games. Also, we advance efforts to bring simulation, gaming, course of action analysis, and training into fruitful contact with one another within the general domain of the study of political violence and counterinsurgency.

Support for the research reported in this paper was received from the Office of Naval Research (Contract # N00014-12-C-0042). Researchers interested in replicating findings should contact the authors for access to the necessary platform, templates, and protocols. Open-source software for constructing agent-based models, as well as published documentation and theoretical footnoting for the models appearing in this paper, is available at www.lustickconsulting.com.

Introduction

We all have experienced the feeling of being carried away by reading an engrossing book, playing an exciting game, or listening to a spellbinding speaker. In each instance we come to exist within a construct of reality produced by others as if it were the world itself. This process entails the operation of a psychological mechanism known as “transportation,” identified by social psychologists as responsible for high-fidelity absorption by subjects of arguments, narrative, or data (Green & Brock, 2000). Of course, as the popularity of Dungeons and Dragons, comic books, and Donald Trump shows, subjects can be “transported” by narratives that have relatively little empirical validity. But to the extent that the book, game, or speaker do correspond to elements and relationships in the real world, transportation can be important, if not critical, to training (from one point of view) and learning (from another). In this paper we report on how a complex, validated, agent-based model can be presented as a serious game capable of transporting users and sharpening insights relevant to improving counterinsurgency policies.

What is a Game?

A game is simulated competition limited by rules accepted by the players. Accordingly, two fundamental components of games are players and rules. Players are decision-makers attempting to achieve an objective or set of objectives as individuals or a team (Abt, 1987). Rules, or “limiting contexts” (Abt, 1987) are the structure of the game that condition or limit player actions (including the actions that “nature” can take in a game in which nature is a player) and/or the outcomes of those actions. Rules can govern a vast set of circumstances regarding a game’s environment. Take the wizarding world’s most popular game, Quidditch, for example. Rules govern the dimensions of the playing field (an open-air stadium with three goal hoops on either end of the pitch), the means of transportation for each player (flying racing brooms), the number of players on each side (seven), the roles that those players have (three chasers, a keeper, two beaters, and a seeker), the rewards for player accomplishments (each goal earns 10 points, capturing the Golden Snitch is 150 points), and forbidden actions (700 documented types of foul) (Rowling & Whisp, 2001). With players assembled and rules agreed upon, the “game” itself is the sequence of moves and countermoves interpretable as a competitive interaction (from zero-sum to pure coordination).

When is an Agent-Based Model a Game?

These concepts of players and rules map neatly on to the social science constructs of actors (agency) and institutions (structure). To complete our framework, we need to add two more concepts, preferences and strategy. Preferences are the desired outcomes of actors; and implicit in the notion of a game is that one of each player’s preferences is to win by playing according to the rules. In some games, like many athletic contests, player preferences are quite straightforward; they want to win the game as an ultimate objective, and have intermediate objectives of trying to score points. However, in many games, a player’s preferences may ultimately not be clear. This is especially true in games with multiple ways to win or multiple objectives within a game. A final concept in a game is strategy (Lustick, 2011). If preferences denote the destinations for the “game state” that players seek, strategies are the legal sequences of actions in response to possible circumstances, available to players as they seek to arrive at

preferred game states. The rules combined with available mechanisms bound the set of strategies or sequences available to players.

Not only do these game elements map easily onto critical social science concepts, they also correspond naturally to the key features of agent-based models. However, that mapping can take multiple forms. This flexibility stems from the recognition that more than one element in an ABM can obtain status as a "player." Only once we define which element is a player, can we use their perspective to define what roles other features in the ABM take on for game purposes.

For example, imagine a detailed ABM of Yemen, a country with complex domestic politics and subject to machinations by numerous foreign powers; most notably Saudi-organized military intervention, Iranian material support for the Houthi faction and U.S. operations targeting Al-Qaeda militants. If we were to think about this "Virtual Yemen" as a game, who would the players be? We can conceive of three possibilities: individual agents, faction commanders, or the operator of the whole simulation.

First consider individual agents as being controlled by potential players fully constrained by the parameters of the model. From the perspective of an individual agent in the model, the virtual Yemen ABM is a multiplayer game featuring numerous tribal, religious, political and terrorist factions all jockeying for power with few constraints from a weak central state. In this game the player has regulated opportunities to alter specified parameters of a single agent, effectively making decisions for that agent. The range of available strategies is determined by the rules of the game (model parameters), available resources (determined in part by the initial game state), and by the changing circumstances in which players find themselves (the game's history). The player may control an elite agent with a particularly strong listening network, therefore allowing decisions to actually impact the rest of the landscape. Other options available to the player might include either violent or non-violent mobilization, or an alteration to which agents it influences. Those decisions in turn may cause changes in the landscape that eventually ripple back. These types of players are analogous to first-person perspective games in which an individual player representing a single person must survive and or prevail in a complicated environment – anything from Grand Theft Auto¹ to America's Army.

Moving up one level, competing players could command the operations and deployment of a political group in Virtual Yemen. Each player would control the strategies, characteristics, and initial pattern of behavior of particular sets of agents, including some features that would be considered immutable parameters (rules) in a game in which the players view the action through the eyes of individual agents. An agreed upon list of parameters would remain unadjustable, thereby constituting the rules of the game. For example, a player could be given the option to increase or decrease the radius surveyed by agents before they update their behavior or make agents under their control more or less sensitive to marginal changes in their environment, rather than what would be a standard type of player decision, such as choosing how to distribute influential agents or how to link agents under the player's control

¹One could imagine that "Grand Theft Auto" with a setting of Aden or Sa'naa would offer an entirely different, yet immensely satisfying list of characters involved in organized crime to interact with than those in Los Angeles, not to mention the broader array of vehicles to steal.

into networks via remote listening rules. Over a batch of runs, the winning player would be determined by which group of agents “prevailed” over others in terms of their aggregate influence, or as measured by any other salient and agreed characteristic of the array at an agreed-upon time step in each model trajectory. For example, instead of having set strategies from engaging with other groups, players might be able to alter the engagement strategies for the factions they respectively control.

Finally, consider the perspective of an operator controlling a stipulated set of adjustable global parameters, in search of a plausible representation of Yemen that achieves or avoids a desired end with a particular probability. This would represent a game between the operator and “nature.” In such a game, the operator “wins” if the parameter settings or characteristics of the initialization produce preferred outcomes and “loses” to nature if the chosen settings fail to achieve desired ends. Alternatively, multiple operators could shift the strategies or deployments of particular agents or classes of agents in the simulation, using a batch of runs to determine which operator prevailed over the other as measured by the salient features in the distribution of outcomes.

It is worth noting that in any of these game modes Virtual Yemen could be tuned to particular important historical moment or contemporary situation in the real Yemen. This feature affords opportunities to gain valuable insights into the many counterfactual versions of Yemen that cannot be directly studied in the real world. As a result, we can study the consequences of particular policy choices or contingent events that are difficult to uncover through traditional modes of analysis.

ABMs and Game Strategy

Thus far we have considered the ABM as a game from the perspective of agents (as individuals or groups) and from the perspective of the entire landscape. The first perspective corresponds to the traditional perspective of “agency” in which individuals within a set of circumstances they cannot control exercise what discretion they do have. Considering the ABM from the perspective of landscape actually shifts to the “structural” level, since operators have the ability to change the circumstances and rules that confront the agents within the array. But as noted earlier, there is a third perspective, often ignored in traditional discussions of “structure and agency.” This is the perspective of individual strategies—whether they are employed by agents engaged in a large-n game with one another, or used by operators as they manipulate parameter settings and initialization scripts while leaving most theoretical beliefs about the world being simulated unadjusted.

A focus on strategies helps address a typical problem in structure/agency type analyses. Structural explanations are often unsatisfyingly underdetermined, leaving us with too little information to connect cause and effect. That is, while outcomes can vary tremendously within one structural context, change in structure could occur without changing an outcome (albeit one would certainly expect the pattern across a distribution of outcomes to be affected by a structural change). In contrast, agency explanations are unsatisfyingly overdetermined. One specific actor and choice is loaded with an enormous and comprehensive causal claims, ruling out the importance of a plethora of other potentially impactful factors. Focus on the strategies used by agents within the context of particular structures

encourages shifts attention away from that which constrains choice (the structure) or from that which chooses (the agent), to that which is chosen – the sequence of moves. (Lustick, 2011).

Considering a game from the point of view of a strategy is the standard technique of evolutionary game theory, including the well-known tournaments among competing strategies for winning iterated prisoner dilemma games pioneered by Robert Axelrod (Axelrod, 1984). This approach opens wonderful new opportunities for study if applied to a range of problems that involve conundrums of “structure and agency.” Evolutionary game theorists move beyond simple studies of which player “wins” in a specific instance of a game. Instead, they discover which sequences of moves are likely to prevail in multiple iterations of the game against others—which ways of playing become regular and prominent and which fade into obscurity. Similarly, in any agent-based model treated as a game, the focus can be on the strategies employed that tend to prevail over multiple iterations of the game rather than on which player is victorious in a particular iteration. We may find, for example, that some strategies prevail at some player skill levels, while other strategies prevail at rates determined by slight changes in rules (the structure) of the game. Consider a game called “counterinsurgency,” which pits a group of state actors against a group or groups of armed actors seeking to overthrow the state. We could analyze this game from the point of view of the “strategy” by asking, for example, which armed tactic or attack vector emerges as most prevalent, given a players with particular kinds of endowments and under specific sets of rules.²

Evaluating the emergence of strategies is extremely difficult with traditional tools of statistical analysis and game theory. Both provide useful insights, but returns are limited. Statistical analysis generally can usefully focus on beginning and end states, but is poorly suited for directly analyzing processes or paths, both of which are central to the study of strategies. Standard (non-evolutionary) game theoretic analysis, examines the logic behind the interplay of two players (with nature often figured as a single player) and helps to identify the circumstances under which multiple equilibria are present. However, game theory is limited in the number of dimensions, players, payoff schedules, and strategies it can integrate before the algebraic requirements associated with closed-formed solutions are overwhelmed by multidimensionality and multi-body problems.

In contrast, ABMs are well suited to investigate the emergence of strategies in a game because they track the actions and results of each agent’s action during each time step of the simulation. This record makes it easy to observe the sequences of moves and divine the emergence of strategies for agents, each of whom can be a potential player. As importantly, we can also track the interaction of strategies among players as they evolve in a dynamic, organic fashion. This level of complexity far exceeds the capacity of a game-theoretic model both in the number of players and strategies available, as well as in the ability of new strategies to emerge from interactions. Because the ABM keeps a detailed record tracking the development of interactions among players, its data also allow us to make causal claims about which strategies emerge and lead to successful results under a given set of circumstances.

² Note that this conception of strategies is still multilayered. Executing a specific football play, like a play-action pass, is a strategy. However, so is using a broader offensive philosophy, like the West Coast Offense.

Since ABMs collect relevant data at each time step, it is also quite easy for the model to aggregate data and keep “score” to measure the final outcome of a given game. The ABM can measure several different types of outcomes as well, which allows any players to pursue different goals within the context of the game. Finally, we can rerun an ABM’s scenario multiple times to demonstrate the robustness of a particular strategy under a wide variety of initial conditions. This ability is useful to help a player calibrate the distribution of the risk and rewards for a particular strategy on a number of dimensions. For example, in the counterinsurgency game, providing large amounts of open aid instead of small amounts of discreet aid may increase the percentage of simulations under which the favored group triumphs, but also may increase the chances of a coordinated violent backlash against the group, leading to its marginalization. Incorporating multiple potential objectives also will show how changes in aid may maximize the chance of one desired outcome (say dominance for a preferred group) at the cost of a less desirable result on the dimension of another outcome (an increase in violence).

The Role of Serious Games

Whatever its substance, a good game transports players to think, act, and feel as if they are not “playing a game,” but living a part of life real enough to matter to them. In this sense, every game is a simulation of life, and every effective simulation induces a sense that engagement with the simulation as if it were the real world. To the degree these simulations are effective real-world proxies they are natural opportunities for training. Whether we consider flight simulators or Turbo Tax’s simulated IRS tax code, it is possible for users (players) to experience the same kind of thrills, satisfactions, disappointments, fears, surprises, and learning that could come about by repeatedly flying real airplanes or submitting actual claim forms to the IRS at a much higher cost. When games are consciously designed as simulations for training, teaching or investigation purposes, they may be deemed “serious games.” (Abt, 1987) (Michael, 2005).

Games may teach, train, and illuminate, but what are we learning? Assuming the game or simulation is based on good science, i.e. that the models involved have been internally verified and externally validated, there still remains the question of how concrete or abstract are the game’s depictions relative to the world of interest to the operator. A flight simulator that presented basic options involving acceleration, velocity, wind speed, and mass might help impart useful general skills to a user. However, without incorporating aircraft-specific details, the simulator would not help the user land a Boeing 787. Becoming an expert chess player might help a commander think strategically about integrating considerations of opponent intentions and perspectives, and these skills might be broadly useful. However, chess-playing per se would not supply any guidance for carrying out a specific mission to storm a beach protected by coral reefs. Accordingly, for a game or simulation to be serious it must not only be like the world, it must be matched in the user’s mind to the correct level of abstraction with respect to the domain of interest. This proper matching is key to avoid errors of misplaced concreteness.

Complexity, Games and ABMs

In this paper we focus on a game involving counterinsurgency and in particular on the problem of explaining the prevalence of Improvised Explosive Devices (IEDs) as an insurgency tactic in relation to

circumstances and counterinsurgent strategies. For this game to be helpful to analysts or policy makers with respect to real counterinsurgencies, it must align the level of the game's engagement with strategies and tactics to user understanding of how abstract and general, or concrete and operationally detailed, the game is relative to a real counterinsurgency. To achieve this alignment, we use the Lustick-Miodownik three-part typology of ABM simulations, sorting them as to the level of abstraction with which they are designed to engage the complexity of the world into abstractions, ensembles, and virtualizations (Lustick & Miodownik, 2009).

Abstractions are the simplest level of the typology. They manipulate only one or a few variables to model highly generic situations. Abstractions allow for the emergence of complex phenomenon, but only describe in broad terms that can't be readily applied to specific situations. Axelrod's model of culture (1997) and Schelling's examination of the dynamics behind segregation (2006 [1978]) are two famous examples of abstractions that involve agents differing on two or three variables that have limited values.

Ensembles are the middle level of complexity. Unlike abstractions, they provide higher levels of detail – in some cases extremely high levels of detail. The key is that they examine a generic question that is applicable to a specific class of cases studies. For example Lustick, Miodownik, and Eidelson (Lustick, et al., 2004) developed a virtual country called "Beita" to study the results of different policies of a central government toward potential secessionist movements in multicultural countries. Beita itself was a detailed virtual country involving thousands of agents with more than a dozen identities and complicated rules for interaction based on constructivist theories of political science. However, Beita was not representative of a specific county, but of a class of countries facing potential secessionist movements from minority linguistic or ethnic groups.

Virtualization models provide the finest level of granularity. The detail not only deploys thousands of agents with dozens of potential identities and interactions, as some ensembles do, but accurately models a specific country or polity at a precise point in real time. The results from a virtualization model will not necessarily give broad insight into a class of problems, but may provide analysts and policymakers with specific information about how an individual country is likely to perform on measures of interest like governmental stability, ethnic violence, or the progression of a civil war. Examples of virtualizations remain rare in political science work, though recent work has showcased this potential in states like Egypt. (Reichert, et al., 2014)

Finally, note how ABMs featuring complex ensembles and virtualizations have space to incorporate outside influences. If these detailed models also possess a high-level of verisimilitude, they may prove relevant for real-life public policy in multiple ways. First, such games allow the introduction of rare events into a model to see how a system reacts to a large exogenous shock such as a natural disaster or mass-casualty terrorist strike. In parallel fashion, analysts could also examine the system's resilience to equally rare, relatively small but highly consequential perturbations. More relevant to our purposes here is the opportunity to introduce levers for comparing the effectiveness of competing policy options. With respect to foreign aid, for example, we could experiment with which factions in a country we give aid to,

how many factions receive aid, the amount of aid, or the type of aid. We could then evaluate the effects of these aid variations on the level of violence in the country, the stability of the government, and which factions gain power.

Allowing analysts to experiment with these sorts of levers on broader ensemble models can serve multiple functions. It can train analysts, develop research that informs us about types of the general types of trade-offs that accompany certain types of policies, and provide a window into the strategies – the sequence of moves – identities take in response to policies. Specific virtualizations will permit analysts to perform course-of-action analyses to supplement their memos and illustrate their recommendations to policymakers.

We could even imagine a training or research simulations in which it is possible for seasoned analysts to develop what they feel to be the best sequence of policy strategies and enter them into a virtual model. The model could then evaluate each of the strategies against a preset goal – say reducing violence. The resulting competition would be an extension of Axelrod's early game theoretic tournaments in which entrants developed strategies to use in a prisoner's dilemma game (Axelrod, 1984).

In short, then, it is rather straightforward to imagine how ABMs integrating generic and some virtualized features could function as serious games.

Conceptualizing the IED in Counterinsurgency Environment

We now move on from outlining the deep structural links between games and agent-based models and their potential of such model-based games to be "serious," to discuss one policy challenge that is deadly serious and can be usefully studied with an ABM and turned into a game: the struggle to counter the use of improvised explosive devices (IEDs). This class of weapon is often used by insurgents in many civil wars and rebellions around the world. Beginning with a general discussion of the IED's role in an insurgency, we then present an ABM of an insurgency that captures the dynamics of IED use. Finally, we discuss how to "gamify" the ABM.

Understanding the roots and processes of insurgency and rebellion is a longstanding research priority with high relevance to international policy. Our investigation here concerns how a counterinsurgency force can affect insurgent use of IEDs. IEDs accounted for about half of American casualties (Krepinevich & Wood, 2007) – 3,000 dead and 36,000 wounded (Zoroya, 2013) – during the U.S. invasions of Iraq and Afghanistan, and are a major tool used by ISIS insurgents fighting the regimes in Syria and Iraq. They also play major roles as weapons of choice in numerous other insurgencies around the world, including those in Thailand, Colombia, and the Philippines as well as the recently ended uprisings in Sri Lanka and Nepal.

The U.S. Department of Defense has spent roughly \$75 billion to counter the threat, primarily in new armored vehicles and technology research (Zoroya, 2013), as well as tactical training. However, most of this large investment and aid to allies has focused on details of thwarting or avoiding the technology of the devices. As a result, most research has ignored broader questions of the factors that underlie the decision to use IEDs and the chain of interactions that leads to their successful deployment.

Political science scholarship on insurgency and violence is also thin on the use of particular weapon types though researchers have suggested a useful framework on which we can build. Wood (2009) introduces the idea of a repertoire of violence, which notes that insurgents and state forces possess a wide range of violence types that they can choose to employ against opponents or civilians. To *methods* of violence, we can add the idea of *targets* of violence. Armed groups may target civilians or combatants, both, or certain subsets of each with finely discriminated violence or indiscriminate violence (Kocher, et al., 2011). Some forms of the repertoire of violence are more suited to precise targeting, while others, such as artillery strikes or area bombing, are indiscriminate. Closely related is the dimension of target discrimination, the *ability* to deliver an attack accurately against a target of choice. This ability is partially conditioned on the method of violence employed. However, it is also dependent on the skill and training of the individual combatant, as the discipline and command and control features of the forces employing the tactic (Weinstein, 2007).

In general, the implications for counterinsurgency of the prevalence of one type of attack vector over others and the counter question of the impact of different counterinsurgency strategies on the prevalence of different attack vectors are natural topics for investigation via an ABM-realized serious game. The IED is one particular method of violence available to the actors in an insurgency—one that has attracted enormous attention, largely because of the high military and civilian casualties inflicted by insurgents using the weapon in Afghanistan, Iraq, and elsewhere. Gill, Horgan and Lovelace offer the following useful definition an IED.

An explosive device is considered an IED when any or all of the following—explosive ingredient, initiation, triggering or detonation mechanism, delivery system—is modified in any respect from its original expressed or intended function. An IED's components may incorporate any or all of military grade munitions, commercial explosives or homemade explosives. The components and device design may vary in sophistication from simple to complex and IEDs can be used by a variety of both state and non-state actors. Non-state actors can include (but not be limited to) terrorists, insurgents, drug traffickers, criminals and nuisance pranksters. (Gill, et al., 2011, p. 742)

Two features of this definition are particularly noteworthy. First, it screens out any conventional weapon that has not been modified or jerry-rigged in some way. For example, a land mine deployed as such would not be an IED, but a rewired artillery shell deployed as a roadside bomb would be an IED. Second, notice how an IED by definition is an adaptive device. This feature makes it an especially dynamic part of any conflict, which suggests the utility of studying it through a research method such as ABM that directly models adaptation and the interaction between forces over time.

Building A Tactical Counterinsurgency Model

The model we have built to stage our counterinsurgency game, and to examine conditions under which IEDs can be expected to be more or less prominent, is based on Lustick Consulting's well-established PS-I agent-based modeling framework. Over the past six years, we have used PS-I successfully to model conditions in several countries of interest. Here, we offer a brief summary of the model's mechanisms. For detailed descriptions of the framework and examples of countries we have modeled, see Reichert et

al. (2014, pp. 18-20), Alcorn and Garces (2012), and Lustick et al. (2012)

PS-I country models are country-specific virtualizations that feature key generic modules applied to every country. The most important of these modules is the Dynamic Political Hierarchy (DPH). The operation of this module is governed by the activated and subscribed identities of agents, the changing relationship of each identity—which is measured by patterns of overlap in the identity subscriptions of agents, to the most influential and prevalent identity in the array—and the actions that any agent can take based on its position in the hierarchy formed and tracked by this module's operation. Below, we explain the operation of this module in detail.

Identities: Based on theories of constructivism (Lustick, 2012), an agent in our models can have (be *subscribed* to) several different identities, though it will only emphasize or advertise one of those at a given time. For example, an agent can be an Arab, a Shia Muslim, a member of a particular tribe, and an army officer. We call the collection of identities to which the agent is subscribed its *repertoire*. The identity shown publicly to other agents and drawn from the agent's repertoire is that agent's *activated* identity. An agent chooses which identity to activate based on what it has done in the past, conformity to what other agents it observes are doing, and other signals exogenous to agent behavior indicating the relative attractiveness of different identities. We adjust neighborhood sight range, elite networks, listening rules, and range and volatility of exogenous perturbations, known as *biases*, based on relevant data on the country being modeled.

We then aggregate the mix of identities into a power structure called the Dominant Political Hierarchy (DPH). When a country is divided into zones of political contestation where the dominant political constellations are radically different across regions—like in a well-developed insurgency—we build the model with multiple DPH. Within each DPH zone, the number of agents activated on and subscribed to each identity is used to identify the *dominant* identity in that zone. This identity is then in turn used to determine the pecking order of the various other identities. Identities whose subscriptions significantly overlap with the dominant identity make up the *incumbent* elements of the ruling coalition. Activated identities with less overlap to either dominant or incumbent are classified as further from the center of power and the loyalties proximity to the ruling group produces. These categories of identities are the *regime*, *system*, and *non-system* levels. As shorthand, think of agents with activated identities in the dominant or incumbent positions (which are the top tiers of the regime level) as members of a governing coalition, those at the broad regime level as the loyal opposition, and those below the regime level as potential insurgents.

We can use the DPH concept to produce a “state zone” and a separate “insurgent zone,” though their boundaries can change in response to agent behavior and patterns of affiliation among agents. In general, insurgents will be on the margins in state-dominant zone, and agents with state-associated activated identities will be on the margins in insurgent-dominant areas.

When an agent is activated on an identity that according to exogenous signals is substantially less attractive than another identity in its repertoire, it is dissatisfied with the behavior it has been constrained to adopt. These “angry” agents may mobilize, and how disruptive this mobilization is for other agents depends on the agent's position in the DPH. Agents closely tied to the current government in the dominant and incumbent zones lobby, while regime-level agents will protest. Outright violence can come from two groups: the dominant identity can strike out at marginal minorities at the system level, while outcasts at the system level can symmetrically attack the dominant identity.

To study the specific phenomenon of counterinsurgencies and IEDs, we take our generic country-level ABM and graft onto it a module that incorporates the basic give-and-take of an insurgency's tactical environment. Using Wood's idea of methods of violence, we can incorporate IEDs into a tactical module as one of many tactics available for use. As IEDs are flexible weapons that come in innumerable variations, we defined tactics quite broadly as a category of methods with similar features, all requiring a common base of required materiel and knowledge to use. First, we discuss the tactics implemented in the model. Next, we outline the features that operationalize real-world distinctions between them. And finally, we disclose how tactics are integrated into the ABM. For a full discussion on the sources on which we base the validity of the conceptualization of tactics in general and IEDs in particular in this model, see Lustick, O'Mahen, Garces and McCauley (2015), but for this paper we will briefly describe the tactics used, their distinctive features as operationalized, and finally how the tactics module is integrated with the base ABM.

Tactics

We model four distinct categories of violence as typical to insurgencies. Note that not all are available to all agents, and that the state and insurgent variations of the same tactic have slightly differing levels of effectiveness, reflecting the higher levels of training and equipment usually available to state forces in conventional tactics. Brief verbal descriptions of each tactic follow below. Table 1 provides a simple matrix showing the base levels of effectiveness of each tactic vis-a-vis all the others. The asymmetric nature of the graph shows the different effectiveness that the same tactic may have when used by state or insurgent forces. Table 2 shows the basic costs and range characteristics of each tactic for participants.³

Engagement: A high risk, low range tactic that has the benefit of being the cheapest available. Even the most cash-strapped organization can field a squad of men with AK-47s, capable of inflicting newsworthy harm upon the enemy. State infantry, while a step up in training and equipment, are still fielded for comparable objectives.

IED: A cheap (being built from whatever is readily available), low-risk (as the expert builders are nowhere near the battlefield), and reasonably effective tactic, particularly against vehicles. They are only available to insurgents.

Mechanized: A tactic that is higher cost than engagement, though with greater range, the mechanized tactic covers armored vehicles, tanks, and artillery. This tactic is generally effective against infantry, but is vulnerable to air power. For insurgents, this is the most expensive tactic available.

Air: In return for a very high cost, air power provides state forces with unparalleled range at minimal risk.

³ We recognize formal sensitivity tests will be likely required in follow-up research to justify our use of these specific values regarding tactics. However, based on our extensive research of IEDs and other insurgent tactics, we believe these are reasonable settings relative to one another and adequate to establish the plausibility of this approach for studying insurgent and counterinsurgent tactics.

Attacker Tactic	Defender Tactic					
	I.ENG	I.MCH	I.IED	S.ENG	S.MCH	S.AIR
I.ENG	0.5	0.4	0.7	0.4	0.3	0.1
I.MCH	0.7	0.5	0.3	0.6	0.3	0.0
I.IED	0.4	0.6	0.1	0.5	0.6	0.1
S.ENG	0.6	0.4	0.7	0.5	0.3	0.2
S.MCH	0.7	0.6	0.3	0.6	0.5	0.3
S.AIR	0.6	0.8	0.2	0.6	0.8	0.6

Key:
I= Insurgent Forces
S= State Forces
ENG= Engagement/Infantry
IED= Improvised Explosive Device
MCH=Mechanized/Armored forces
AIR= Air power

*Note that insurgents and state forces can theoretically attack within their own zones (gray values). Multiple insurgent groups might attack each other, or insurgents or state forces could discipline rogue units. We do not allow intra-DPH zone attacks in this model, however.

Tactic Attributes

Besides effectiveness, we note three other salient features of tactics at our model’s level of analysis: Cost, Range, and Backfire Risk, which are enumerated for each tactic in Table 2. Cost is an obvious metric—airplanes cost more than AK-47s. Range is equally so—airplanes can reach more distant targets during an operation than a squad of infantry. Backfire risk recognizes that failed attacks have negative effects on the attacker – losses of soldiers, experts, and/or materiel that were destroyed in the fight or abandoned in retreat, or more subtly in the exposure of methods and capabilities to the enemy.

However, certain tactics (IEDs, airstrikes) are designed to keep the attacker and key equipment away from the presumably armed target. An effective infantry unit can have its discipline and morale shattered by losses in the field, but a failed missile attack only costs the attacker the price of the ammunition. Therefore, each tactic has a separate risk of backfiring - inflicting penalties on the attacker. On a failed attack, the attacker has a probability (as listed in Table 2) of suffering a backfire, inflicting penalties to their influence and future ability to use tactics. Because even a successful attack can reveal valuable information to the enemy, our model allows a backfire to occur even in the event of a successful attack, albeit with one-third of the probability of a backfire after a failure. Finally, because of the conception of a backfire as the exposure of methods and communication links, it can also penalize agents linked to the attacking agent, though less severely than originator of the backfire.

Module Integration

In addition to their chosen tactic, our agents have three other attributes that further condition the use of tactics: expertise, resources, and flexibility. Expertise measures an agent’s technical ability to use a given tactic effectively through attacking or defending. An agent with high expertise can make even a weak tactic effective, and will be less likely to experiment with new tactics. Agents also have a measure

of resources, which is compared against a tactic's cost to determine what methods of warfare an agent can afford to use and maintain. We use the size of a group as a good proxy for the support base that funds mobilization. Finally, agents also have a tactic flexibility variable that describes how willing an agent is to discard old tactics and try something new.

The results of combat within the tactics module graft onto the ideas of elite networks built into the generic ABM. Successful attacks result not only in a gain of expertise to the attacker, but also a possibility of sharing expertise in that tactic with other agents in the attacking agent's network who share that agent's identity or allowing those agents to switch onto the more effective tactic. This process represents the diffusion of successful tactics across different units. Unsuccessful attacks result in a decrease in influence for the attacking agent. This training/learning process also occurs for defenders. Additionally, the tactics module's distinction between state and insurgent tactics pairs very well with the DPH Zone module's two political regions.

	Tactic	Cost	Range	Backfire Risk
Insurgent	Infantry	0	1	100%
	Mechanized	2.4	2	100
	IED (Insurgent only)	1.1	2	5
State	Infantry	0.5	1	100
	Mechanized	3.2	2	70
	Air (State Only)	5	4	5

In studying tactics, we are less interested in a particular country than the phenomenon of IEDs in general. As a result, we develop a generic model representing features of a number of countries in which a long-running insurgency battles an established counterinsurgency. Therefore, the default landscape we use to model an insurgency has two separate hierarchies, reflecting a strong insurgency that has managed to create and/or dominate the political system of a region within the country. An agent sees the whole map through the lens of the hierarchy it belongs to, so different agents may have very different ideas about which identity is dominant.

In order to initialize our model with real-world characteristics, we have selected a number of countries that are both in the World Values Survey dataset and have experienced a long-running insurgency to use as a foundation for identities model.⁴ For each randomized model run, two of these countries are selected. Three quadrants of the map are seeded according to three districts from the first country, while the insurgent zone is seeded according to a random district from the second. The three-quarters model the larger state DPH zone, with the final quadrant set as the separate insurgent zone. Certain generic identities (state and military) are artificially weakened in the insurgent zone, but still guarantee some overlap of identities between the two zones, along with international identities such as religions,

⁴ For validation, we use *Janes' Terrorism and Insurgency* state and insurgent attack data for Colombia, Egypt, India, Nigeria, Pakistan, the Philippines, and Thailand. Since the most recent World Values Survey was not available for all seven countries, we expanded the list for identity initialization to similar countries including Colombia, Iraq, Mexico, Nigeria, Pakistan, Peru, the Philippines, Tunisia and Turkey.

languages, and political parties.

Our case selection includes countries with a well-established insurgency that is the country's largest source of violence, so we only model violence for which the attacker and victim are in different DPH zones and ignore other violence not associated with the state-insurgent clash.

A Hypothetical Example of the Model in Action

To illustrate how all of these concepts work together, consider the representative process of a single agent attacking. Our hypothetical attacking insurgent has some experience with the IED tactic and is activated on the dominant identity of the insurgent zone. Its actual range is calculated from the size of the agent's group, and the IED tactic's base range.

Any agents that our attacker perceives as being low in the political hierarchy are valid targets - and this likely includes some agents in the nearby state zone. This first set of potential victims is weighted by the agents' expertise in order to introduce goals that are strategic (reduce enemy capacity) and emotional (revenge and reprisal). In our example, the agent targets a state agent with high expertise and the engagement tactic.

Both agents have some expertise, so there is no net effect on the effectiveness of IEDs against engagement and it remains at around 50%. In this case the attack succeeds, which means the target agent loses all of its expertise, as well as significantly reducing the victim's influence with respect to identities. The victim's identity then becomes temporarily toxic and nearby agents may shift off of it.

Because the attack was successful, the attacker gains some expertise. Other nearby agents gain benefits as well, if they share the attacker's network and subscribe to the attacking identity. Of those agents in the network, some share the attacker's tactic and gain some expertise from the event, learning from the attacker any new discoveries of what worked and what didn't. Instead of gaining expertise, agents in the network using a different tactic have the opportunity to change tactics, learning the recently successful one. They are more likely to learn the new tactic if they have high tactic flexibility and low expertise with a different tactic (a sign their current tactic has not been successful). If the attack had failed (i.e. the defender succeeded), then the defender would be the one advertising its tactic, as well as gaining and sharing expertise with its local network in the manner described.

In our hypothetical example, even though the attack succeeded, it revealed key information about the experts, techniques, and network behind the attack. Therefore, the attacker and agents in its network lose a set percentage of their expertise. Backfires can occur after either successful or failed attacks, but a failed attack has a higher likelihood of backfiring that levies a small and temporary influence penalty on the would-be attacker, representing the social and political weakness of a cell that is both exposed and unsuccessful.

Research Design: Turning the Tactics Model into a Research Game

We can now adapt this ABM framework to our counterinsurgency game. Our player is nominally on the side of the state forces attempting to quell an insurgency, and has one move to make. This move has five different levers to pull that implement five different types of counterinsurgency campaign.

We derive four of these strategies from Lalwani's typology of counterinsurgency strategy (Lalwani, 2014), which argues that states employ four types of military responses depending on goals, available resources and the structural characteristics of an insurgency. These four strategies are can be subdivided by two variables: the amount of effort required by the state and the amount of violence the state uses.

Attrition is a high-effort, high-violence strategy that involves deployment of large numbers of military personnel and weapons. The strategy maximizes the application of violence and repression in an attempt to annihilate the insurgency's military capacity to function. This strategy spreads violence liberally in both attacks on insurgent military targets and attacks on the base of material and popular support for the insurgents, with little regard for the consequences of collateral damage.

Population control, a high-effort low-violence strategy, also requires a massive investment in personnel. However, this campaign style focuses on establishing the state's legitimacy in the hearts and minds of people instead of asserting the state's authority through force. State forces do engage in operations against hardened insurgents. However, the primary effort is to protect the community, provide economic and political support to redress grievances and improve development, and recruit support from the community both through intelligence and training of local forces to help provide security.

Enfeeblement, a low-effort high violence strategy, does not attempt to directly defeat an insurgency, but rather to contain it to acceptable levels of violence and in a limited area. This strategy liberally employs violence, but keeps deployments and major confrontations to a minimum in an attempt to limit costs.

Cooptation is the low-effort low (direct) violence strategy that involves containing an insurgency through the support of local elites who will oppose insurgents in return for a suitable support. This outsourcing process limits the ability of the state to assert sovereignty, and the principal-agent problem may limit the effectiveness of indirect rule. Direct violence from the state is low, but violence delivered by proxies may vary.

The fifth option retains the status quo interactions of the model and serves as a baseline counterinsurgency campaign that is both a medium-effort and medium-violence strategy.

Our challenge was to integrate these broader theories on counterinsurgency types by operationalizing them within the tactics model. A decomposition of Lalwani's typology of strategies suggests they are based on six lower-level variables: Military Effort, Air to Engagement ratio, Public Goods, Saturation of Control, Intel-driven Operations, and Cooptation.

Military Effort is a measure of the resources the state puts behind the counterinsurgency campaign – money, available troops and other resources. In our model, we implemented effort as the number of state agents with a tactic they can use to attack. Since the model only tracks violence between the state and insurgency, the agents with tactics are the agents the state has made available for counterinsurgency actions.

A critical counterinsurgency question is the tradeoff between **Air and Engagement** tactics. What kinds of costs can the state incur, lives or money? Air tactics are generally more expensive but engagement tactics are more likely to put soldiers' lives at risk. In the model, Air to Engagement Ratio adjusts the costs of the two tactics relative to one another and as a result the strategic predispositions of the state shift as well.

States often deploy **public goods** and services to bolster legitimacy, and to tie populations into the government system. Particularly in the context of counterinsurgencies, states use public goods as a reward for those groups and identities that show loyalty to the state or opposition to the insurgents. This is implemented as a global positive shift to the bias range of identities that are high in the state DPH zone. Where most identities have a bias that fluctuates randomly between -3 and +3, these 'loyal' identities range from -1 to +5, which likely will encourage their spread in insurgent areas.

Saturation of control is the direct occupation of the insurgent territory, measuring the extent to which state elite structures attempt to dominate the system of the insurgent zone. To implement this feature in our model, we create several new state elite agents and place them within the insurgent zone. As state elites, they still belong to the state DPH zone (despite geographically being in the insurgent zone), have state tactics, are directly linked to the wider state elite network, and are initialized with some expertise. Think of the presence of these groups of state elites as the state building massive military bases to attempt to control insurgent regions.

Intel-driven operations reflect an interest in targeting the insurgent network without massive investment of troops and weapons for direct confrontations. More precisely, the strategy targets the key leadership elements and experts that make the insurgent network powerful. The idea of network vulnerabilities is modeled with backfires, so we operationalized an intel-driven counterinsurgency campaign as one that generates broader insurgency backfires. Instead of just affecting the attacker and its direct connections, this strategy results in backfires that inflict penalties on the attacker and agents within two steps, reflecting more efficient extraction of value from intelligence.

Finally, **co-optation** influences pre-existing elites to work with the state and against the insurgent. To simulate this strategy, we tie some pre-existing insurgent elites into the state network and place the state-friendly identities in their repertoire. Unlike the transplanted elites from saturation of control, these elites act as insurgent zone agents, use insurgent tactics, and have a repertoire of identities corresponding to the insurgent zone.

These variables are used to craft the six strategies, as seen in Table 3.

State Strategy	Modeling Variables					
	Effort	Air vs. Engagement	Public Goods	Saturation of Control	Intel-Driven Operations	Co-optation
Attrition	+	+				
Population Control			+	+		
Enfeeblement	-				+	

Co-optation	-						+
+ represents increased values on variable for counterinsurgency strategy							
- represents decreased values on variable for counterinsurgency strategy							

Enfeeblement and cooptation are Lalwani's low-effort strategies, each implemented as having fewer state agents available to attack, but having an increase in their corresponding variables of intel-driven operations and cooptation. Attrition and population control are both high effort, though for population control this effort is reflected not in the number of attack-capable state agents, but in the effort to create and maintain elites in enemy territory through public goods and saturation of control. To model the default baseline strategy, we set military effort at a level in between attrition on one hand and enfeeblement and cooptation on the other.

The other values that defined the baseline were set based on real-world data on the relative and absolute rates of usage of various tactics, as collected from Jane's Terrorism and Insurgency Database (2015). With the baseline set, and the other tactics defined as specific adjustments to that baseline, we ran 3,000 randomized trials (600 runs per strategy) to evaluate the effects of our player's choices on the course of the insurgency.

Results

The model output from our experiment includes 753,000 rows (3,000 runs with 251 timesteps each) and about 413 variables collected for each timestep, ranging from the activation and DPH level of identities, the tactics available (or used) by agents, and successful defenses or backfires that occur during the model run. Although this data is only a small fraction of the total amount of information that could be collected, it is enough to provide us with general results that provide validity to the model output as well as material for discussion and debate of effective counterinsurgency strategies.

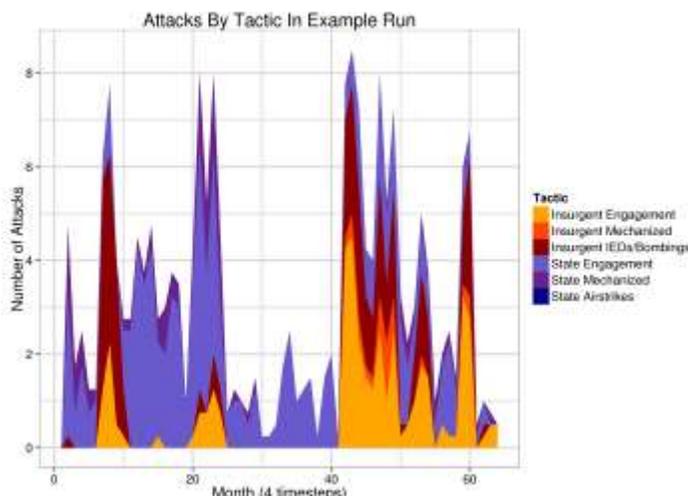


Figure 1 - An example run in our generic tactic model.

There are several ways that we could treat these results as a game, perhaps the most obvious being that we imagine a user experiencing a single run over time, measuring what's happening in the landscape such as the number of attacks per timestep broken down by tactic type. Figure 1 shows an sample run from our experiment where there is an insurgent attack spike early in the run, followed by a large increase in state attacks (quelling the insurgency temporarily), and later another spike in insurgent attacks that does not appear to be dying down by the end of the run. What we cannot see from this figure is that there were two large shifts in the DPH landscape (one before each insurgent attack spike) that lead to both changes in the number of available insurgent attackers or the number of state attackers. These shifts are caused by changes in the favorability of identities, which

then lead to new overlaps between identity-groups, and finally to agents with a higher or lower propensity for violence in the next time step.

This specific example run used the enfeeblement strategy, but how well did the strategy work overall? It's clear that this is only an illustration of what a single simulation run can provide. Players interested in the controlled effect of our four counterinsurgency strategies, would need to dramatically increase sample size. In the next figure we show the effect of our four strategies on the average number of state attacks, insurgent attacks, civilian attacks (as subset of insurgent attacks), and overall attacks.⁵ We were surprised to find that three of our four strategies actually increased the average number of insurgent attacks and only attrition helped to reduce the insurgent violence.⁶

Although the attrition strategy does appear to be most effective if our goal is to reduce insurgent violence, we also see the largest spike in state violence (top left olive bar). The population control strategy, on the other hand, reduces civilian attacks modestly, but increases in state and non-civilian insurgent attacks means that the population control strategy creates the most violence overall (bottom right pink bar). The cooptation and enfeeblement strategies are both low-effort, so it is unsurprising that the number of state attacks decreases, although it is perhaps interesting that the number of insurgent attacks does not decrease in turn.

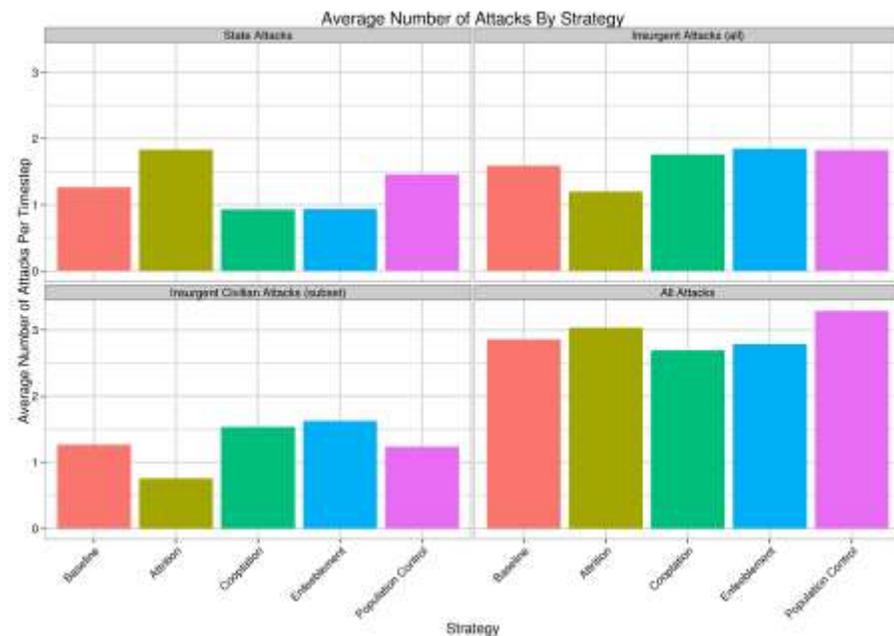


Figure 2 - Average levels of different attack types per time step in our distribution of all runs.

We emphasize that this is merely one specific way to evaluate results. Insurgent violence, while important, is not the only goal by which a player's success could be measured. One objective might be

⁵ This particular model only allows civilian attacks to be carried out by insurgents because we treat all agents in the insurgent zone as potential attackers but only a subset of the state zone. This is clearly an oversimplification but can perhaps highlight the difficulty of measuring levels of civilian violence in complex insurgencies when each side often defines "civilian" differently.

⁶ It is worth reiterating that the environment within which we are testing these strategies is meant to represent a generic well-established conflict without modeling any particular country or insurgency directly. Although some strategies are more or less effective overall, the particular conditions of Thailand or Colombia might see different directional changes in the effectiveness of each strategy. That question could of course be explored using a virtualized models of those countries, but not by using the type of generic model deployed here.

to minimize level of effort while holding insurgent attacks below a certain level. With the development of proper metrics, the same model-based game could be used to study the relative success of each strategy with respect to very different sorts of goals. We could then use these evaluations of effectiveness to analyze the trade-offs between different metrics of success depending on strategies, which would help maximize benefits across a wide variety of metrics while recognizing that all good things often do not go together.

In principle, of course, each of the four strategies is meant to reduce insurgent violence through the use of the particular levers of counterinsurgency strategy. However, the ABM output shows how well-intended strategies may have unintended consequences and provides a clear step-by-step sequence of how actions unfold into consequences. For example, one way to reduce insurgent violence is to target insurgent *expertise*. The enfeeblement strategy takes this tack, but it largely fails. Although we do see a reduction in overall insurgent expertise as intended (not shown) the enfeeblement strategy mostly targets engagement agents because they are more susceptible to backfire. As a result, in our models, enfeeblement effectively destroys the ability of insurgents to wage conventional warfare. However, realizing that conventional warfare is ineffective, insurgents then turn to unconventional means, which *increases* IED use, as shown in Figure 3. One takeaway from our model then may be that a raw focus of developing tactical intelligence may not be an effective way of disrupting insurgent IED use or reducing overall violence.⁷

Another strategy to reduce insurgent violence might be to reduce the *success rate* of insurgent attacks, which occurs in population control and the attrition strategy. Although reducing the success rate also reduces overall attacks under the attrition strategy, the total number of attacks actually increases under population control. This finding means that there is an increase in violent attempts that outpaces the decrease in the success rate, which is another example of how modeling complex processes through an ABM can showcase the interdependence of the many variables that come into play in managing a counterinsurgency.

One final important finding has to do with the spread of IED use. Overall we found that fewer agents tend to have access to the IED tactic, but those that do generally accumulate much higher levels of expertise. The greater cost of implementing the tactic explains the low rate of access, but the tactic pays off in the long run if agents can manage to use it early and build up enough expertise to fend off attacks from the state.⁸ This is perhaps one of the key reasons that our counterinsurgency strategies perform so poorly in our experiment. The cooptation, enfeeblement, and population control strategies all target engagement agents most heavily, which in turn reduces their overall expertise and effectiveness. However, this destruction of the engagement tactic drives those agents to try other tactics like IEDs that are more costly, but also ultimately more effective. In contrast, the attrition strategy actually tends to encourage engagement tactics, therefore increasing the average expertise of insurgent engagement

⁷ In part, of course, this is an artifact of our model. We could model a form of enfeeblement that specifically increases backfire among IED-using agents to reflect a counterinsurgency campaign focused on eradicating IEDs. Creating a slider to allow might provide an idea of the trade-offs between targeting a specific form of attack and generally pushing intelligence.

⁸ This is exactly what happened in Iraq, where bomb making abilities were already well-honed in Saddam Hussein's intelligence services, which seeded the proto-insurgency with high levels of expertise in IEDs (McFate, 2005). The IEDs in Iraq were also much more technically advanced and effective in a wider variety of contexts than ones used in Afghanistan (which, of course, are still a large problem) (Castner, 2012).

agents and causing them to lock themselves into a local maximum, which conventional forces of the state can easily contain. Figure 3 showcases the different levels of violence.

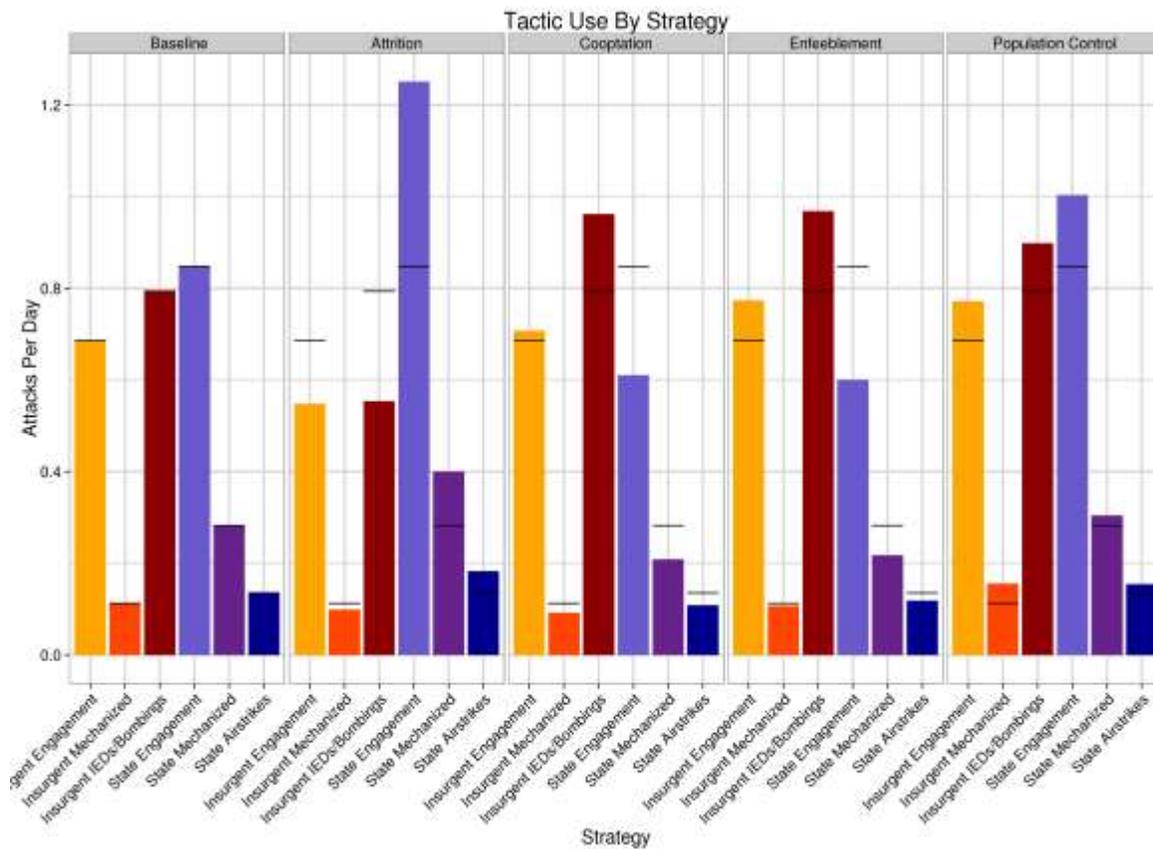


Figure 3 - Average tactic use by strategy in all runs. The black lines represent baseline levels for comparison across strategies.

This is a striking finding. We might ask how it is that the attrition strategy has this indirect but strong effect. We ran further experiments to investigate how the variables within attrition (namely, Military Effort and Air v. Engagement) contributed to the decrease. In doing so we discovered that most of the decrease in violence was a result of the increased Military Effort. More interestingly, we found that while overall insurgent violence decreased, most of the decrease was driven by a decline in IED use. Engagement use did decrease but at a much lower rate. There was a very steep decrease in the total expertise of the IED agents, but we found that the total expertise for engagement agents actually *increased*.

Our analysis of this unexpected result showed no change in number of agents with access to different tactics, the only change was in the expertise that the tactics developed over the course of a run. With a stronger, larger state apparatus, based of course on a high level of effort, the insurgents using IEDs developed less expertise, and so were less successful when attacking. Our hypothesis is that because IEDs tend to cluster and increase expertise at faster rates (due to low backfire likelihood) they attract

more attacks from the state and their relatively high expertise levels protect them - for a time. The state attackers waste many attacks before cutting those clusters down. Meanwhile, engagement has the opportunity to build up steady, broader expertise. Under normal conditions the insurgency might thrive by attracting attention to IED use when there is little payoff to the state's cost of tackling the problem. On the other hand, when state agents mercilessly attack IED users in order to cut down the network, the state eventually succeeds and all that are left are relatively more expert, but isolated and ineffective insurgents specializing in engagement. The major takeaway is that performing effective counterinsurgency operations requires a massive investment of time, resources and soldiers, at least within the parameters of our model.

However, this is only a hypothesis that has grown out of the games initially played with this model. Testing it directly would be challenging, but possible. This perhaps highlights a key criticism of ABM techniques. The complexity of what are in principle fully transparent models can make the exploration of causal mechanisms at work a data- and time-intensive task.

Conclusion

In this paper we have staged a game that can be understood as played by an operator deploying alternative strategies within a structural setting that has important correspondences with actual insurgencies and the actual challenges of counterinsurgency. Alternatively it could be imagined as a game among different insurgent strategies of violence whose outcome varies depending on which counterinsurgency regime those strategies encounter. For the model to be effective as a serious game, its complexity would have to be translated into a user interface capable of allowing the operator, as agent, to "feel" the stakes of changes in parameters or strategic commitments. Even without that transportation effect, however, the model presented here, as a potential game, has yielded findings that are both intriguing and highly suggestive of logics that may well not have been thoroughly explored. IEDs have attracted more attention than other insurgent attack vectors, leading to the commitment of enormous resources to the direct defeat of that tactic. However, a very different but perhaps more promising strategy might be to create conditions that attract insurgents to less effective, easier to defeat, engagement-type attacks. Once those attacks have produced sufficient intelligence, a ramped up attrition strategy could be implemented to defeat the insurgency before it is able to produce difficult to destroy networks of highly effective and adaptive IED-armed insurgents.

Works Cited

- Abt, C. C., 1987. *Serious Games*. Lanham(Maryland): University Press of America.
- Alcorn, B. & Garces, M., 2012. Nesting a data-driven agent-based model within a theoretical framework. *Prcoedia Computer Science*.
- Axelrod, R., 1984. *The Evolution of Cooperation*. New York: Basic Books.
- Axelrod, R., 1997. The Dissemination of Culture: A Model with Local convergence and Global Polarization. *The Journal of Conflict Resolution*, 41(2), pp. 203-226.
- Castner, B., 2012. [Interview] (12 7 2012).
- Gill, P., Horgan, J. & Lovelace, J., 2011. Improvised Explosive Devices: The Problem of Finding a Definition. *Conflict and Terrorism*, 34(9), pp. 732-748.
- Green, M. C. & Brock, T. C., 2000. The Role of Transportation in the Persuasiveness of Public Narratives. *Journal of Personality and Social Psychology*, 79(5), pp. 701-721.
- Holland, J., 2000. *Emergence: From Chaos to Order*. Oxford: Oxford University Press.
- Jane's (subscription required), 2015. *Jane's Terrorism and Insurgency Centre*, s.l.: Jane's.
- Kocher, A., Pepinsky, T. & Kalyvas, S., 2011. Aerial Bombing and Counterinsurgency in the Vietnam War. *American Journal of Political Science*, 55(2), pp. 201-218.
- Krepinevich, A. & Wood, D., 2007. *Of IEDs and MRAPs: Force Protection in Complex Irregular Operations*, Washington DC: Center for Strategic and Budgetary Assessments.
- Lalwani, S., 2014. *Selective Leviathans: Explaining State Strategies of Counterinsurgency and Consolidation*. s.l.:Ph.D dissertation, Massachusetts Institute of Technology.
- Lustick, I., 2012. Deploying Constructivism for the Analysis of Rare Events: How Possible is the Emergence of "Punjabistan"?. In: K. Chandra, ed. *Constructivist Theories of Ethnic Politics*. New York: Oxford University Press, pp. 422-451.
- Lustick, I., Alcorn, B., Garces, M. & Ruvinsky, A., 2012. From Theory to Simulation: the Dynamic Political Hierarchy in Virtualisation Models. *Journal of Experimental & Theoretical Artificial Intelligence*, 24(3), pp. 279-299.
- Lustick, I. & Miodownik, D., 2009. Abstractions, Ensembles and Virtualizations: Simplicity and Complexity in Agent-Based Modeling.. *Comparative Politics*, Issue January, pp. 223-244.
- Lustick, I., Miodownik, D. & Eidelson, R., 2004 . Secessionism in Multicultural States: Does Sharing Power Prevent it or Encourage It?. *American Political Science Review*, 98(2), pp. 209-229.

- Lustick, I., O'Mahen, P., Garces, M. & McCauley, T., 2015. *Agent-Based Simulations as a Modeling Solution for the use of Improvised Explosive Devices in Insurgencies and Rebellions*. Philadelphia, Paper to be presented at the Northeastern Political Science Association Annual Meeting.
- Lustick, I. S., 2011. Taking Evolution Seriously: Historical Institutionalism and Evolutionary Theory. *Polity*, 43(2), pp. 179-209.
- McFate, M., 2005. Iraq: The Social Context of IEDS. *Military Review*, Issue May-June, pp. 37-40.
- Michael, D. R. C. S. L., 2005. *Serious Games: Games that Educate, Train and Inform*. s.l.:Muska&Lipman/Premier Trade.
- Reichert, M., Garces, M. & Lustick, I., 2014. *Studying the Edge of the State Space: An Agent-Based Modeling Approach to Studying Rare Events*. s.l., s.n.
- Rowling, J. K. & Whisp, K., 2001. *Quidditch Through the Ages*. New York: Arthur A. Levine Books.
- Schelling, T. C., 2006 [1978]. *Micromotives and Macrobehavior*. New York: W.W. Norton.
- Weinstein, J., 2007. *Inside Rebelllion: The Politics of Insurgent Violence*. New York: Cambridge University Press.
- Wood, E. J., 2009. Armed Groups and Sexual Violence: When is Wartime Rape Rare?. *Politics and Society*, 37(1), pp. 131-162.
- Zoroya, G., 2013. How the IED Changed the US Military: Homemade bombs caused death and agony for troops and prompted changes in military equipment. *USA Today*, 13 12.